

# GCC Internals



Diego Novillo  
dnovillo@google.com

November 2007



## 1. Overview

- Features, history, development model

## 2. Source code organization

- Files, building, patch submission

## 3. Internal architecture

- Pipeline, representations, data structures, alias analysis, data flow, code generation

## 4. Passes

- Adding, debugging

Internal information valid for GCC mainline as of 2007-11-20

# 1. Overview

- Major features
- Brief history
- Development model

## Availability

- Free software (GPL)
- Open and distributed development process
- System compiler for popular UNIX variants
- Large number of platforms (deeply embedded to big iron)
- Supports all major languages: C, C++, Java, Fortran 95, Ada, Objective-C, Objective-C++, etc

## Code quality

- Bootstraps on native platforms
- Warning-free
- Extensive regression testsuite
- Widely deployed in industrial and research projects
- Merit-based maintainership appointed by steering committee
- Peer review by maintainers
- Strict coding standards and patch reversion policy

## Analysis/Optimization

- SSA-based high-level global optimizer
- Constraint-based points-to alias analysis
- Data dependency analysis based on chains of recurrences
- Feedback directed optimization
- Inter-procedural optimization
- Automatic pointer checking instrumentation
- Automatic loop vectorization
- OpenMP support

- Major features

- **Brief history**

- Development model

## **GCC 1 (1987)**

- Inspired on Pastel compiler (Lawrence Livermore Labs)
- Only C
- Translation done one statement at a time

## **GCC 2 (1992)**

- Added C++
- Added RISC architecture support
- Closed development model challenged
- New features difficult to add

## EGCS (1997)

- Fork from GCC 2.x
- Many new features: Java, Chill, numerous embedded ports, new scheduler, new optimizations, integrated libstdc++

## GCC 2.95 (1999)

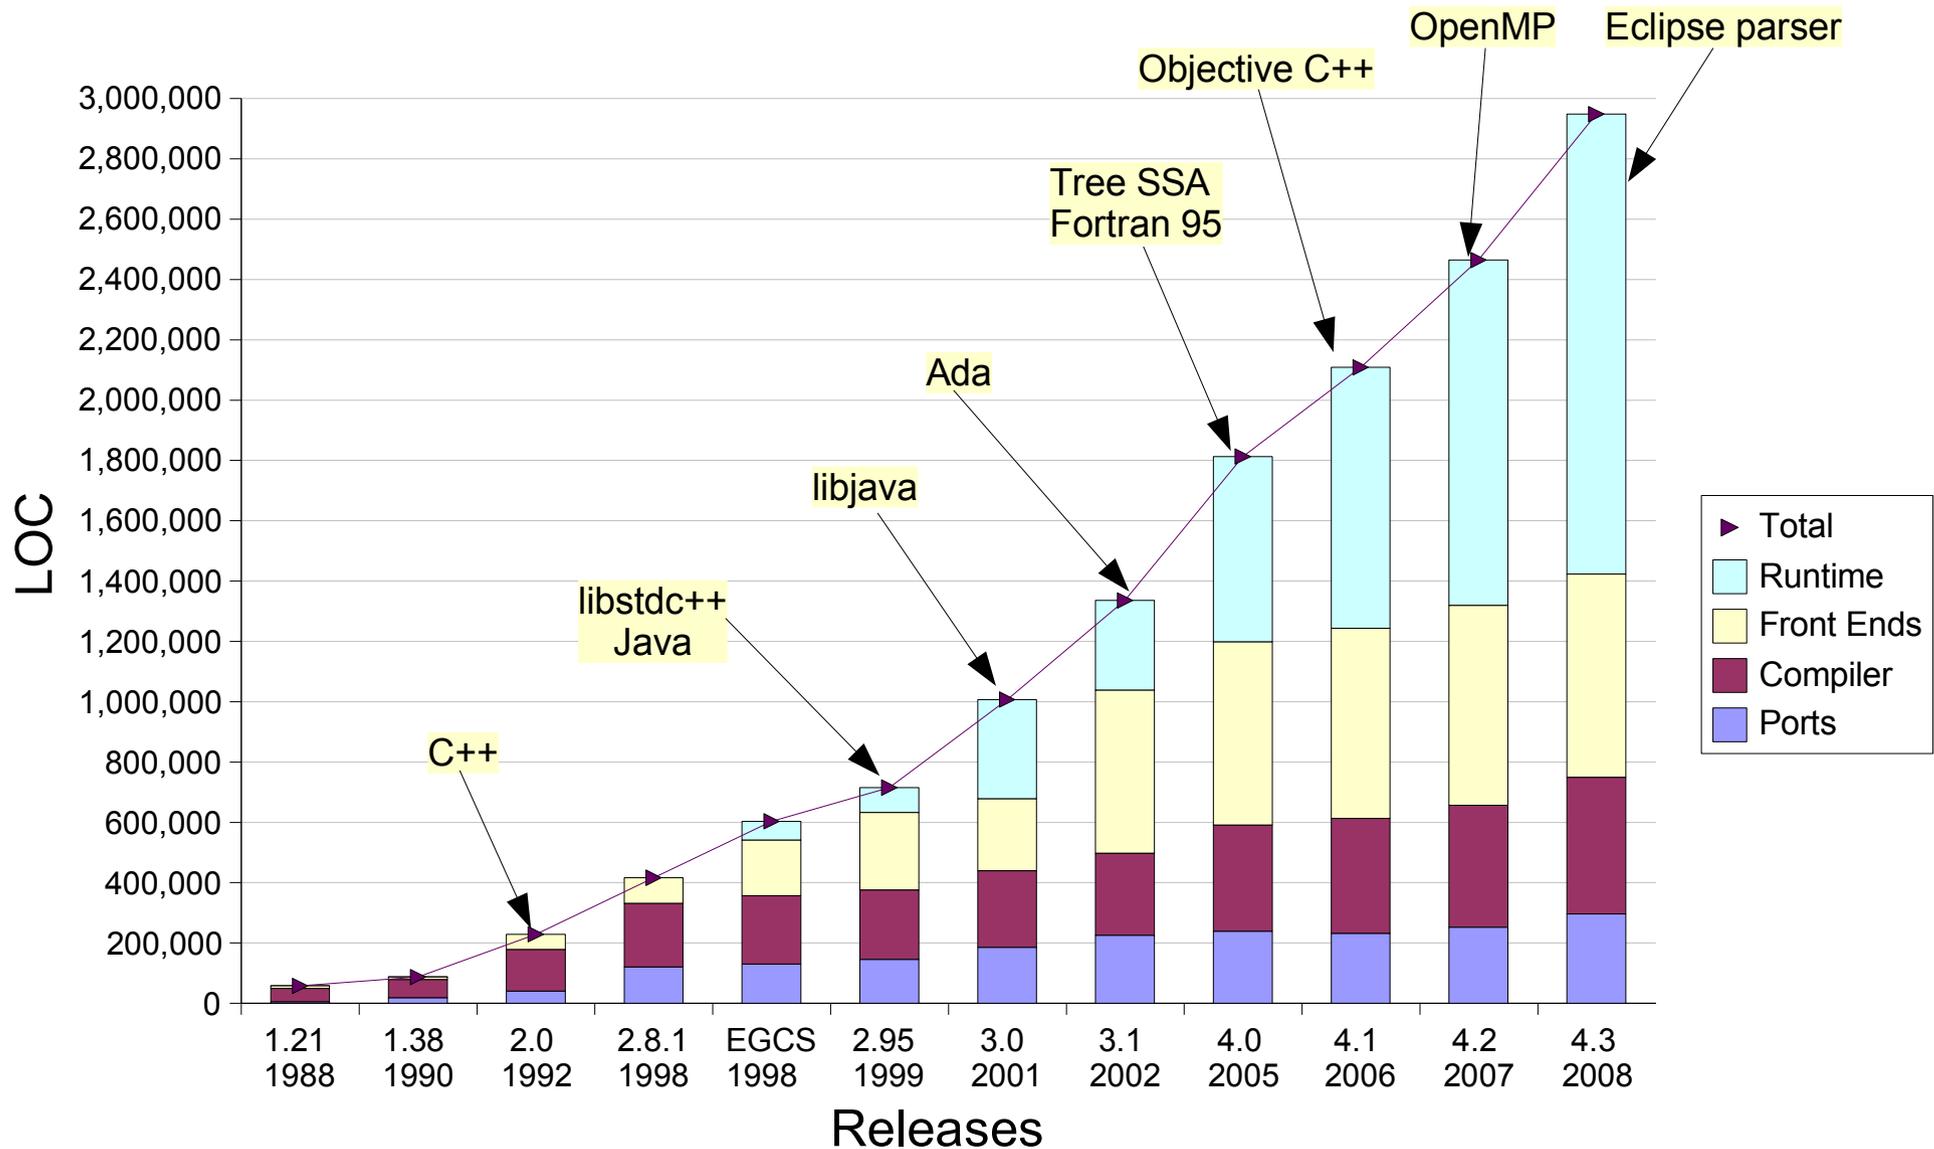
- EGCS and GCC2 merge into GCC
- Type based alias analysis
- Chill front end
- ISO C99 support

