

How The Flang Frontend Works

Introduction to the interior of the Open–Source Fortran frontend for LLVM

Paul Osmialowski*

Arm

Development Solutions Group

Manchester, UK

pawel.osmialowski@arm.com

ABSTRACT

In May 2017, PGI® publicized Flang [16][7], an Open–Source Fortran frontend for LLVM along with a complementary runtime library. The ultimate goal set for Flang is to make it part of the whole LLVM ecosystem with level of support and attention equal to that experienced by the Clang frontend. To come closer to this goal it is important to make Flang widely known and more visible. A good introduction to the frontend interior could serve such a purpose and the intention of this paper is to describe how Flang works and how its source code is structured.

ACM Reference Format:

Paul Osmialowski. 2017. How The Flang Frontend Works. In *Proceedings of LLVM-HPC'17: Fourth Workshop on the LLVM Compiler Infrastructure in HPC, Denver, CO, USA, November 12–17, 2017 (LLVM-HPC'17)*, 14 pages. <https://doi.org/10.1145/3148173.3148183>

1 INTRODUCTION AND MOTIVATION

Through the decades of its existence, the Fortran programming language has proven to be an important language for numerical and scientific computing [2][8]. Among its advantages the most important are:

- A large code base created during decades of investment in scientific software.
- Wide acceptance in scientific environment resulting in a large number of available experts.
- Good expressiveness for describing numerical algorithms.
- Good efficiency of compiled code.
- Portability across different platforms.

These features guarantee further development of the Fortran applications. Utilising the LLVM compiler family in the High Performance Computing (HPC) world demands the existence of a Fortran compiler frontend for it.

*Osmialowski has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement N° 671697.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
LLVM-HPC'17, November 12–17, 2017, Denver, CO, USA

© 2017 Copyright held by the owner/author(s). Publication rights licensed to Association for Computing Machinery.

ACM ISBN 978-1-4503-5565-0/17/11...\$15.00

<https://doi.org/10.1145/3148173.3148183>

2017-10-06 13:08 page 1 (pp. 1-14)

There were a few different attempts to introduce Fortran frontend into the LLVM compilers infrastructure [18][9][4]. The most complete is the Flang project developed by NVIDIA’s PGI® [6][12]. After it was made publicly available [16][7], the next logical step would be to make it a part of the LLVM project on equal rights with the Clang compiler. This appears to be uneasy task since both of these projects were developed apart from each other for most of the time with totally different design goals in mind. The purpose of this article is to reveal the extent of the differences – this could help the wider audience to decide about the possibility of Flang being accepted as a member of LLVM family.

Being a part of the LLVM ecosystem and using its backend, Flang can target multiple hardware architectures. Nowadays support is present for AArch64, 64-bit Power and x86_64 architectures, but it should be fairly easy to add any other 64-bit hardware architecture supported by LLVM to this list.

The Flang project is open for communication with the wider public. The GitHub *flang-compiler* account [12] holds all of the source code and documentation, it is also used for bug tracking – this proved to be a good communication channel with the developers. To improve the communication even further, The *flang-compiler* channel was opened on the Slack on–line service [13]. This is an invitation only service, but anyone can join [14].

Since it was made public, Flang gained lots of attention and became tested for conformance with Fortran standards and performance of generated code. With its source code based on PGI®’s commercial Fortran compiler, Flang inherited most of the interior concepts: two–pass parser, semantic analyzer, two levels of internal representation (ILM, ILI) and internal optimizer. Having the code base developed for more than two decades, such heritage offers maturity and good Fortran standards conformance [1].

Figures 1 and 2 depict how the choice of Fortran compiler (between GFortran and Flang) affects performance of compiled workloads. Note that in case of Flang, most of the optimization is done within LLVM passes and this arrangement (Fortran frontend combined with LLVM middle–end and backend) is still very young, there is room for further improvement.

The rest of the article is organized as follows. Section 2 describes the Flang pipeline operation between its invocation by the frontend driver (Clang) and the LLVM IR code generation. Section 3 focuses more on the source code and its structure. Section 4 provides a *Hello World* example and describes how it is processed. Section 5 guides through the command line parameters accepted while invoking the two stages of the Flang frontend. The article is concluded in Section 6.




Figure 1: Execution time differences between workloads compiled by Flang and workloads compiled by GFortran. Workloads were from the Polyhedron [17] suite of benchmarks (downloaded from <http://www.fortran.uk>) executed on single core of an AArch64 machine equipped with Cortex-A57 cores. Both compilers were invoked with flags `-O3 -ffp-contract=fast -funroll-loops`.

2 THE PIPELINE – HOW FLANG FITS INTO LLVM

The Flang project is split into two major parts. The first part is the frontend driver which comprises of a series of patches held on top of Clang’s git repository fork on GitHub [3]. This part is responsible for parsing Flang and Fortran specific command line options and establishing the whole of the pipeline.

The second part is the frontend itself and its sources are held in a separate git repository on GitHub [11]. When compiled from the sources, it builds two executables (*flang1* and *flang2*), the first reads Fortran code and produces ILM (the first Internal Representation),




Figure 2: Execution time differences between workloads compiled by Flang and workloads compiled by GFortran. Workloads were from the PolyBench/Fortran [15] suite of benchmarks executed on single core of an AArch64 machine equipped with Cortex-A57 cores. Both compilers were invoked with flags `-O3 -ffp-contract=fast -funroll-loops`.

the second reads ILM and produces ILI internally (the second Internal Representation which also undergoes optimization steps) from which LLVM IR is finally generated.

Figure 3 presents a broad view of the Fortran program compilation process. A more detailed list of functional blocks is as follows:

- Flang1 phases:
 - scanner*, turns Fortran code into tokens.
 - parser*, turns tokens into an AST and a symbol table.
 - transformer*, turns the AST into a canonical AST.
 - output*, turns the canonical AST into ILM.
- Flang2 phases:
 - expander*, turns ILM into ILI.
 - optimizer*, turns ILI into optimized ILI.

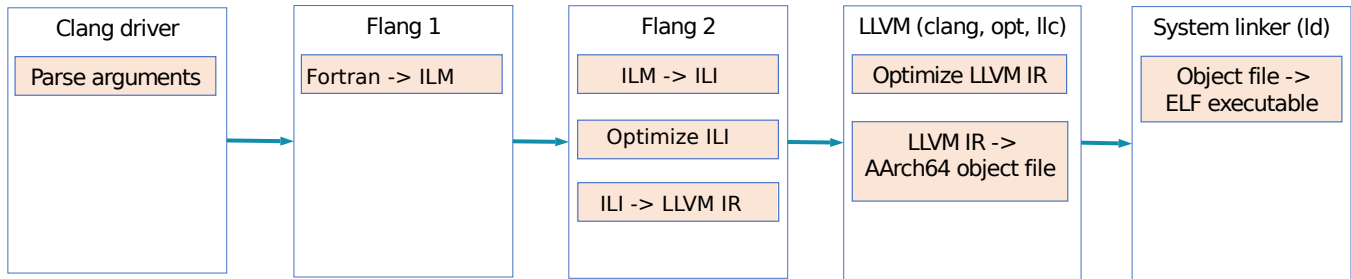


Figure 3: The Flang’s pipeline.

the bridge, turns optimized ILI into LLVM IR.

The Fortran program is read from an input file by the *flang1* executable started by frontend driver. Within *flang1*, the scanner turns Fortran program into tokens, then the parser turns those tokens into an Abstract Syntax Tree (AST) and symbol table (syntab). The AST is then transformed into canonical form which is used to generate temporary a text file with ILM code – the first Intermediate Representation used by Flang.

The ILM code is read from the temporary input text file by the *flang2* executable. During the import stage its text content is fed into data structures held in the process memory. This form is turned into ILI, the second Intermediate Representation used by Flang, which is then optimized by internal optimizer. The optimized ILI is then used during the LLVM bridge phase in which final output text file with LLVM IR is generated.

2.1 Fortran 90 Modules

For each module defined in a given Fortran 90 source file, a *.mod* file is created along with the binary *.o* file produced for a given Fortran source file. If a source file defines more than one module, a single *.o* file is created accompanied by several *.mod* files. The role of Fortran 90 modules is similar to C/C++ header files (except that no preprocessor is involved and they are imported by Fortran 90 programs with a *use* statement) and they are typically installed into *include* directories along with C/C++ header files (while compiled library binaries containing actual executable code are held in *lib* directories). In the case of GFortran, *.mod* files are binaries. In the case of Flang, *.mod* files are text files containing exported symbols definitions (yet still not very human readable) produced by *flang1* (and not processed any further by *flang2*). Modules compiled with one compiler cannot be used by code compiled with the other compiler.