## **GCC 3 (2001)**

- Integrated libjava
- Experimental SSA form on RTL
- Functions as trees

## **GCC 4 (2005)**

- Internal architecture overhaul (Tree SSA)
- Fortran 95
- Automatic vectorization



<sup>1</sup>generated using David A. Wheeler's 'SLOCCount'.

# 1. Overview

- Major features

- Brief history

- **Development model**

- Project organization
  - Steering Committee → Administrative, political
  - Release Manager → Release coordination
  - Maintainers → Design, implementation
- Three main stages (~2 months each)
  - Stage 1 → Big disruptive changes.
  - Stage 2 → Stabilization, minor features.
  - Stage 3 → Bug fixes only (bugzilla).

- Major development is done in branches
  - Design/implementation discussion on public lists
  - Frequent merges from mainline
  - Final contribution into mainline only at stage 1 and approved by maintainers
- Anyone with SVN write-access may create a development branch
- Vendors create own branches from FSF release branches

- All contributors **must** sign FSF copyright release
  - Even when working on branches
- Three levels of access
  - Snapshots (weekly)
  - Anonymous SVN
  - Read/write SVN
- Two main discussion lists
  - [gcc@gcc.gnu.org](mailto:gcc@gcc.gnu.org)
  - [gcc-patches@gcc.gnu.org](mailto:gcc-patches@gcc.gnu.org)

# 2. Source code

- Source tree organization
- Configure, build, test
- Patch submission

- Getting the code for mainline (or trunk)

```
$ svn co svn://gcc.gnu.org/svn/gcc/trunk
```

- Build requirements (<http://gcc.gnu.org/install>)