Listing 1 contains example module code while the listing 2 contains generated *.mod* file for this example module code.

```

1 module somemod
2   implicit none
3   real(kind = 8) :: pi
4   save
5   contains
6   integer(kind = 4) function somefun(a, b, c)
7     integer(kind = 4), intent(in) :: a, b, c
8     somefun = a * b * c
9   end function somefun
10 end module somemod
    
```

Listing 1: Example of Fortran 90 module code.

```

1 V30 :0x4 somemod
2 11 somemod.f90 S624 0
3 07/16/2017 11:05:13
4 enduse
5 S 624 24 0 0 0 8 1 0 5015 10005 0 A 0 0 0 0 B 0 0 0 0 0 0
6   0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 somemod
7 S 626 6 4 0 0 9 1 624 5023 4 0 A 0 0 0 0 B 0 0 0 0 0 0
8   0 0 0 0 0 627 0 0 0 0 0 0 0 0 0 0 624 0 0 0 0 pi
9 S 627 11 0 0 0 8 1 624 5026 40800000 805000 A 0 0 0 0 B 0
10  0 0 0 0 8 0 0 626 626 0 0 0 0 0 0 0 0 0 0 0 0 0 0
11  0 0 0 _somemod$2
12 S 628 23 5 0 0 6 632 624 5037 4 0 A 0 0 0 0 B 0 0 0 0 0 0
13  0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 somefun
14 S 629 1 3 1 0 6 1 628 5045 4 3000 A 0 0 0 0 B 0 0 0 0 0 0
15  0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 a
16 S 630 1 3 1 0 6 1 628 5047 4 3000 A 0 0 0 0 B 0 0 0 0 0 0
17  0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 b
18 S 631 1 3 1 0 6 1 628 5049 4 3000 A 0 0 0 0 B 0 0 0 0 0 0
19  0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 c
20 S 632 14 5 0 0 6 1 628 5037 4 400000 A 0 0 0 0 B 0 0 0 0
21  0 0 0 2 3 0 0 633 0 0 0 0 0 0 0 0 0 0 6 0 624 0 0 0 0
22  somefun
23 F 632 3 629 630 631
24 S 633 1 3 0 0 6 1 628 5037 4 1003000 A 0 0 0 0 B 0 0 0 0
25  0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
26  somefun
27 Z
28 Z
    
```

Listing 2: Example of *.mod* file generated by *flang1* for Fortran 90 module code.

3 THE CODE AND SOURCES DIRECTORY TREE

3.1 Directory tree

The directory tree holding Flang’s source code is depicted on Figure 4. Since the building process is organized by CMake [5], every source code directory contains the usual *CMakeLists.txt* file. The out-of-source building process is recommended for projects utilizing CMake infrastructure [10].

The code is mostly written in ANSI C98, some C++ files do exist and they are all compiled with *-std=c++11* compiler flag. There are also Fortran routines (e.g. *ftn_transpose_real*, *ftn_transpose_cmplx*, *vmmul_real*, *vmmul_cmplx*) in the runtime library. Various definition files are provided in *nroff* format (*.n* files) from which certain C headers (*.h* files) are generated. Some of the build-time utilities serve that purpose (*astutil*, *machar*, *fen2rst*, *n2rst*, *ilmtpl*, *ilitp*, *fsymutil*, *fesymutil*, *fsymini*, *fesymini*, *fmachar*). The documentation source files are also provided in *nroff* format (*.n* files).

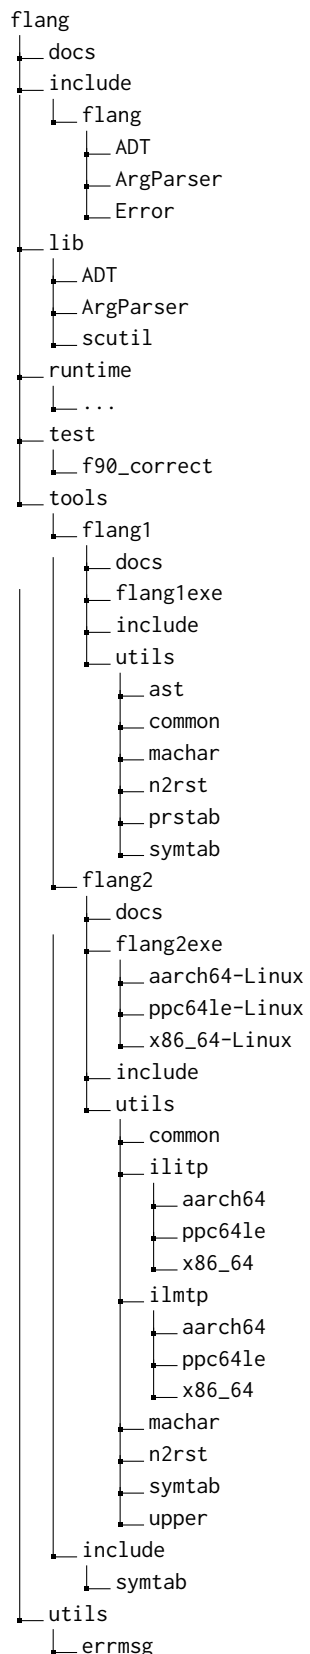


Figure 4: Flang frontend directory structure.

Flang project source code directory structure can be described as follows:

- *lib*, *include* – auxiliary libraries shared by *flang1*, *flang2* and the Fortran runtime library.
- *runtime* – the Fortran runtime library.
- *tools/flang1* – *flang1*, Fortran to ILM translator.
- *tools/flang2* – *flang2*, ILM to LLVM IR translator.
- *tools/include/symtab* – *gbldefs.h* (one of three in the tree) and *global.h* (one of four in the tree).
- *utils* – the *errmsggen* (*errmsg.cpp*) utility, used to generate two *errmsgdf.h* files (for *flang1* and *flang2*) – error message definitions.

The *lib* directory structure presents as follows:

- *ADT* (Abstract Data Types) – contains *hash.c* – open addressing, quadratically probed hash tables.
- *ArgParser* – command line arguments parser routines, common to *flang1* and *flang2*.
- *scutil* – various routines, mostly sophisticated file I/O handling.

The *tools/flang1* directory structure presents as follows:

- *flang1exe* – the frontend, stage 1.
- *include/platform.h.in* – template for *platform.h* header file generated by CMake (platform specific definitions).
- *utils* – utilities required during *flang1* building process:
 - *ast* – *astutil*, generates *ast.h* and *astdf.h* from *ast.in.h*.
 - *common* – *utils.h*, *utils.cpp*, nroff-to-C converter used by *astutil*, *machar* and *fen2rst*.
 - *machar* – *machar.n*, *machar.cpp*, machine characteristics utility definitions, used for generating *machar.h* and *machardf.h*.
 - *n2rst* – *n2rst.cpp*, nroff (.n) to Shpinx (.rst) format conversion utility.
 - *prstab* – parse table generator; grammar definition and parser tables.
 - *symtab* – symbol table generator and utilities.

The *tools/flang2* directory structure presents as follows:

- *flang2exe* – the frontend, stage 2.
- *include/platform.h.in* – template for *platform.h* header file generated by CMake (platform specific definitions).
- *utils* – utilities required during *flang2* building process:
 - *common* – *utils.h*, *utils.cpp*, nroff-to-C converter used by *ilitp*, *ilmtp*, *fsymutil*, *fesymutil*, *fsymini*, *fesymini* and *fmachar*.
 - *ilmtp* – *ilmtp.cpp*, ILM template utility; reads ILM definition file (*ilmtp.n*) and generates *ilmtp.h* and *ilmtpdf.h*; contains architecture specific *.n* files for aarch64, ppc64le and x86_64.
 - *ilitp* – *ilitp.cpp*, ILI template utility; reads ILI definition file (*ilitp.n*) and generates *iliatt.h* and *ilinofof.h*; contains architecture specific *.n* files for aarch64, ppc64le and x86_64.
 - *machar* – see previous.
 - *n2rst* – see previous.
 - *symtab* – see previous.
 - *upper* – *upperl* utility, turns *upperilm.in* into *upperilm.h*; included by *flang2exe/upper.c* which imports the lowered Fortran code.

The directory structure contains some duplication. We can observe some utilities (*machar*, *n2rst*, *syntab*) repeated in both *tools/flang1* and *tools/flang2* directories. It is expected that more general versions of these tools should emerge in time and they will be moved to the top-level *utils* directory (a similar move was already done to the *errmsggen* utility). Detailed lists of files in *tools/flang1/flang1exe* and *tools/flang2/flang2exe* also contain some duplication. From an LLVM developers point of view, the most important difference is the presence of *ll*-prefixed filenames in the *tools/flang2/flang2exe* directory, these files are parts of the LLVM IR bridge code: *ll_builder.c*, *ll_builder.h*, *ll_dbgutil.c*, *ll_ftn.c*, *ll_structure.c*, *ll_structure.h*, *ll_write.c*, *ll_write.h*, *llassem.c*, *llassem.h*, *llassem_common.c*, *lldebug.c*, *lldebug.h*, *llmputil.c*, *llmputil.h*, *llopt.c*, *llsched.c*, *llutil.c* and *llutil.h*.

The LLVM IR generation is performed in terms of explicit operations on text strings. Listing 3 presents an example of code translating an operation identifier into its textual representation.

```

1 static const char *
2 get_op_name (enum LL_Op op)
3 {
4     switch (op) {
5     case LL_FPTRUNC:
6         return "fptrunc";
7     case LL_FPEXT:
8         return "fpext";
9     case LL_SEXT:
10        return "sext";
11    case LL_ZEXT:
12        return "zext";
13    case LL_TRUNC:
14        return "trunc";
15    case LL_BITCAST:
16        return "bitcast";
17    case LL_SITOFP:
18        return "sitofp";
19    case LL_UITOFP:
20        return "uitofp";
21    case LL_FPTOSI:
22        return "fptosi";
23    case LL_FPTOUI:
24        return "fptoui";
25    case LL_ADD:
26        return "add";
27    case LL_FADD:
28        return "fadd";
29    case LL_SUB:
30        return "sub";
31    case LL_FSUB:
32        return "fsub";
33    case LL_MUL:
34        return "mul";
35    case LL_FMUL:
36        return "fmul";
37    case LL_UDIV:
38        return "udiv";
39    case LL_SDIV:
40        return "sdiv";
41    case LL_SREM:
42        return "srem";
43    case LL_UREM:
44        return "urem";
45    case LL_FDIV:
46        return "fdiv";
47    case LL_OR:
48        return "or";
49    case LL_ASHR:
50        return "ashr";
51    case LL_LSHR:
52        return "lshr";
53    case LL_AND:
54        return "and";

```

```

55 case LL_XOR:
56     return "xor";
57 case LL_SHL:
58     return "shl";
59 case LL_INTTOPTR:
60     return "inttoptr";
61 case LL_PTRTOINT:
62     return "ptrtoint";
63 case LL_ICMP:
64     return "icmp";
65 case LL_FCMP:
66     return "fcmp";
67 default:
68     return "thisisnotacceptable";
69 }
70 }

```

Listing 3: Example of code translating operation identifier into its textual representation.

3.2 Documentation

The documentation source files are also provided in *nroff* format (*.n* files). Various command-line utilities (e.g. *groff* or *a2ps*) can be used to generate documentation pages from these source files. The same format is used for various definition files (*.n* files) from which certain C headers (*.h* files) are generated. This ensures that all generated entities are also well documented.

Directory tree contains various *.n* files (*nroff* format) with meaningful names (one chapter - one *.n* file), see Listing 4 for full list of provided documentation source files.

```

1 tools/flang1/docs/ast.n
2 tools/flang1/docs/commopt.n
3 tools/flang1/docs/comms.n
4 tools/flang1/docs/controller.n
5 tools/flang1/docs/dinit.n
6 tools/flang1/docs/grammar.n
7 tools/flang1/docs/intro.n
8 tools/flang1/docs/outconv.n
9 tools/flang1/docs/output.n
10 tools/flang1/docs/parser.n
11 tools/flang1/docs/scanner.n
12 tools/flang1/docs/semant.n
13 tools/flang1/docs/titlep.n
14 tools/flang1/docs/transform.n
15 tools/flang1/utils/ast/ast.n
16 tools/flang1/utils/machar/machar.n
17 tools/flang1/utils/syntab/symini_ftn.n
18 tools/flang1/utils/syntab/syntab.n
19 tools/flang2/docs/coding.n
20 tools/flang2/docs/controller.n
21 tools/flang2/docs/dinit.n
22 tools/flang2/docs/error.n
23 tools/flang2/docs/expander.n
24 tools/flang2/docs/fin.n
25 tools/flang2/docs/ili.n
26 tools/flang2/docs/ilm.n
27 tools/flang2/docs/intro.n
28 tools/flang2/docs/register.n
29 tools/flang2/docs/xflag.n
30 tools/flang2/docs/xref.n
31 tools/flang2/utils/ilirp/aarch64/ilirp.n
32 tools/flang2/utils/ilirp/aarch64/ilirp_longdouble.n
33 tools/flang2/utils/ilirp/ilirp_atomic.n
34 tools/flang2/utils/ilirp/ppc64le/ilirp.n
35 tools/flang2/utils/ilirp/ppc64le/ilirp_float128.n
36 tools/flang2/utils/ilirp/ppc64le/ilirp_longdouble.n
37 tools/flang2/utils/ilirp/x86_64/ilirp.n
38 tools/flang2/utils/ilirp/x86_64/ilirp_longdouble.n
39 tools/flang2/utils/ilirp/aarch64/ilirp.n
40 tools/flang2/utils/ilirp/aarch64/ilirp_atomic.n
41 tools/flang2/utils/ilirp/aarch64/ilirp_longdouble.n

```

```

42 tools/flang2/utils/ilmtpppc64le/ilmtpp.n
43 tools/flang2/utils/ilmtpppc64le/ilmtpp_atomic.n
44 tools/flang2/utils/ilmtpppc64le/ilmtpp_longdouble.n
45 tools/flang2/utils/ilmtppx86_64/ilmtpp.n
46 tools/flang2/utils/ilmtppx86_64/ilmtpp_atomic.n
47 tools/flang2/utils/ilmtppx86_64/ilmtpp_longdouble.n
48 tools/flang2/utils/machar/machar.n
49 tools/flang2/utils/symtab/symmini_ftn.n
50 tools/flang2/utils/symtab/symtab.n

```

Listing 4: List of provided documentation source files.

3.3 Coding convention

As declared in point 1.3 of *tools/flang2/docs/coding.n* document, the established coding convention is as follows:

- ANSI C98 with some K&R code that will eventually be removed.
- Declared proper use of static keyword for limiting visibility of file globals and file functions.
- Naming convention for files: files that define external symbols have *df* prefix, headers that declare external symbols also have *df* prefix, e.g. *ilidfc*, *flgdf.h*.
- Naming convention for code: variable names are lower case, macros and *typedefs* are typically upper case, *enums* begin with a capital letter, then are lower case (no explicit rule stated for function names though!).
- The first non-comment line of each *.c* file is `#include "gbldefs.h"` (note that there are three *gbldefs.h* files in the directory tree!).
- All dynamic storage allocation and freeing is done through the macros *NEW*, *NEED* and *FREE* defined in *gbldefs.h*.

Although the listed items seem to be reasonable, they do not prevent surprising constructs from appearing in the Flang code. One example of those surprises is the re-defined *assert* macro (see Listing 5). It does not only reuse usual *assert* macro name, it also changes its interface by introducing additional parameters.

```

1 /* \brief Assert that cond is true, and emit an internal
2    compiler error
3    * otherwise.
4    *
5    * Note that unlike the C standard <assert.h> assert()
6    * macro, this version is
7    * active in both debug and release builds, and expands
8    * to a statement,
9    * not an expression.
10   */
11 #define assert(cond, txt, val, sev) \
12   if (cond) \
13     ; \
14   else \
15     interr((txt), (val), (sev))

```

Listing 5: Re-defined assert macro.

3.4 Multi-platform compatibility

In order to achieve portability across hardware architectures, the *platform.h* file is generated by the CMake from *platform.h.in* template file (listing 6 presents example *platform.h* file generated by CMake for a GNU/Linux system running on AArch64 hardware).

```

1 #include "define.h"
2
3 #define VHOST "Linux"
4 #define TARGET "aarch64"
5

```

```

6 /* FIXME:
7    * - filtered for pop
8    * - true for all arch/arch
9    */
10 #define BIGOBJ 1
11 #define PTRCOMMON 1
12 #define PGF90 1
13 #define NATIVE 1
14 #define NOLZ 1
15
16 #define HOST_ARM 1
17 #define HOST_LINUX 1
18 #define HOST_LINUX_ARM 1
19 #define HOST_LINUX_ARM64 1
20 #define TARGET_ARM 1
21 #define TARGET_LLVM 1
22 #define TARGET_LLVM_ARM 1
23 #define TARGET_LLVM_ARM64 1
24 #define TARGET_LLVM_64 1
25 #define TARGET_LINUX 1
26 #define TARGET_LINUX_ARM 1
27 #define TARGET_LINUX_ARM64 1

```

Listing 6: Example *platform.h* file generated for an AArch64 system.

It should be noted that despite workable support for three hardware architectures (AArch64, 64-bit Power and x86_64), the original Flang code evolved with strict ties to x86_64 and recently added support for other architectures is implemented with reuse of x86_64 features. Listing 7 shows example of x86_64 definitions reused for AArch64 features. Somewhat detached from the list of supported architectures is the array of known target triples, see Listing 8.

```

1 /* **** ARM -- recycle FEATURE_x64/x86 manifests **** */
2 #if defined(TARGET_LLVM_ARM)
3 #define FEATURE_SCALAR_NEON FEATURE_SCALAR_SSE
4 #define FEATURE_NEON FEATURE_SSE
5 #define FEATURE_FMA FEATURE_FMA3
6 #endif

```

Listing 7: Example of x86_64 features reuse.

```

1 static const struct triple_info known_triples[] = {
2     /* These prefixes are tried in order against
3     target_triple. */
4     {"nvptx64-", "e-p:64:64:64-i1:8:8-i8:8:8-i16:16:16-
5     i32:32:32-i64:64:64"
6     "-f32:32:32-f64:64:64-v16:16:16-v32
7     :32:32-v64:64:64-v128:128:"
8     "128-n16:32:64"},
9     {"spir64-", "e-p:64:64:64-i1:8:8-i8:8:8-i16:16:16-i32
10    :32:32-i64:64:64"
11    "-f32:32:32-f64:64:64-v16:16:16-v24
12    :32:32-v32:32:32-v48:64:64"
13    "-v64:64:64-v96:128:128-v128:128:128-v192
14    :256:256"
15    "-v256:256:256-v512:512:512-v1024
16    :1024:1024"},
17    {"i386", "e-p:32:32-f64:32:64-f80:32-n8:16:32-S128"},
18    {"x86_64-", "e-p:64:64-i64:64-f80:128-n8:16:32:64-
19    S128"},
20    {"armv7-", "e-p:32:32-i64:64-v128:64:128-n32-S64"},
21    {"aarch64-", "e-m:e-i64:64-i128:128-n32:64-S128"},
22    {"powerpc64le", "e-p:64:64-i64:64-n32:64"},
23    {"", ""}
24 };

```

Listing 8: Array of target triples.