- ISO C90 compiler

- GMP library

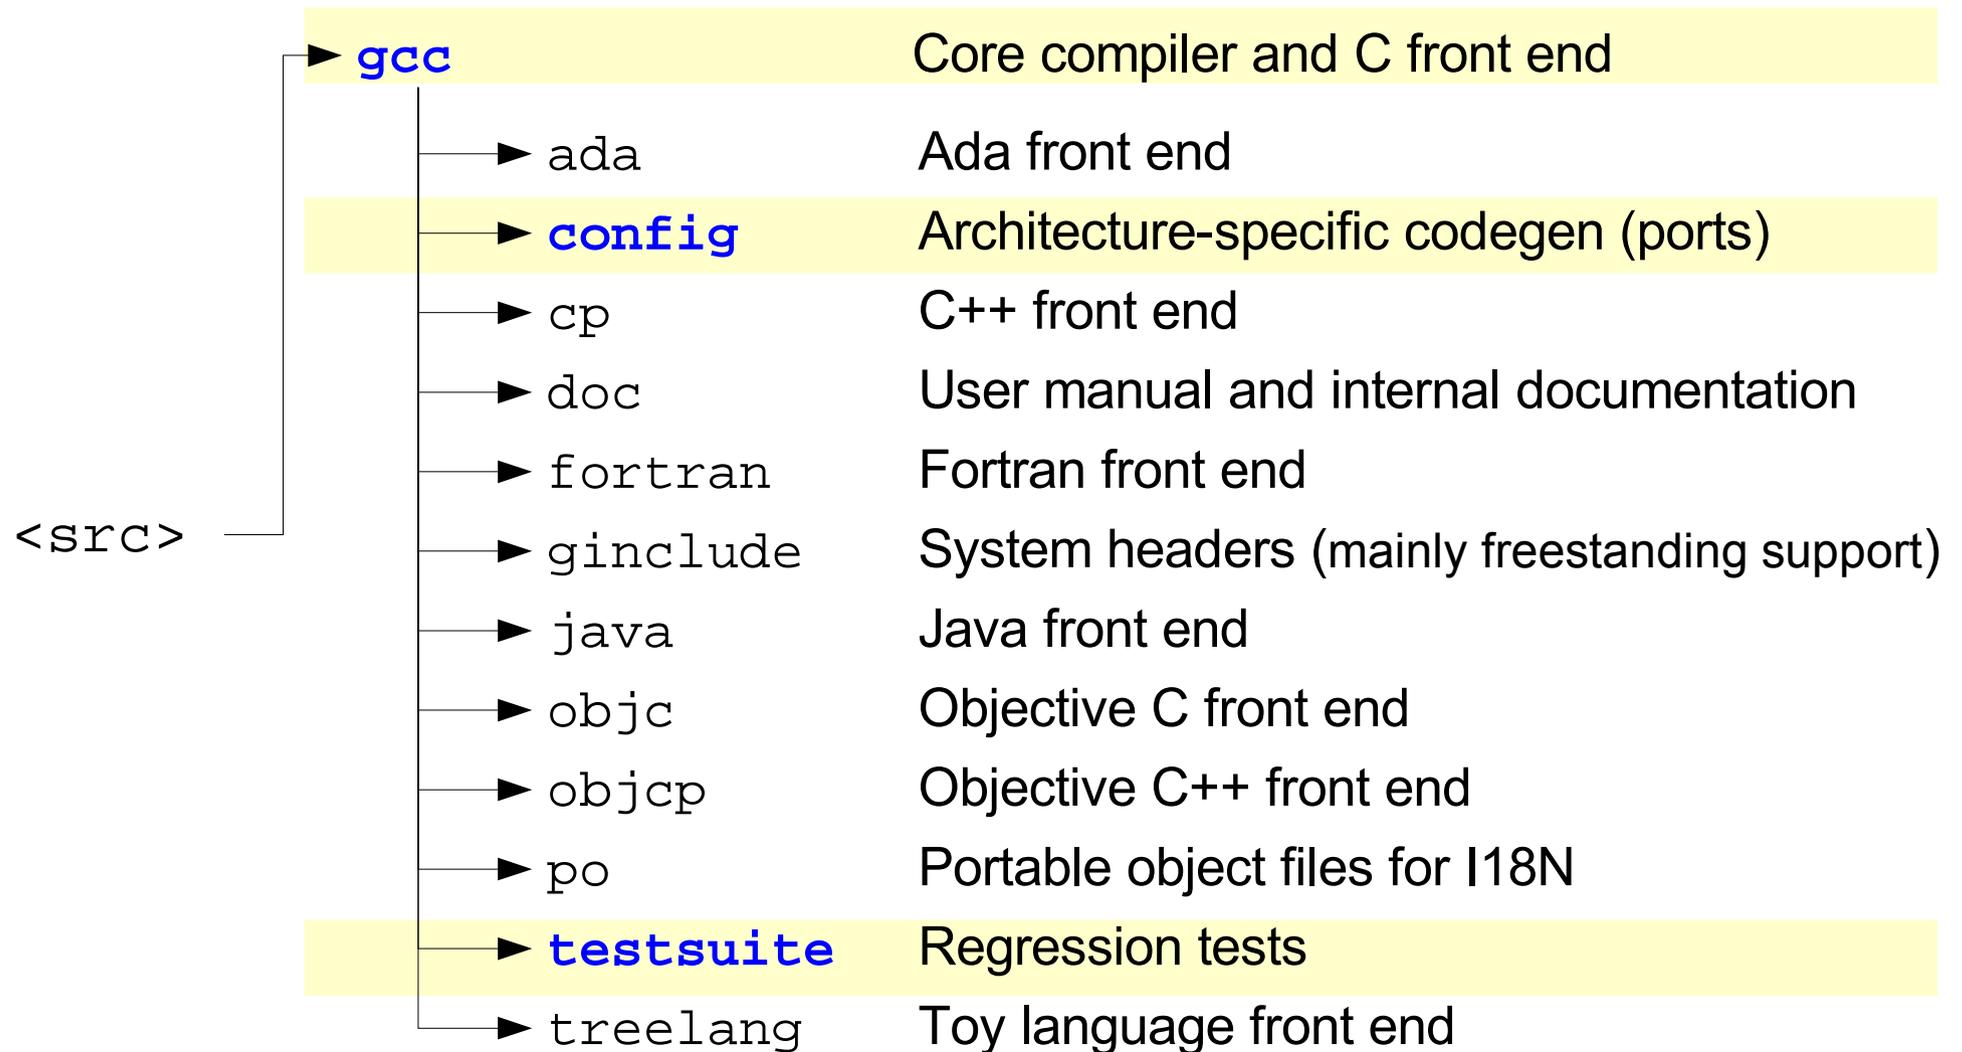
- MPFR library

- GNAT (only if building Ada)

Multiple precision floating point libraries

- Source code includes runtimes for all languages and extensive regression testsuite.

<code>&lt;src&gt;</code>	▶ <code>gcc</code>	Front/middle/back ends
	▶ <code>libgcc</code>	Internal library for missing target features
	▶ <code>libcpp</code>	Pre-processor
	▶ <code>libada</code>	Ada runtime
	▶ <code>libstdc++-v3</code>	C++ runtime
	▶ <code>libgfortran</code>	Fortran runtime
	▶ <code>libobjc</code>	Objective-C runtime
	<code>boehm-gc</code>	Java runtime
	▶ <code>libffi</code>	
	▶ <code>libjava</code>	
	<code>zlib</code>	
	▶ <code>libiberty</code>	Utility functions and generic data structures
	▶ <code>libgomp</code>	OpenMP runtime
	▶ <code>libssp</code>	Stack Smash Protection runtime
▶ <code>libmudflap</code>	Pointer/memory check runtime	
▶ <code>libdecnumber</code>	Decimal arithmetic library	



- Alias analysis
- Build support
- C front end
- CFG and callgraph
- Code generation
- Diagnostics
- Driver
- Profiling
- Internal data structures
- Mudflap
- OpenMP
- Option handling
- RTL optimizations
- Tree SSA optimizations

# 2. Source code

- Source tree organization
- **Configure, build, test**
- Patch submission

```
$ svn co svn://gcc.gnu.org/svn/gcc/trunk
```

```
$ mkdir bld && cd bld
```

```
$ ../trunk/configure --prefix=`pwd`
```

```
$ make all install
```

- Bootstrap is a 3 stage process
  - Stage 0 (host) compiler builds Stage 1 compiler
  - Stage 1 compiler builds Stage 2 compiler
  - Stage 2 compiler builds Stage 3 compiler
  - Stage 2 and Stage 3 compilers must be binary identical

## **--prefix**

- Installation root directory

## **--enable-languages**

- Comma-separated list of language front ends to build

- Possible values

`ada, c, c++, fortran, java, objc, obj-c++, treelang`

- Default values

`c, c++, fortran, java, objc`

## **--disable-bootstrap**

- Build stage 1 compiler only

## **--target**

- Specify target architecture for building a cross-compiler
- Target specification form is (roughly)  
cpu-manufacturer-os  
cpu-manufacturer-kernel-os  
e.g. x86\_64-unknown-linux-gnu  
arm-unknown-elf
- All possible values in `<src>/config.sub`

## `--enable-checking=list`

- Perform compile-time consistency checks
- List of checks: `assert fold gc gcac misc rtl rtlflag runtime tree valgrind`

### – Global values:

`yes` → `assert,misc,tree,gc,rtlflag,runtime`

`no` → Same as `--disable-checking`

`release` → Cheap checks `assert,runtime`

`all` → Everything except `valgrind`

**SLOW!**



## **-j N**

- Usually scales up to 1.5x to 2x number of processors

## **all**

- Default `make` target. Knows whether to bootstrap or not

## **install**

- Not necessary but useful to test installed compiler
- Set `LD_LIBRARY_PATH` afterward

## **check**

- Use with `-k` to prevent stopping when some tests fail

- Staged compiler binaries

- ① `<bld>/stage1-{gcc,intl,libcpp,libdecnumber,libiberty}`

- ② `<bld>/prev-{gcc,intl,libcpp,libdecnumber,libiberty}`

- ③ `<bld>/{gcc,intl,libcpp,libdecnumber,libiberty}`

- Runtime libraries are not staged, except `libgcc`

- `<bld>/<target-triplet>/lib*`

- Testsuite results

- `<bld>/gcc/testsuite/*.{log,sum}`

- `<bld>/<target-triplet>/lib*/testsuite/*.{log,sum}`

- Compiler is split in several binaries

<code>&lt;bld&gt;/gcc/xgcc</code>	Main driver
<code>&lt;bld&gt;/gcc/cc1</code>	C compiler
<code>&lt;bld&gt;/gcc/cc1plus</code>	C++ compiler
<code>&lt;bld&gt;/gcc/jc1</code>	Java compiler
<code>&lt;bld&gt;/gcc/f951</code>	Fortran compiler
<code>&lt;bld&gt;/gcc/gnat1</code>	Ada compiler