4 HELLO WORLD EXAMPLE

In order to demonstrate how example Fortran code is processed by Flang, a short program can be used as seen in Listing 9.

```
1 program hello
2   print *, 'hello world ', 12.34
3 end
```

Listing 9: *hello.f90* Example Fortran code.

Using the `-v` parameter we can instruct the frontend driver to print out all the commands with their arguments that it is about to execute in order to proceed with the compilation (see example in Listing 10). The meanings of particular *flang1* and *flang2* arguments are described in Section 5.

```
1 $ flang -v -c hello.f90
2 clang version 4.0.1 (https://github.com/flang-compiler/
3 clang.git bcdf99e52b47e13a64504a5783ce4eed40833835)
4 (http://llvm.org/git/llvm.git
5 f3d3277bb713bb8aced9a7ac2e9b05c52d2844ee)
6 Target: aarch64-unknown-linux-gnu
7 Thread model: posix
8 Candidate multilib: .;@m64
9 Selected multilib: .;@m64
10
11 "flang1" hello.f90 -opt 0 -terse 1 -inform warn -nohpf -
12 nostatic -y 129 2 -inform warn -x 19 0x400000 -quad
13 -x 59 4 -x 15 2 -x 49 0x400004 -x 51 0x20 -x 57 0x4c
14 -x 58 0x10000 -x 124 0x1000 -tp px -x 57 0xfb0000 -
15 x 58 0x78031040 -x 47 0x08 -x 48 4608 -x 49 0x100 -
16 stdinc $(dirname which flang ) ../include:/usr/
17 local/include:$(dirname which flang ) ../lib/clang
18 /4.0.1/include:/usr/include/aarch64-linux-gnu:/
19 include:/usr/include -def unix -def __unix__ -def
20 __unix__ -def linux -def __linux__ -def __linux__ -def
21 __NO_MATH_INLINES -def __LP64__ -def __x86_64__ -def
22 __x86_64__ -def __LONG_MAX__=9223372036854775807L -
23 def __SIZE_TYPE__=unsigned long int -def
24 __PTRDIFF_TYPE__=long int -def __THROW= -def
25 __extension__= -def __amd64__amd64__ -def __k8__ -def
26 __k8__ -def __PGLLVLM__ -freeform -vect 48 -y 54 1 -
27 x 70 0x40000000 -y 163 0xc0000000 -x 189 0x10 -
28 stbfile /tmp/hello-f2fb4a.stb -modexport /tmp/hello-
29 f2fb4a.cmod -modindex /tmp/hello-f2fb4a.cidx -output
30 /tmp/hello-f2fb4a.ilc
31
32 "flang2" /tmp/hello-f2fb4a.ilc -ieee 1 -x 6 0x100 -x 42 0
33 x400000 -y 129 4 -x 129 0x400 -fn hello.f90 -opt 0 -
34 terse 1 -inform warn -y 129 2 -inform warn -x 51 0
35 x20 -x 119 0xa10000 -x 122 0x40 -x 123 0x1000 -x 127
36 4 -x 127 17 -x 19 0x400000 -x 28 0x40000 -x 120 0
37 x10000000 -x 70 0x8000 -x 122 1 -x 125 0x20000 -quad
38 -x 59 4 -tp px -x 120 0x1000 -x 124 0x1400 -y 15 2
39 -x 57 0x3b0000 -x 58 0x48000000 -x 49 0x100 -astype
40 0 -x 183 4 -x 121 0x800 -x 54 0x10 -x 70 0x40000000
41 -x 249 40 -x 124 1 -y 163 0xc0000000 -x 189 0x10 -y
42 189 0x4000000 -x 183 0x10 -stbfile /tmp/hello-f2fb4a
43 .stb -asm /tmp/hello-f2fb4a.ll
44
45 "clang -4.0" -cc1 -triple aarch64-unknown-linux-gnu -emit-
46 obj -mrelax-all -disable-free -disable-llvm-verifier
47 -discard-value-names -main-file hello.f90 -
48 mrelocation-model static -mthread-model posix -
49 mdisable-fp-elim -fmath-errno -masm-verbose -
50 mconstructor-aliases -fuse-init-array -target-cpu
51 generic -target-feature +neon -target-abi aapcs -v -
52 dwarf-column-info -debugger-tuning=gdb -coverage-
53 notes-file ./hello.gcno -resource-dir $(dirname
54 which flang ) ../lib/clang/4.0.1 -fdebug-compilation-
55 dir . -ferror-limit 19 -fmessage-length 316 -fallow-
56 half-arguments-and-returns -fno-signed-char -fobjc-
57 runtime=gcc -fdiagnostics-show-option -fcolor-
58 diagnostics -o hello.o -x ir /tmp/hello-f2fb4a.ll
```

```
14 clang -cc1 version 4.0.1 based upon LLVM 4.0.1 default
15 target aarch64-unknown-linux-gnu
```

Listing 10: Example execution of Flang invoked with `-v` parameter.

Figure 5 depicts how the most significant line of this example program is processed by Flang; it presents relevant fragments of AST, ILM, ILI and LLVM IR.

Listing 11 shows how this example program is turned into tokens (Section 5 explains how to instruct *flang1* and *flang2* to print human-readable representation of their internal structures).

```
1 tkntyp: PROGRAM tknval: 0 lineno: 1
2 tkntyp: <id name> tknval: 0 (hello) lineno: 1
3 tkntyp: END tknval: 0 lineno: 1
4 tkntyp: PRINT tknval: 0 lineno: 2
5 tkntyp: * tknval: 0 lineno: 2
6 tkntyp: , tknval: 0 lineno: 2
7 tkntyp: <quoted string> tknval: 626 lineno: 2
8 tkntyp: , tknval: 0 lineno: 2
9 tkntyp: <real> tknval: 1095069860 lineno: 2
10 tkntyp: END tknval: 0 lineno: 2
11 tkntyp: <END stmt> tknval: 0 lineno: 3
12 tkntyp: END tknval: 0 lineno: 3
```

Listing 11: Tokens extracted from example Fortran program.

The Fortran *print* instruction from the example program is turned during compilation into the following sequence of runtime library calls:

- Subroutine *f90io_src_info03* – initialize internal globals; starts every use of the *print* instruction.
- Function *f90io_print_init* – initialize output buffer.
- Function *f90io_sc_ch_ldw* – place character string into output buffer.
- Function *f90io_sc_f_ldw* – place floating point number into output buffer.
- Function *f90io_ldw_end* – end adding to output buffer and print its content; ends every use of the *print* instruction.

The sequence of calls above is reflected in the AST, Symbol Table and Intermediate Representation. Listing 12 shows the relevant part of the AST (calls to the *f90io_sc_ch_ldw* and the *f90io_sc_f_ldw* functions) as printed in human-readable form by *flang1*.

```
1 ident      hshlk/std:  0  type:integer  alias:  0
2   callfg:0  opt=(0,0)
3 aprt: 20  sprt:  631 (z_io)
4 assign     hshlk/std:  2  type:integer  opt=(0,0)
5 aprt: 21  dest:  20  src:  18
6
7 constant   hshlk/std:  0  type:integer  opt=(0,0)
8 aprt: 22  sprt:  632 (12)
9
10 constant   hshlk/std:  0  type:character*12  opt=(0,0)
11 aprt: 23  sprt:  633 ("hello world ")
12
13 constant   hshlk/std:  0  type:real  opt=(0,0)
14 aprt: 24  sprt:  634 (-9.2615258789E+02)
15
16 constant   hshlk/std:  0  type:integer  opt=(0,0)
17 aprt: 25  sprt:  637 (14)
18
19 unaryop    hshlk/std:  0  type:integer  alias:  0
20   callfg:0  opt=(0,0)
21 aprt: 26  lop:  25  optype:28
22
23 func-call  hshlk/std:  0  type:alias:  0  callfg:1
24   opt=(0,0)
```

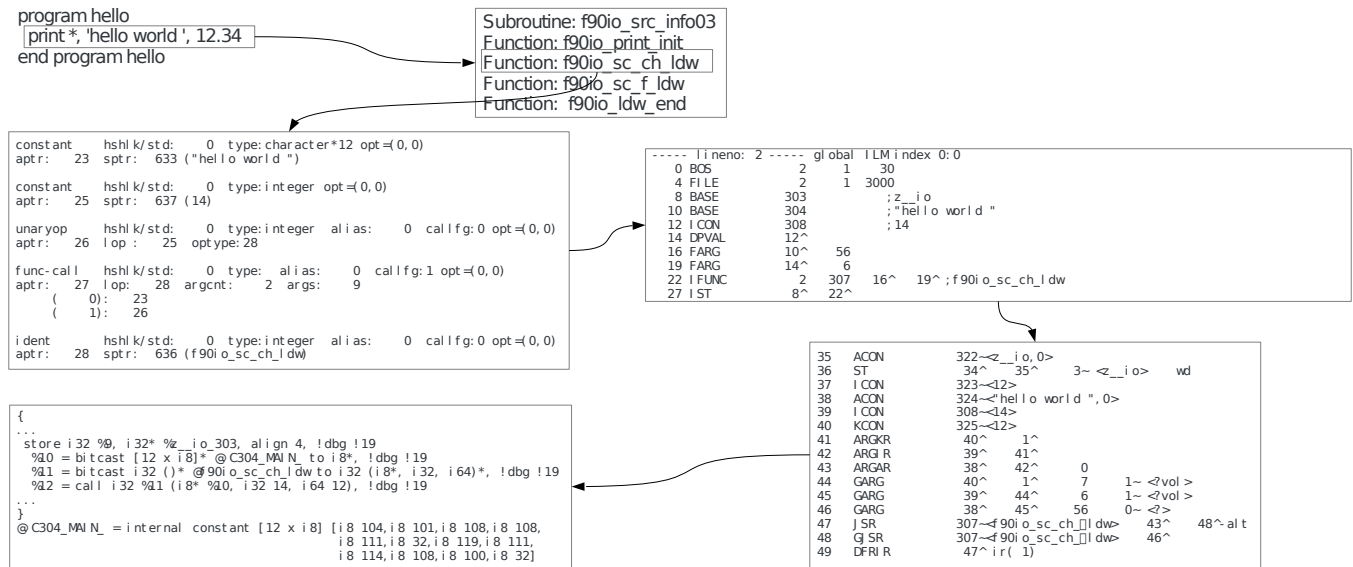


Figure 5: Hello World example.

```
23 aptr: 27 lop: 28 argcnt: 2 args: 9
24 ( 0): 23
25 ( 1): 26
26
27 ident hshlk/std: 0 type:integer alias: 0
callfg:0 opt=(0,0)
28 aptr: 28 sptr: 636 (f90io_sc_ch_ldw)
29
30 assign hshlk/std: 3 type:integer opt=(0,0)
31 aptr: 29 dest: 20 src: 27
32
33 unaryop hshlk/std: 0 type:real alias: 0
callfg:0 opt=(0,0)
34 aptr: 30 lop: 24 optype:28
35
36 constant hshlk/std: 0 type:integer opt=(0,0)
37 aptr: 31 sptr: 639 (27)
38
39 unaryop hshlk/std: 0 type:integer alias: 0
callfg:0 opt=(0,0)
40 aptr: 32 lop: 31 optype:28
41
42 func-call hshlk/std: 0 type:alias: 0 callfg:1
opt=(0,0)
43 aptr: 33 lop: 34 argcnt: 2 args: 12
44 ( 0): 30
45 ( 1): 32
46
47 ident hshlk/std: 0 type:integer alias: 0
callfg:0 opt=(0,0)
48 aptr: 34 sptr: 638 (f90io_sc_f_ldw)
```

Listing 12: Relevant part of Abstract Syntax Tree of example Fortran program.

The *sptr* values are references to Symbol Table, the relevant parts of which are presented in human-readable form in Listing 13.

```
1 z__io integer
2 sptr:631 dtype:6 sc:1=LOCAL stype:6=variable
3 enclfunc:0 hashlk:0 scope:624=hello symlk:1=NOSYM
4 address:0
5 autobj:0 cmlbk:0 cvlen:0 adjstrlk:0 descr:0 midnum:0
newarg:0 devcopy:0 nmcnst:0 ptroff:0 sdsc:0 slnk:0
6 assn dcl d hccsym seq
7 *** b4: 15 0xf
```

```
8
9 12 integer
10 sptr:632 dtype:6 sc:0=NONE stype:3=constant
11 enclfunc:0 hashlk:0 scope:1=NOSYM symlk:0
12 address:0 conval1g:0 conval2g:12 conval3g:0 conval4g:0
13
14 "hello world " character*12
15 sptr:633 dtype:58 sc:0=NONE stype:3=constant
16 enclfunc:0 hashlk:0 scope:1=NOSYM symlk:0
17 address:0 conval1g:5071 conval2g:0 conval3g:0 conval4g:0
18
19 -9.7322831928E+18 real
20 sptr:634 dtype:8 sc:0=NONE stype:3=constant
21 enclfunc:0 hashlk:0 scope:1=NOSYM symlk:0
22 address:0 conval1g:0 conval2g:1095069860 conval3g:0
conval4g:0
23
24 f90io_ldw integer
25 sptr:635 dtype:6 sc:5=EXTERN stype:14=subroutine
26 enclfunc:0 hashlk:0 scope:0 symlk:1=NOSYM
27 dpdsc:0 intent:in paramct:0
28 ggame:0 fval:0 altname:0 slnk:0
29 dcl d func hccsym indep nodesc sdscsafe sequent typd
30
31 f90io_sc_ch_ldw integer
32 sptr:636 dtype:6 sc:5=EXTERN stype:14=subroutine
33 enclfunc:0 hashlk:0 scope:0 symlk:1=NOSYM
34 dpdsc:0 paramct:0
35 ggame:0 fval:0 altname:0 slnk:0
36 dcl d func hccsym nodesc pure
37
38 14 integer
39 sptr:637 dtype:6 sc:0=NONE stype:3=constant
40 enclfunc:0 hashlk:0 scope:1=NOSYM symlk:0
41 address:0 conval1g:0 conval2g:14 conval3g:0 conval4g:0
42
43 f90io_sc_f_ldw integer
44 sptr:638 dtype:6 sc:5=EXTERN stype:14=subroutine
45 enclfunc:0 hashlk:636=f90io_sc_ch_ldw scope:0 symlk:1=
NOSYM
46 dpdsc:0 paramct:0
47 ggame:0 fval:0 altname:0 slnk:0
48 dcl d func hccsym nodesc pure
49
50 27 integer
51 sptr:639 dtype:6 sc:0=NONE stype:3=constant
```



```
52 enclfunc:0 hashlk:370=ibset scope:1=NOSYM symlk:0
53 address:0 conval1g:0 conval2g:27 conval3g:0 conval4g:0
```

Listing 13: Relevant part of Symbol Table of example Fortran program.

As can be seen, Symbol Table contains character strings, numeric constants and runtime library function names.

Listing 14 shows the part of ILM (as generated by *flang1* for the *print* instruction in the example Fortran program) responsible for emitting the "hello world" string into a print buffer. In this example, runtime function *f90io_sc_ch_ldw* (*s636*) is called by the *IUFUNC* instruction. We can observe that the string "hello world" (*s633*) and data type 14 (*s637*; *__STR*, Fortran character, see definition in Listing 15) should be passed to this function.

```
1 i40: -----
2 i0: BOS l2 n1 n0
3 i4: FILE n2 n1 n3000
4 i8: BASE s631
5 i10: BASE s633
6 i12: ICON s637
7 i14: DPVAL i12
8 i16: FARG i10 t58
9 i19: FARG i14 t6
10 i22: IUFUNC n2 s636 i16 t58 i19 t6
11 i27: IST i8 i22
12 i30: -----
```

Listing 14: Part of ILM emitting "hello word" into a print buffer as generated for example Fortran program.

```
1 typedef enum {
2 ...
3   __STR = 14,          /* F character */
4 ...
5   __REAL4 = 27,       /* F real*4, real */
6 ...
7 } _DIST_TYPE;
```

Listing 15: Definitions of Fortran character and Fortran real*4 data types in the runtime library.

Similarly, listing 16 shows the part of ILM (as generated by *flang1* for the *print* instruction of example Fortran program) responsible for emitting the floating point numeric constant into a print buffer. In this example, the runtime function *f90_io_sc_f_ldw* (*s638*) is called by the *IUFUNC* instruction. We can observe that the encoded numeric constant (*s634*) and data type 27 (*s639*; *__REAL4*, Fortran real*4, see definition in Listing 15) should be passed to this function.

```
1 i30: -----
2 i0: BOS l2 n1 n0
3 i4: FILE n2 n1 n4000
4 i8: BASE s631
5 i10: RCON s634
6 i12: DPVAL i10
7 i14: ICON s639
8 i16: DPVAL i14
9 i18: FARG i12 t8
10 i21: FARG i16 t6
11 i24: IUFUNC n2 s638 i18 t8 i21 t6
12 i29: IST i8 i24
13 i32: -----
```

Listing 16: Part of ILM emitting the floating point numeric constant into a print buffer as generated for example Fortran program.

Listing 17 shows the part of ILI (as generated by the *flang2* for *print* instruction of example Fortran program) responsible for emitting "hello world" and the floating point numeric constant. Two calls to runtime functions (*f90io_sc_ch_ldw* and *f90io_sc_f_ldw*) can be easily noticed. The code preparing arguments for those functions is also similar. Note that the Symbol Table identifiers which previously were integer numbers around 600 now are integer numbers around 300 (the Symbol Table was regenerated during import step of *flang2*). This listing presents a human-readable dump of ILI code (see Section 5 for a guidance on dumping internal data structures in a human-readable form).

```
1 35 ACON          322~<z__io,0>
2 36 ST           34^  35^      3~ <z__io>   wd
3 37 ICON        323~<12>
4 38 ACON        324~<"hello world ",0>
5 39 ICON        308~<14>
6 40 KCON        325~<12>
7 41 ARGKR       40^  1^
8 42 ARGIR       39^  41^
9 43 ARGAR       38^  42^      0
10 44 GARG        40^  1^      7      1~ <?vol>
11 45 GARG        39^  44^      6      1~ <?vol>
12 46 GARG        38^  45^      56     0~ <?>
13 47 JSR         307~<f90io_sc_ch_ldw>  43^  48^~
    alt
14 48 GJSR        307~<f90io_sc_ch_ldw>  46^
15 49 DFRIR      47^  ir( 1)
16 50 ST           49^  35^      3~ <z__io>   wd
17 51 FCON        305~< 1.2340000153E+01>
18 52 ICON        310~<27>
19 53 ARGIR       52^  1^
20 54 ARGSP       51^  53^
21 55 GARG        52^  1^      6      1~ <?vol>
22 56 GARG        51^  55^      8      1~ <?vol>
23 57 JSR         309~<f90io_sc_f_ldw>  54^  58^~alt
24 58 GJSR        309~<f90io_sc_f_ldw>  56^
25 59 DFRIR      57^  ir( 1)
```

Listing 17: Part of ILI emitting "hello word" and a floating point numeric constant into a print buffer as generated for the example Fortran program.