- Main driver forks one of the \*1 binaries
- `<bld>/gcc/xgcc -v` shows what compiler is used

- The best way is to have two trees built
  - pristine
  - pristine + patch

- Pristine tree can be recreated with

```
$ cp -a trunk trunk.pristine
```

```
$ cd trunk.pristine
```

```
$ svn revert -R .
```

- Configure and build both compilers with the exact same flags

- Use `<src>/trunk/contrib/compare_tests` to compare individual `.sum` files

```
$ cd <bld>/gcc/testsuite/gcc
```

```
$ compare_tests <bld.pristine>/gcc/testsuite/gcc/gcc.sum gcc.sum
```

```
Tests that now fail, but worked before:
```

```
gcc.c-torture/compile/20000403-2.c -Os (test for excess errors)
```

```
Tests that now work, but didn't before:
```

```
gcc.c-torture/compile/20000120-2.c -O0 (test for excess errors)
```

```
gcc.c-torture/compile/20000405-2.c -Os (test for excess errors)
```

# 2. Source code

- Source tree organization
- Configure, build, test
- **Patch submission**

- Non-trivial contributions require copyright assignment
- Code should follow the GNU coding conventions
  - [http://www.gnu.org/prep/standards\\_toc.html](http://www.gnu.org/prep/standards_toc.html)
  - <http://gcc.gnu.org/codingconventions.html>
- Submission should include
  - ChangeLog describing what changed (not how nor why)
  - Test case (if applicable)
  - Patch itself generated with `svn diff` (context or unified)

- When testing a patch

1. Disable bootstrap
2. Build C front end only
3. Run regression testsuite
4. Once all failures have been fixed
  - Enable all languages
  - Run regression testsuite again
5. Enable bootstrap
6. Run regression testsuite

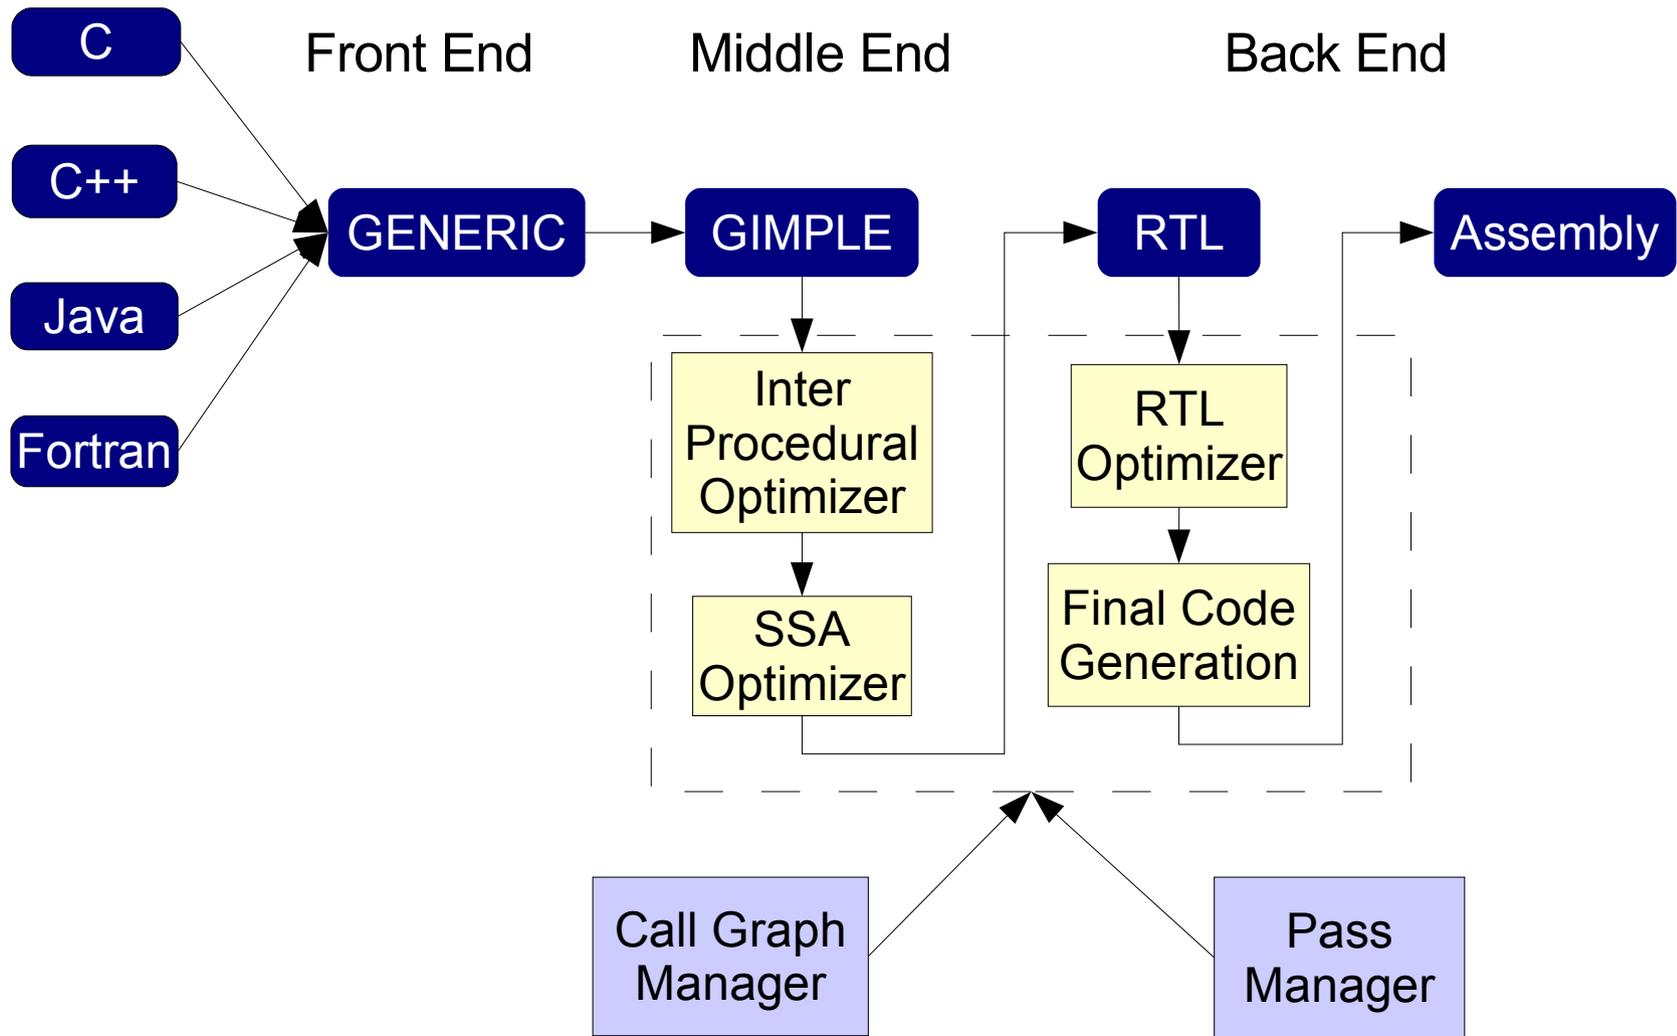
1-4 not strictly necessary, but recommended

- Patches are only accepted after #5 and #6 work on 1+ major platform (more than one sometimes).