Finally, Listing 18 shows a part of LLVM IR (without metadata) as generated by *flang2* for the example Fortran program.

```
1 ; ModuleID = 'hello.f90'
2 target datalayout = "e-m:e-i64:64-i128:128-n32:64-S128"
3 target triple = "aarch64-unknown-linux-gnu"
4 define void @MAIN() !dbg !14 {
5   L.entry:
6     %z__io_303 = alloca i32, align 4
7
8     %0 = bitcast i32* @.C283_MAIN_ to i8*, !dbg !18
9     %1 = bitcast void (*) @fort_init to void (i8)*, !dbg
    !18
10    call void %1 (i8* %0), !dbg !18
11    br label %L.LB1_316
12
13  L.LB1_316:
14    %2 = bitcast i32* @.C300_MAIN_ to i8*, !dbg !19
15    %3 = bitcast [9 x i8]* @.C298_MAIN_ to i8*, !dbg !19
16    %4 = bitcast void (*) @f90io_src_info03 to void (i8*,
    i8*, i64)*, !dbg !19
17    call void %4 (i8* %2, i8* %3, i64 9), !dbg !19
18    %5 = bitcast i32* @.C301_MAIN_ to i8*, !dbg !19
19    %6 = bitcast i32* @.C283_MAIN_ to i8*, !dbg !19
20    %7 = bitcast i32* @.C283_MAIN_ to i8*, !dbg !19
21    %8 = bitcast i32 (*) @f90io_print_init to i32 (i8*, i8*,
    i8*, i8)*, !dbg !19
22    %9 = call i32 %8 (i8* %5, i8* null, i8* %6, i8* %7), !
    dbg !19
23    store i32 %9, i32* %z__io_303, align 4, !dbg !19
```

```

24 %10 = bitcast [12 x i8]* @.C304_MAIN_to i8*, !dbg !19
25 %11 = bitcast i32 (*)* @f90io_sc_ch_ldw to i32 (i8*, i32
    , i64)*, !dbg !19
26 %12 = call i32 @%11 (i8* %10, i32 14, i64 12), !dbg !19
27 store i32 %12, i32* %z__io_303, align 4, !dbg !19
28 %13 = bitcast i32 (*)* @f90io_sc_f_ldw to i32 (float ,
    i32)*, !dbg !19
29 %14 = call i32 @%13 (float 0x4028AE1480000000, i32 27),
    !dbg !19
30 store i32 %14, i32* %z__io_303, align 4, !dbg !19
31 %15 = call i32 @f90io_ldw_end (), !dbg !19
32 store i32 %15, i32* %z__io_303, align 4, !dbg !19
33 ret void, !dbg !20
34 }
35
36 @.C310_MAIN_ = internal constant i32 27
37 @.C305_MAIN_ = internal constant float 0x4028AE1480000000
38 @.C308_MAIN_ = internal constant i32 14
39 @.C304_MAIN_ = internal constant [12 x i8] [i8 104,i8
    101,i8 108,i8 108,i8 111,i8 32,i8 119,i8 111,i8 114,
    i8 108,i8 100,i8 32]
40 @.C284_MAIN_ = internal constant i64 0
41 @.C301_MAIN_ = internal constant i32 6
42 @.C298_MAIN_ = internal constant [9 x i8] [i8 104,i8 101,
    i8 108,i8 108,i8 111,i8 46,i8 102,i8 57,i8 48]
43 @.C300_MAIN_ = internal constant i32 2
44 @.C283_MAIN_ = internal constant i32 0

```

Listing 18: Part of LLVM IR (without metadata) as generated by *flang2* for the example Fortran program.

As we can observe, the global scope of a Fortran program becomes wrapped by a compiler-generated *MAIN_* function.

5 FRONTEND COMMAND LINE OPTIONS (FLANG1 AND FLANG2)

Although both *flang1* and *flang2* executables can be invoked separately without being controlled by frontend driver, no help is provided describing command line arguments accepted by either of them. The best way to find the full list of accepted arguments is to examine the source code.

Listing 19 shows what arguments are registered for *flang1* using the *ArgParser* library API.

```

1  arg_parser_t *arg_parser;
2
3  create_arg_parser(&arg_parser, true);
4
5  /* Register two ways for supplying source file argument
    */
6  register_filename_arg(arg_parser, &sourcefile);
7  register_string_arg(arg_parser, "src", &sourcefile, NULL
    );
8  /* Output file (.ilm) */
9  register_combined_bool_string_arg(arg_parser, "output",
    (bool *)&(flg.output),
10     &outfile_name);
11 /* Other files to input or output */
12 register_string_arg(arg_parser, "stbfile", &stboutfile,
    NULL);
13 register_string_arg(arg_parser, "modexport", &
    modexport_val, NULL);
14 register_string_arg(arg_parser, "modindex", &
    modindex_val, NULL);
15 register_string_arg(arg_parser, "qfile", &dbgfile, NULL)
    ;
16
17 /* Optimization level */
18 register_integer_arg(arg_parser, "opt", &(flg.opt), 1);
19
20 /* Allocate space for command line macro definitions */
21 flg.def = (char **)getitem(8, argc * sizeof(char *));
22 flg.undef = (char **)getitem(8, argc * sizeof(char *));

```

```

23 flg.idir = (char **)getitem(8, argc * sizeof(char *));
24 module_dirs = (char **)getitem(8, argc * sizeof(char *))
    ;
25 /* Command line macro definitions */
26 register_string_list_arg(arg_parser, "def", flg.def);
27 register_string_list_arg(arg_parser, "undef", flg.undef)
    ;
28 register_string_list_arg(arg_parser, "idir", flg.idir);
29 register_string_list_arg(arg_parser, "moddir",
    module_dirs);
30
31 /* x flags */
32 register_xflag_arg(arg_parser, "x", flg.x,
    (sizeof(flg.x) / sizeof(flg.x[0])));
33
34 register_yflag_arg(arg_parser, "y", flg.y);
35 /* Debug flags */
36 register_qflag_arg(arg_parser, "q", flg.dbg,
    (sizeof(flg.dbg) / sizeof(flg.dbg[0])));
37
38 register_action_map_arg(arg_parser, "qq", phase_dump_map
    , dump_map);
39
40 /* Other flags */
41 register_boolean_arg(arg_parser, "mp", (bool *)&(flg.smp
    ), false);
42 register_boolean_arg(arg_parser, "preprocess", &
    arg_preproc, true);
43 register_boolean_arg(arg_parser, "reentrant", &
    arg_reentrant, false);
44 register_integer_arg(arg_parser, "terse", &(flg.terse),
    1);
45 register_inform_level_arg(arg_parser, "inform",
    (inform_level_t *)&(flg.inform
    ), LV_Inform);
46 register_boolean_arg(arg_parser, "hpf", (bool *)&(flg.
    hpf), true);
47 register_boolean_arg(arg_parser, "static", (bool *)&(flg
    .doprelink), true);
48 register_boolean_arg(arg_parser, "quad", (bool *)&(flg.
    quad), true);
49 register_boolean_arg(arg_parser, "freeform", &
    arg_freeform, false);
50 register_string_arg(arg_parser, "tp", &tp, NULL);
51 register_string_arg(arg_parser, "stvinc", &(flg.stvinc),
    (char *)1);
52 register_integer_arg(arg_parser, "vect", &(vect_val), 0)
    ;
53 register_boolean_arg(arg_parser, "standard", (bool *)&(
    flg.standard), false);
54 register_boolean_arg(arg_parser, "save", (bool *)&(flg.
    save), false);
55 register_boolean_arg(arg_parser, "extend", &arg_extend,
    false);
56 register_boolean_arg(arg_parser, "recursive", (bool *)&(
    flg.recursive),
57     false);
58 register_string_arg(arg_parser, "cmdline", &(flg.cmdline
    ), NULL);
59
60 /* Set values form command line arguments */
61 parse_arguments(arg_parser, argc, argv);
62

```

Listing 19: Registration of *flang1* command line arguments.

Closer examination of *flang1* source code reveals the meaning of these arguments:

- *-filename* – source file name (method 1 of 2).
- *-src* – source file name (method 2 of 2).
- *-output* – output *.ilm* text file name (will contain STB + ILM).
- *-stbfile* – output *.stb* symbol table text file name (will contain STB only).
- *-modexport* – output *.cmod* text file name, collective content of all modules exported by given Fortran source file (note that separate *.mod* files will be also generated by *flang1!*).

- `-modindex` – output `.cmdx` text file containing index of exported modules (including offsets in `.cmod` file where each of the modules begins).
- `-qfile` – output debug `.qdbg` text file name.
- `-opt` – optimization level.
- `-def` – same as clang's `-D`.
- `-undef` – same as clang's `-U`.
- `-idir` – same as clang's `-I`.
- `-moddir` – same as flang's `-J`.
- `-x` – the `-x` parameter (see below).
- `-y` – the `-y` parameter (see below).
- `-q` – the `debug` parameter (see further below).
- `-qq` – the `dump` parameter (see further below).
- `-mp/-nomp` – process OpenMP directives.
- `-preprocess/-nopreprocess` – process C preprocessor directives in Fortran code.
- `-reentrant/-noreentrant` – local variables should have *RECURSIVE* semantics instead of *SAVE* semantics (same as *recursive/norecursive* below); also inhibit special treatment of terminus functions.
- `-terse` – reduce verbosity of error messages (where 0 = do not reduce).
- `-inform` – increase verbosity level of error messages.
- `-hpf/-nohpf` – handle HPF (High Performance Fortran) – off by default, does not seem to be fully implemented!
- `-static/-nostatic` – enable IPA (inter-procedural analysis), does not seem to do anything, always set to *OFF*.
- `-quad/-noquad` – quad align (round-up size) *unconstrained objects* if size is greater or equal to 16.
- `-freeform/-nofreeform` – allow free form code.
- `-tp` – x86-specific target architecture name, needs to be dropped.
- `-stdinc` – colon-separated list of include directories.
- `-vect` – vectorizer settings (bitmask), does not seem to be examined anywhere in the frontend.
- `-standard/-nostandard` – be strict about standard, e.g. disallow use of `<>` instead of `/=` or `.ne.`
- `-save/-nosave` – local variables should have *SAVE* semantics instead of *RECURSIVE* semantics (typically used for overriding OpenMP's default using `-Msave`).
- `-extend/-noextend` – allow lines up to 132 characters length (instead of 72 characters).
- `-recursive/-norecursive` – local variables should have *RECURSIVE* semantics instead of *SAVE* semantics (OpenMP's default, can be forced by `-Mrecursive`).
- `-cmdline` – override command line used to invoke the compiler.

Listing 20 shows what arguments are registered for `flang2` using the `ArgParser` library API.

```

1  /* Argument parser */
2  arg_parser_t *arg_parser;
3
4  /* Create argument parser, force errors on unknown flags
5   (last argument is
6   * true) */
7  create_arg_parser(&arg_parser, true);
8
9  /* Input file name */
10 register_filename_arg(arg_parser, &sourcefile);

```

2017-10-06 13:08 page 11 (pp. 1-14)

```

10 /* Fortran source file */
11 register_string_arg(arg_parser, "fn", &(gbl.file_name),
12 NULL);
13 /* Other files to input or output */
14 register_string_arg(arg_parser, "stbfile", &stboutfile,
15 NULL);
16 register_combined_bool_string_arg(arg_parser, "asm", (
17 bool *)&(flg.asmcode),
18 &asmfile);
19
20 /* Register version arguments */
21 register_string_arg(arg_parser, "vh", &(version.host), "
22 ");
23
24 /* x flags */
25 register_xflag_arg(arg_parser, "x", flg.x,
26 (sizeof(flg.x) / sizeof(flg.x[0])));
27 register_yflag_arg(arg_parser, "y", flg.y);
28 /* Debug flags */
29 register_qflag_arg(arg_parser, "q", flg.dbg,
30 (sizeof(flg.dbg) / sizeof(flg.dbg[0])));
31 register_action_map_arg(arg_parser, "qq", phase_dump_map,
32 dump_map);
33
34 /* Other flags */
35 register_integer_arg(arg_parser, "opt", &(flg.opt), 1);
36 register_integer_arg(arg_parser, "ieee", &(flg.ieee), 0);
37 register_inform_level_arg(arg_parser, "inform",
38 (inform_level_t *)&(flg.inform),
39 LV_Inform);
40 register_integer_arg(arg_parser, "endian", &(flg.endian),
41 0);
42 register_boolean_arg(arg_parser, "mp", (bool *)&(flg.smp),
43 false);
44 register_boolean_arg(arg_parser, "reentrant", &
45 arg_reentrant, false);
46 register_integer_arg(arg_parser, "terse", &(flg.terse),
47 1);
48 register_boolean_arg(arg_parser, "quad", (bool *)&(flg.
49 quad), false);
50 register_boolean_arg(arg_parser, "save", (bool *)&(flg.
51 save), false);
52 register_string_arg(arg_parser, "tp", &tp, NULL);
53 register_integer_arg(arg_parser, "astype", &(flg.astype),
54 0);
55 register_boolean_arg(arg_parser, "recursive", (bool *)&(
56 flg.recursive),
57 false);
58 register_integer_arg(arg_parser, "vect", &(vect_val), 0);
59
60 register_string_arg(arg_parser, "cmdline", &(cmdline),
61 NULL);
62 register_boolean_arg(arg_parser, "debug", (bool *)&(flg.
63 debug), false);
64
65 /* Run argument parser */
66 parse_arguments(arg_parser, argc, argv);

```

Listing 20: Registration of `flang2` command line arguments.

Closer examination of `flang2` source code reveals the meaning of these arguments:

- `-filename` – source `.ilm` file name.
- `-fn` – (the original) Fortran source file name.
- `-stbfile` – input `.stb` symbol table file name.
- `-asm` – output `.ll` LLVM IR text file name.
- `-vh` – build host; does not seem to have any use currently.
- `-x` – the `-x` parameter (see below).
- `-y` – the `-y` parameter (see below).
- `-q` – the `debug` parameter (see further below).
- `-qq` – the `dump` parameter (see further below).

- `-opt` – optimization level.
- `-ieee` – enable or disable IEEE division.
- `-inform` – increase verbosity level of error messages.
- `-endian` – force endianness, where 0 = little, 1 = big.
- `-mp/-nomp` – process OpenMP pragmas.
- `-reentrant/-noreentrant` – local variables should have *RECURSIVE* semantics instead of *SAVE* semantics (same as *recursive/norecursive* below); also inhibit special treatment of terminus functions.
- `-terse` – reduce verbosity of error messages (where 0 = do not reduce).
- `-quad/-noquad` – quad align (round-up size) *unconstrained objects* if size is greater or equal to 16.
- `-save/-nosave` – local variables should have *SAVE* semantics instead of *RECURSIVE* semantics (typically used for overriding OpenMP's default using `-Msave`).
- `-tp` – x86-specific target architecture name, needs to be dropped.
- `-astype` – assembler syntax, where 0 = ELF, 1 = COFF.
- `-recursive/-norecursive` – local variables should have *RECURSIVE* semantics instead of *SAVE* semantics (OpenMP's default, can be forced by `-Mrecursive`).
- `-vect` – vectorizer settings (bitmask), does not seem to be examined anywhere in the frontend.
- `-cmdline` – override command line used to invoke the compiler.
- `-debug/-nodebug` – debug *ON/OFF* flag.

5.1 -x and -y arguments

These flags (also called *xflags*) were added to allow more detailed control (than flags listed above) of the compiler's behavior. The `-y` flag is the inversion of `-x` flag (whenever given `-x` flag instructs to do some action, the `-y` flag with the same arguments instructs not to do it). Multiple `-x` and `-y` arguments can be passed to *flang1* or *flang2* during their invocation. In order to ease the use of this flag in code, the *XBIT* macro was defined as shown in Listing 21. Example of use can be observed in Listing 22.

General usage of these flags when invoking *flang1* or *flang2* directly is: `-x <number><value>` or `-y <number><value>`. Numbers and corresponding values are described in documentation that could be generated from file *xflag.n*. For example, number 124 and value 0x10 for which the *if* condition is checking in Listing 22 is described as follows:

- 124: F77 implementation-defined behavior modifications.
- 0x10: treat INTEGER as INTEGER*8 and LOGICAL as LOGICAL*8

```
1 #define XBIT(n, m) (flg.x[n] & m)
```

Listing 21: *XBIT* macro definition.

```
1 if (XBIT(124, 0x10)) {
2   dt_int = DT_INT8; /* -i8 */
3 } else {
4   dt_int = DT_INT;
5 }
```

Listing 22: *XBIT* macro in use.

Any of these `-x` and `-y` flags can be passed to *flang1* and *flang2* from *flang* command line options: `-Hx,<number>,<value>` passes `-x <number><value>` to *flang1*, `-Mx,<number>,<value>` passes `-x <number><value>` to *flang2* (e.g. `-Mx,183,0x20000`). The same is valid for `-y`. Multiple `-x`, `-y` (and `-Hx`, `-Hy`, `-Mx`, `-My` for the compiler invocation) arguments can be specified within one command line.