# 3. Internal architecture

- **Compiler pipeline**
- Intermediate representations
- CFG, statements, operands
- Alias analysis
- SSA forms
- Code generation

# Compiler pipeline



- Operate on GIMPLE
- Around 100 passes
  - Vectorization
  - Various loop optimizations
  - Traditional scalar optimizations: CCP, DCE, DSE, FRE, PRE, VRP, SRA, jump threading, forward propagation
  - Field-sensitive, points-to alias analysis
  - Pointer checking instrumentation for C/C++
  - Interprocedural analysis and optimizations: CCP, inlining, points-to analysis, pure/const and type escape analysis

- Around 70 passes
- Operate closer to the target
  - Register allocation
  - Scheduling
  - Software pipelining
  - Common subexpression elimination
  - Instruction recombination
  - Mode switching reduction
  - Peephole optimizations
  - Machine specific reorganization

# 3. Internal architecture

- Compiler pipeline
- **Intermediate representations**
- CFG, statements, operands
- Alias analysis
- SSA forms
- Code generation

- **GENERIC** is a common representation shared by all front ends
  - Parsers may build their own representation for convenience
  - Once parsing is complete, they emit **GENERIC**
- **GIMPLE** is a simplified version of **GENERIC**
  - 3-address representation
  - Restricted grammar to facilitate the job of optimizers

## GENERIC

```
if (foo (a + b, c))  
    c = b++ / a  
endif  
return c
```

## High GIMPLE

```
t1 = a + b  
t2 = foo (t1, c)  
if (t2 != 0)  
    t3 = b  
    b = b + 1  
    c = t3 / a  
endif  
return c
```

## Low GIMPLE

```
t1 = a + b  
t2 = foo (t1, c)  
if (t2 != 0) <L1, L2>  
L1:  
    t3 = b  
    b = b + 1  
    c = t3 / a  
    goto L3  
L2:  
L3:  
return c
```

- Register Transfer Language  $\approx$  assembler for an abstract machine with infinite registers
- It represents low level features
  - Register classes
  - Memory addressing modes
  - Word sizes and types
  - Compare-and-branch instructions
  - Calling conventions
  - Bitfield operations
  - Type and sign conversions

`b = a - 1`



```
(set (reg/v:SI 59 [ b ])  
      (plus:SI (reg/v:SI 60 [ a ]  
                (const_int -1 [0xffffffff]))))
```

- It is commonly represented in LISP-like form
- Operands do not have types, but type modes
- In this case they are all `SImode` (4-byte integers)

# 3. Internal architecture

- Compiler pipeline
- Intermediate representations
- **Control/data structures**
- Alias analysis
- SSA forms
- Code generation

- Every internal/external function is a node of type `struct cgraph_node`
- Call sites represented with edges of type `struct cgraph_edge`
- Every `cgraph` node contains
  - Pointer to function declaration
  - List of callers
  - List of callees
  - Nested functions (if any)
- Indirect calls are not represented

- Callgraph manager drives intraprocedural optimization passes
- For every node in the callgraph, it sets `cfun` and `current_function_decl`
- IPA passes must traverse callgraph on their own
- Given a `cgraph` node

```
DECL_STRUCT_FUNCTION (node->decl)
```

points to the `struct function` instance that contains all the necessary control and data flow information for the function

- Built early during lowering
- Survives until late in RTL
  - Right before machine dependent transformations (`pass_machine_reorg`)
- In GIMPLE, instruction stream is physically split into blocks
  - All jump instructions replaced with edges
- In RTL, the CFG is laid out over the double-linked instruction stream
  - Jump instructions preserved

- Every CFG accessor requires a `struct function` argument
- In intraprocedural mode, accessors have shorthand aliases that use `cfun` by default
- CFG is an array of double-linked blocks
- The same data structures are used for GIMPLE and RTL
- Manipulation functions are callbacks that point to the appropriate RTL or GIMPLE versions

# 3. Internal architecture

- Compiler pipeline
- Intermediate representations
- Control/data structures
- **Alias analysis**
- SSA forms
- Code generation

- GIMPLE represents alias information explicitly
- Alias analysis is just another pass
  - Artificial symbols represent memory expressions (virtual operands)
  - FUD-chains computed on virtual operands → Virtual SSA
- Transformations may prove a symbol non-addressable
  - Promoted to GIMPLE register
  - Requires another aliasing pass

- Points-to alias analysis (PTAA)
  - Based on constraint graphs
  - Field and flow sensitive, context insensitive
  - Intra-procedural (inter-procedural in 4.2)
  - Fairly precise
- Type-based analysis (TBAA)
  - Based on input language rules
  - Field sensitive, flow insensitive
  - Very imprecise

- Two kinds of pointers are considered
  - Symbols: Points-to is flow-insensitive
    - Associated to Symbol Memory Tags (SMT)
  - SSA names: Points-to is flow-sensitive
    - Associated to Name Memory Tags (NMT)
- Given pointer dereference  $*ptr_{42}$ 
  - If  $ptr_{42}$  has NMT, use it
  - If not, fall back to SMT associated with  $ptr$

- Separate structure fields are assigned distinct symbols

```
struct A
{
  int x;
  int y;
  int z;
};
```

```
struct A a;
```

- Variable a will have 3 sub-variables  
{ SFT.1, SFT.2, SFT.3 }
- References to each field are mapped to the corresponding sub-variable

```
foo (i, a, b, *p)
{
  p =(i > 10) ? &a : &b
  *p = 3
  return a + b
}
```

```
foo (i, a, b, *p)
{
  p = (i > 10) ? &a : &b
```

```
# a = VDEF <a>
# b = VDEF <b>
*p = 3
```

```
# VUSE <a>
t1 = a
```

```
# VUSE <b>
t2 = b
```

```
t3 = t1 + t2
return t3
}
```

- Pure query system
- Pairwise disambiguation of memory references
  - Does store to A affect load from B?
  - Mostly type-based (same predicates used in GIMPLE's TBAA)
- Very little information passed on from GIMPLE

- Some symbolic information preserved in RTL memory expressions
  - Base + offset associated to aggregate refs
  - Memory symbols
- Tracking of memory addresses by propagating values through registers
- Each pass is responsible for querying the alias system with pairs of addresses

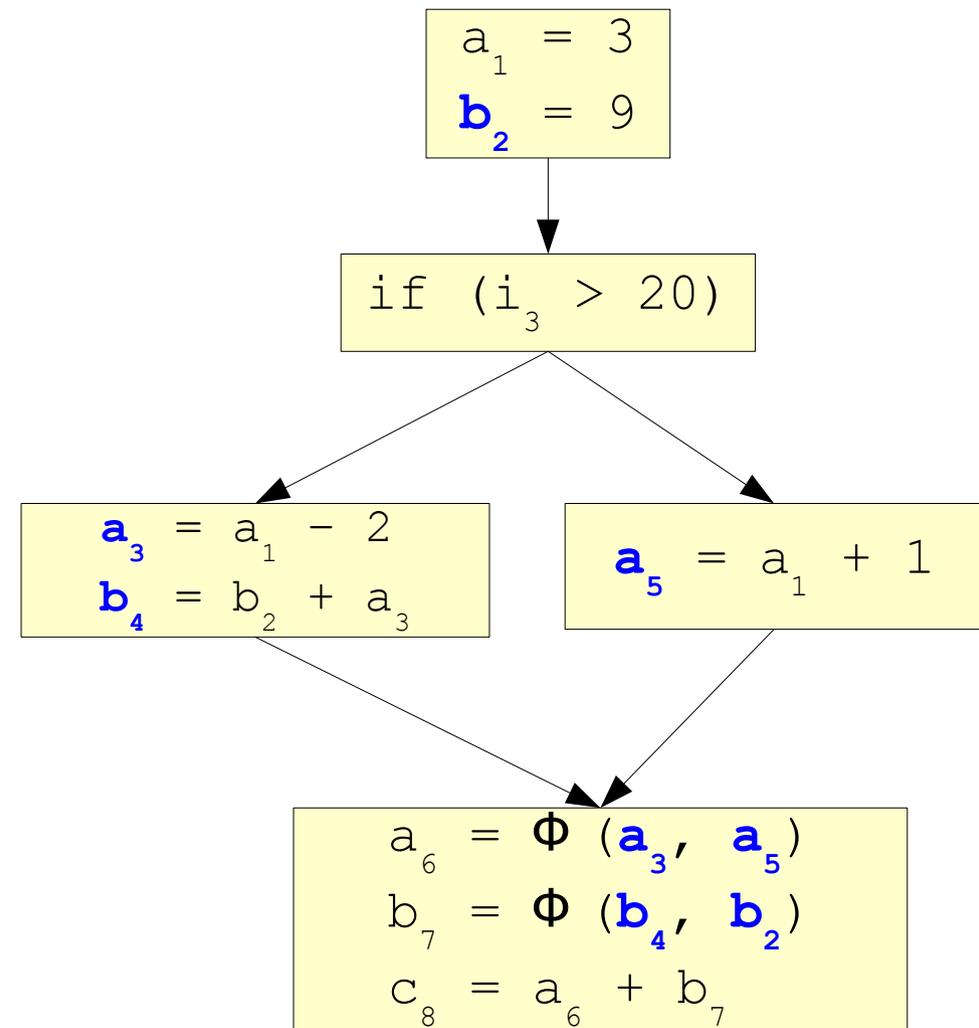
- Big impedance between GIMPLE and RTL
  - No/little information transfer
  - Producers and consumers use different models
  - GIMPLE → explicit representation in IL
  - RTL → query-based disambiguation
- Work underway to resolve this mismatch
  - Results of alias analysis exported from GIMPLE
  - Adapt explicit representation to query system

# 3. Internal architecture

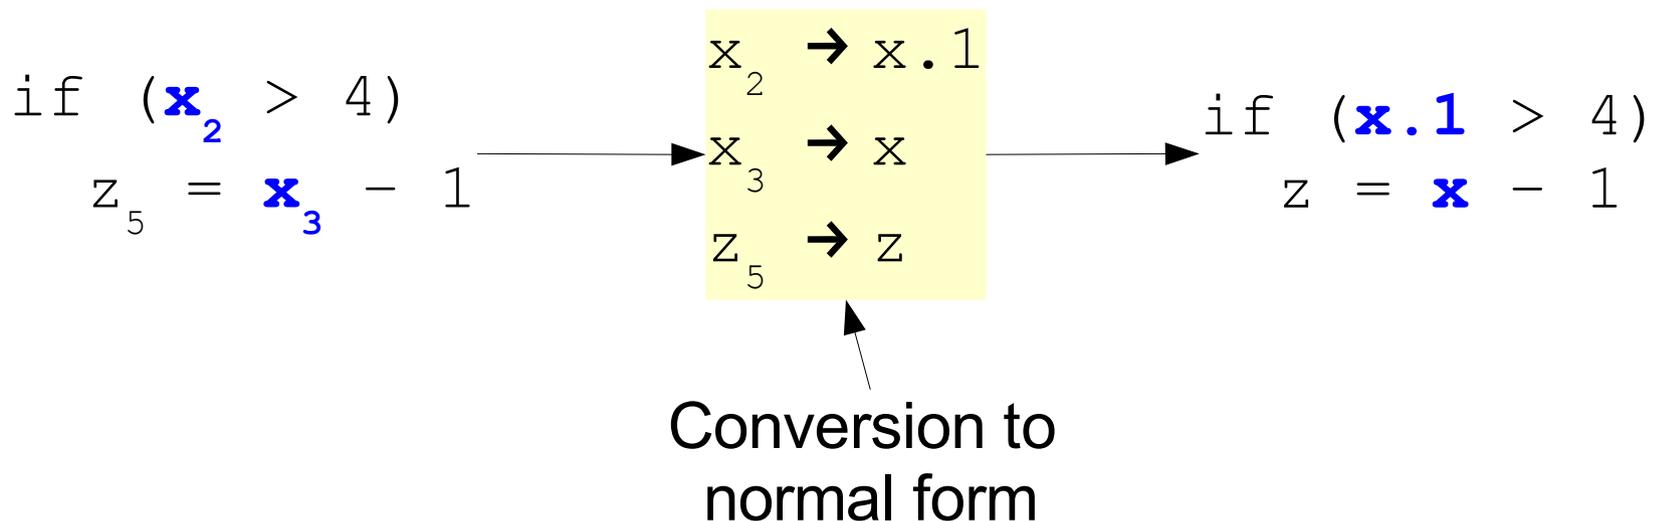
- Compiler pipeline
- Intermediate representations
- Control/data structures
- Alias analysis
- **SSA forms**
- Code generation