Most of the *xflags* values are bit vectors, except those that should hold actual (plain) values (e.g. *xflags* 9 and 10). Listing 23 reveals how these two classes are distinguished.

```
1 /* Does this index correspond to a bitvector (true) or
2    plain value (false) */
3 bool
4 is_xflag_bitvector(int index)
5 {
6     switch (index) {
7         case 9: /* max cnt for unroller */
8         case 10: /* # of times to unroll */
9         case 16: /* lower bound loop iter count for vectorizer */
10        /*
11         case 27: /* default/max overlap size */
12         case 30: /* lower limit on iteration count */
13         case 31: /* lower limit on strip size */
14         case 32: /* cache size */
15         case 33: /* upper limit on strip size in complex-
16            vector loops */
17         case 35: /* max iteration count */
18         case 38: /* MAXVPUS */
19         case 40: /* # array loads/stores */
20         case 41: /* # fp operations */
21         case 43: /* # CPUs */
22         case 44: /* min parallel loop count */
23         case 64: /* code straightening optimizations */
24         case 79: /* CSE of DP loads: distance threshold */
25         case 100: /* break blocks */
26         case 101: /* ST processor SI information */
27         case 105: /* upper limit on unrolling with unroll & jam
28            */
29         case 106: /* set unroll factor for scalar unroll & jam
30            */
31         case 130: /* Levels of VLIW scheduling */
32         case 131: /* Levels of predication */
33         case 133: /* VLIW density */
34         case 134: /* register stall limit */
35         case 138: /* Limit of the #of prefetches for a loop */
36         case 139: /* Limit of #vili for vectorizing single
37            precision loop */
38         case 140: /* Limit of #vili for vectorizing double
39            precision loop */
40         case 157: /* count for unroller for multiblock loops */
41         case 160: /* nest level for intensity count */
42         case 181: /* 3D tile size */
43         case 188: /* openacc default vector length */
44         case 199: /* max blocks in loop to be fused */
45         case 249: /* LLVM version number */
46             return false;
47         default:
48             break;
49     }
50     return true;
51 }
```

Listing 23: *xflags*: bit vectors versus plain values.

5.2 -q and -qq arguments

The `-q` flag configures compile-time debugging behavior. Multiple `-q` (and `-Hq` and/or `-Mq` can be specified for the compiler invocation, similar to `-x` and `-y` shown above) parameters can be specified within one command line.

They share the same syntax as with `-x` parameters: `-q <number><value>`, for example: `-q 0 1` is a very useful setting, which in documentation generated from `coding.n` file is described as follows:

```
0:
1 Use stderr for debug file (instead of ".qdbg").
```

The `-q` parameter not only allows tracing of frontend activity (due to verbose debug output) but also allows to display contents of important data structures. For example, `-Mq,0,1 -Mq,4,1` causes `flang2` to print human-readable ILM representation right after it is imported from `flang1`'s output file. In order to print human-readable ILI representation, one can type: `-Mq,0,1 -Mq,10,2 -Mq,10,1`.

Additionally to the above, `-qq` parameter causes debug routines to dump human-readable views of various internal data structures. The syntax is: `-qq <phase_name><object_name>`. Listing 24 presents how phase names and various object names are defined. For example, in order to dump `flang1`'s AST, one can pass `-Hq,0,1 -Hqq,parser,ast` to the compiler. In order to dump `flang1`'s `syntab`, one can pass `-Hq,0,1 -Hqq,parser,syntab` to the compiler.

```
1 static char *who[] = {"init", "parser", "bblock", "
2 vectorize", "optimize",
3 "schedule", "assemble", "xref", "
4 unroll"};
5
6 /* Create a datastructure of various dump actions and
7 their names */
8 action_map_t *dump_map; /* Deallocated after arguments
9 are parsed */
10 create_action_map(&dump_map);
11 add_action(dump_map, "ast", dump_ast);
12 add_action(dump_map, "dtype", dumpdpts);
13 add_action(dump_map, "std", dump_stdts);
14 add_action(dump_map, "sstd", dump_sstdts);
15 add_action(dump_map, "stdp", dump_stdps);
16 add_action(dump_map, "sym", dump_symbols);
17 add_action(dump_map, "syms", dump_symbols);
18 add_action(dump_map, "syntab", dump_symbols);
19 add_action(dump_map, "allsym", dump_all_symbols);
20 add_action(dump_map, "stats", dump_stg_stats);
21 add_action(dump_map, "area", report_area);
22 add_action(dump_map, "olddtype", dmp_dtype);
23 add_action(dump_map, "odtype", dmp_dtype);
24 add_action(dump_map, "oldsym", dump_old_symbols);
25 add_action(dump_map, "osym", dump_current_symbols);
26 add_action(dump_map, "hsym", dump_current_symbols);
27 add_action(dump_map, "hsyms", dump_current_symbols);
28 add_action(dump_map, "common", dcommons);
29 add_action(dump_map, "commons", dcommons);
30 add_action(dump_map, "nast", dumpasts);
31 add_action(dump_map, "stdtree", dumpstdtrees);
32 add_action(dump_map, "stdtrees", dumpstdtrees);
33 add_action(dump_map, "shape", dumpshapes);
34 add_action(dump_map, "aux", dumplists);
35
36 /* Initialize the map that will be used by dump handler
37 later */
38 create_action_map(&phase_dump_map);
```

Listing 24: phase names and object names that can be passed with `-qq` parameter.

6 CONCLUSIONS

This article concludes as follows:

- The Flang compiler project has successfully introduced a Fortran compiler frontend that utilizes LLVM compilers infrastructure.

- Now that it is publicly available, the next logical step would be to work towards having it accepted as a part of the LLVM project on equal footing with the Clang compiler.
- Being a part of the LLVM ecosystem and using its backend, Flang can target multiple hardware architectures. Nowadays the support is present for AArch64, 64-bit Power and x86_64 architectures, but it should be fairly easy to add any other 64-bit hardware architecture to this list.
- With its source code based on the commercial PGI® Fortran compiler, Flang inherited most of the interior concepts: two-pass parser, semantic analyzer, two levels of internal representation (ILM, ILI) and internal optimizer.
- The source code is written mostly in ANSI C98 with some utilities written in C++ and some of runtime library routines written in Fortran.
- Documentation can be generated from provided `.n` files (`nroff` format).
- Some may criticize that Flang source code still contains too many references to its commercial predecessor. We must keep in mind that this can change in time and is completely outweighed by consequences of having the code base developed for more than two decades – such heritage offers maturity and good Fortran standards conformance [1].
- The generated LLVM IR is outputted into a plain text file. All text operations are performed literally, no LLVM API is used for this process.
- Performance comparison of workloads compiled by Flang and GFortran show that although Flang is doing its job well, there is still work to do in this area. Flang performs most of its optimizations within LLVM passes and this arrangement (Fortran frontend combined with LLVM middleend and backend) is still very young and there is a room for further improvement.

REFERENCES

- [1] Jason Blevins. 2017. Fortran 2003 status. (3 Mar 2017). <http://fortranwiki.org/fortran/show/Fortran+2003+status>
- [2] Walt Brainerd. 2003. The Importance Of Fortran In The 21st Century. *Journal of Modern Applied Statistical Methods* 2, Article 3 (2003). Issue 1.
- [3] LLVM Contributors. 2017. C Language Family Front-end (fork). (2017). <https://github.com/flang-compiler/clang>
- [4] Hal Finkel. 2013. A Fortran language frontend for LLVM. (2013). <https://github.com/hfinkel/lfort>
- [5] Kitware Inc. Ongoing project. CMake. (Ongoing project). <https://cmake.org>
- [6] Don Johnston. 2015. NNSA, national labs team with Nvidia to develop open-source Fortran compiler technology. (13 Nov 2015). <https://www.llnl.gov/news/nnsa-national-labs-team-nvidia-develop-open-source-fortran-compiler-technology>
- [7] Michael Larabel. 2017. FLANG: NVIDIA Brings Fortran To LLVM. (18 May 2017). https://www.phoronix.com/scan.php?page=news_item&px=LLVM-NVIDIA-Fortran-Flang
- [8] Eugene Loh. 2010. The Ideal HPC Programming Language. Maybe it's Fortran. Or maybe it just doesn't matter. *ACM Queue* 8 (2010). Issue 6.
- [9] Aleksei Lorenz. 2013. End of GSoC report. (23 Sep 2013). <http://flang-gsoc.blogspot.co.uk>
- [10] Jussi Pakkanen. 2013. Why you should consider using separate build directories. (16 Apr 2013). <http://voices.canonical.com/jussi.pakkanen/2013/04/16/why-you-should-consider-using-separate-build-directories>
- [11] PGI. 2017. (2017). <https://github.com/flang-compiler/flang>
- [12] PGI. 2017. flang-compiler. (2017). <https://github.com/flang-compiler>
- [13] PGI. 2017. flang-compiler on Slack. (2017). <https://flang-compiler.slack.com/messages>
- [14] PGI. 2017. Join flang-compiler on Slack. (2017). https://join.slack.com/t/flang-compiler/shared_invite/MjExOTEyMzQ3MjIxLkTE0OTk4NzQyNzUtODQzZWEyMjkWYw

- [15] Louis-Noël Pouchet, Mohanish Narayan, and Uday Bondugula. 2012. PolyBench-Fortran, the Polyhedral Benchmark suite. (2012). <http://web.cse.ohio-state.edu/~pouchet.2/software/polybench/polybench-fortran.html>
- [16] James Price. 2017. [llvm-dev] LLVM Fortran front-end. (18 May 2017). <http://lists.lvm.org/pipermail/llvm-dev/2017-May/113131.html>
- [17] Fortran UK. 2015. The Polyhedron Fortran Benchmarks Suite. (2015). <https://www.fortran.uk/fortran-compiler-comparisons/the-polyhedron-solutions-benchmark-suite>
- [18] Bill Wendling. 2012. Fortran Front-End. (2012). <https://github.com/isanbard/flang>