## Static Single Assignment (SSA)

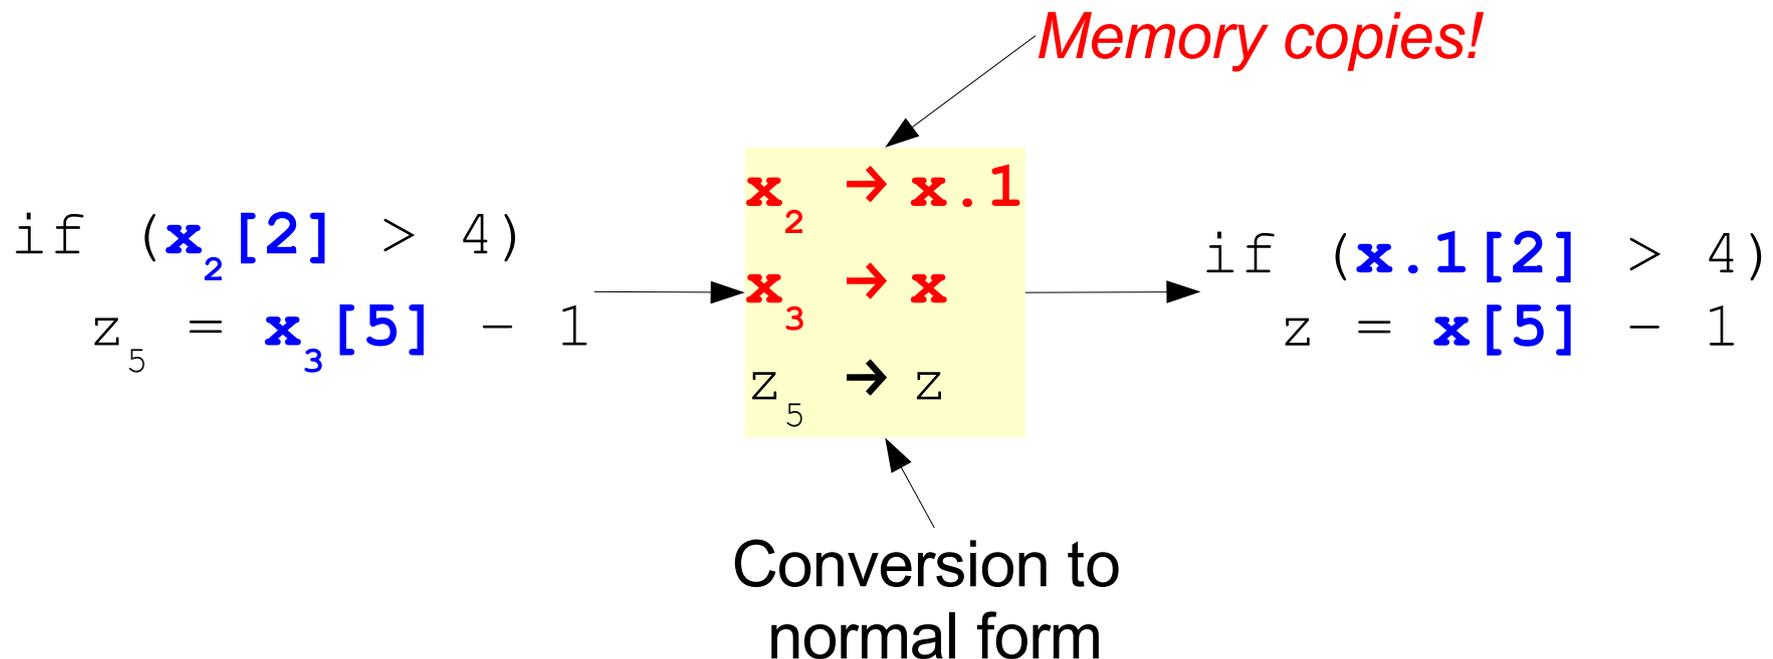
- Versioning representation to expose data flow explicitly
- Assignments generate new versions of symbols
- Convergence of multiple versions generates new one ( $\Phi$  functions)



- Rewriting (or standard) SSA form
  - Used for real operands
  - Different names for the same symbol are *distinct objects*
  - overlapping live ranges (OLR) are allowed
  - Program is taken out of SSA form for RTL generation (new symbols are created to fix OLR)



- Factored Use-Def Chains (FUD Chains)
  - Also known as Virtual SSA Form
  - Used for virtual operands.
  - All names refer to the *same object*.
  - Optimizers may **not** produce OLR for virtual operands.



- VDEF operand needed to maintain DEF-DEF links
- They also prevent code movement that would cross stores after loads
- When alias sets grow too big, static grouping heuristic reduces number of virtual operators in aliased references

```
foo (i, a, b, *p)
{
    p_2 = (i_1 > 10) ? &a : &b

    # a_4 = VDEF <a_11>
    a = 9;

    # a_5 = VDEF <a_4>
    # b_7 = VDEF <b_6>
    *p_2 = 3;

    # VUSE <a_5>
    t1_8 = a;

    t3_10 = t1_8 + 5;
    return t3_10;
}
```

SSA forms are kept up-to-date incrementally

## Manually

- As long as SSA property is maintained, passes may introduce new SSA names and PHI nodes on their own
- Often this is the quickest way

## Automatically using `update_ssa`

- Marking individual symbols (`mark_sym_for_renaming`)
- name → name mappings (`register_new_name_mapping`)
- Passes that invalidate SSA form must set `TODO_update_ssa`

# 3. Internal architecture



- Compiler pipeline
- Intermediate representations
- Control/data structures
- Alias analysis
- SSA forms
- **Code generation**

- Code is generated using a rewriting system

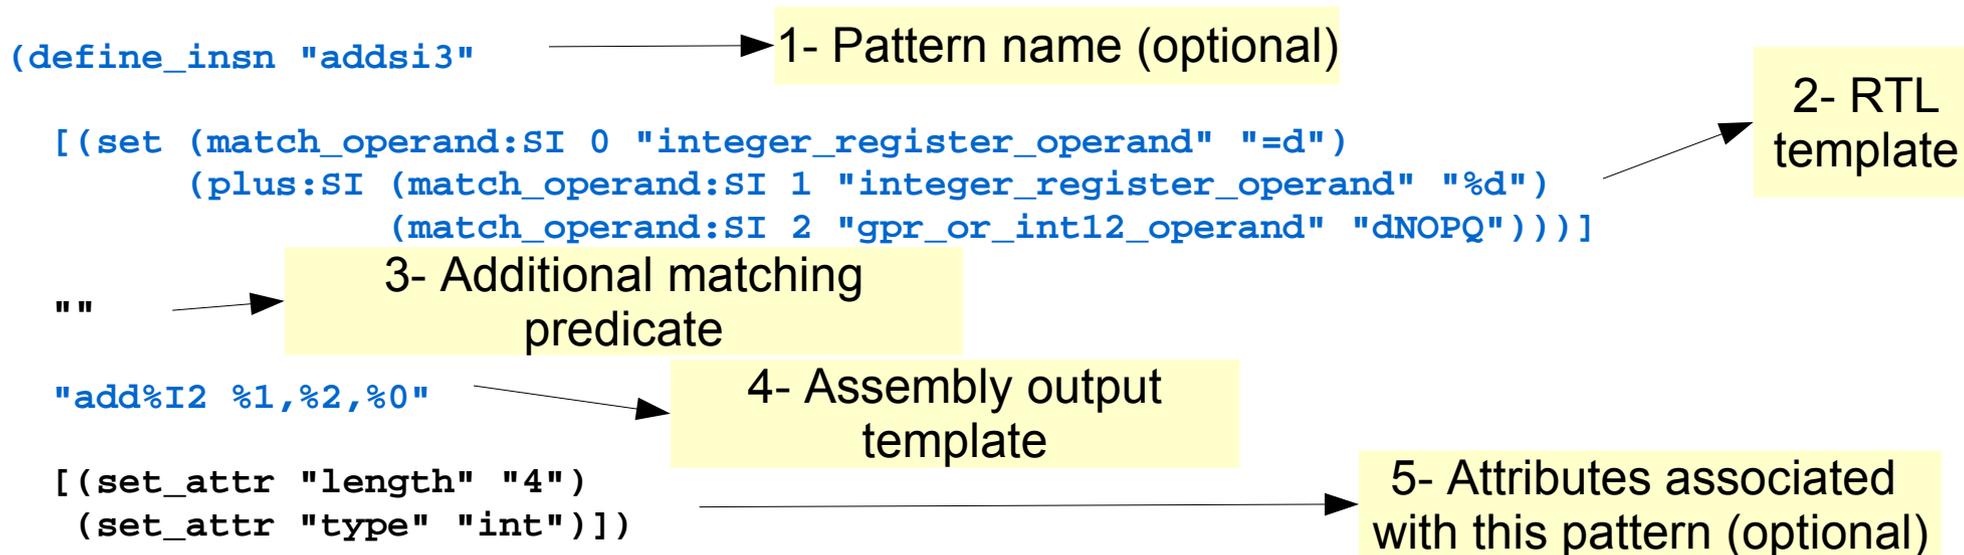
- Target specific configuration files in

`gcc/config/<arch>`

- Three main target-specific files

- `<arch>.md`      Code generation patterns for RTL insns
- `<arch>.h`      Definition of target capabilities (register classes, calling conventions, type sizes, etc)
- `<arch>.c`      Support functions for code generation, predicates and target variants

- Two main types of rewriting schemes supported
  - Simple mappings from RTL to assembly (`define_insn`)
  - Complex mappings from RTL to RTL (`define_expand`)
- `define_insn` patterns have five elements



```
define_insn    addsi3
```

- Named patterns
  - Used to generate RTL
  - Some standard names are used by code generator
  - Some missing standard names are replaced with library calls (e.g., `divsi3` for targets with no division operation)
  - Some pattern names are mandatory (e.g. move operations)
- Unnamed (anonymous) patterns do not generate RTL, but can be used in insn combination

```
[ (set (match_operand:SI 0 "integer_register_operand" "=d,=d")  
      (plus:SI (match_operand:SI 1 "integer_register_operand" "%d,m")  
              (match_operand:SI 2 "gpr_or_int12_operand" "dNOPQ,m" ))) ]
```

## Matching uses

Machine mode (SI, DI, HI, SF, etc)

Predicate (a C function)

Both operands and operators can be matched

Constraints provide second level of matching

Select best operand among the set of allowed operands

Letters describe kinds of operands

Multiple alternatives separated by commas

```
"add%I2 %1, %2, %0"
```

- Code is generated by emitting strings of target assembly
- Operands in the insn pattern are replaced in the `%n` placeholders
- If constraints list multiple alternatives, multiple output strings must be used
- Output may be a simple string or a C function that builds the output string

- Some standard patterns cannot be used to produce final target code. Two ways to handle it
  - Do nothing. Some patterns can be expanded to libcalls
  - Use `define_expand` to generate matchable RTL
- Four elements
  - The name of a standard insn
  - Vector of RTL expressions to generate for this insn
  - A C expression acting as predicate to express availability of this instruction
  - A C expression used to generate operands or more RTL

```
(define_expand "ashlsi3"  
  [(set (match_operand:SI 0 "register_operand" "")  
        (ashift:SI  
          (match_operand:SI 1 "register_operand" "")  
          (match_operand:SI 2 "nonmemory_operand" "")))]  
  ""  
  "{  
    if (GET_CODE (operands[2]) != CONST_INT  
        || (unsigned) INTVAL (operands[2]) > 3)  
      FAIL;  
  }")
```

- Generate a left shift only when the shift count is [0...3]
- **FAIL** indicates that expansion did not succeed and a different expansion should be tried (e.g., a library call)
- **DONE** is used to prevent emitting the RTL pattern. C fragment responsible for emitting all insns.

# 4. Passes

- Adding a new pass
- Debugging dumps

- To implement a new pass
  - Add a new file to `trunk/gcc` or edit an existing pass
  - Add a new target rule in `Makefile.in`
  - If a flag is required to trigger the pass, add it to `common.opt`
  - Create an instance of `struct tree_opt_pass`
  - Declare it in `tree-pass.h`
  - Sequence it in `init_optimization_passes`
  - Add a gate function to read the new flag
  - Document pass in `trunk/gcc/doc/invoke.texi`

- APIs available for
  - CFG: block/edge insertion, removal, dominance information, block iterators, dominance tree walker.
  - Statements: insertion in block and edge, removal, iterators, replacement.
  - Operands: iterators, replacement.
  - Loop discovery and manipulation.
  - Data dependency information (scalar evolutions framework).

- Other available infrastructure
  - Debugging dumps (`-fdump-tree-...`)
  - Timers for profiling passes (`-ftime-report`)
  - CFG/GIMPLE/SSA verification (`--enable-checking`)
  - Generic value propagation engine with callbacks for statement and  $\Phi$  node visits.
  - Generic use-def chain walker.
  - Support in test harness for scanning dump files looking for specific transformations.
  - Pass manager for scheduling passes and describing interdependencies, attributes required and attributes provided.

# 4. Passes

- Adding a new pass

- **Debugging**

Most passes understand the `-fdump` switches

`-fdump-<ir>-<pass>[-<flag1>[-<flag2>]...]`

ipa  
tree  
rtl

- details, stats, blocks, ...
- all enables all flags
- Possible values taken from array `dump_options`

- inline, dce, alias, combine ...
- all to enable all dumps
- Possible values taken from `name` field in struct `tree_opt_pass`

- Adding dumps to your pass
  - Specify a name for the dump in `struct tree_opt_pass`
  - To request a dump at the end of the pass add `TODO_dump_func` in `todo_flags_finish` field
- To emit debugging information during the pass
  - Variable `dump_file` is set if dumps are enabled
  - Variable `dump_flags` is a bitmask that specifies what flags were selected
  - Some common useful flags: `TDF_DETAILS`, `TDF_STATS`

- Never debug the `gcc` binary, that is only the driver
- The real compiler is one of `cc1`, `jc1`, `f951`, ...

```
$ <bld>/bin/gcc -O2 -v -save-temps -c a.c
```

```
Using built-in specs.
```

```
Target: x86_64-unknown-linux-gnu
```

```
Configured with: [ ... ]
```

```
[ ... ]
```

```
End of search list.
```

```
<path>/cc1 -fpreprocessed a.i -quiet -dumpbase a.c  
-mtune=generic -auxbase a -O2 -version -o a.s
```

```
$ gdb --args <path>/cc1 -fpreprocessed a.i -quiet -dumpbase  
a.c -mtune=generic -auxbase a -O2 -version -o a.s
```

- The build directory contains a `.gdbinit` file with many useful wrappers around debugging functions
- When debugging a bootstrapped compiler, try to use the stage 1 compiler
- The stage 2 and stage 3 compilers are built with optimizations enabled (may confuse debugging)
- To recreate test suite failures, cut and paste command line from

```
<bld>/gcc/testsuite/{gcc,gfortran,g++,java}/*.log
```

# Current and Future Projects

- Delay optimization until link time
  - IR for each compilation unit (CU) is streamed out
  - Multiple CUs are read and combined together
- Enables “whole program mode” optimization
  - Increased optimization opportunities
- Challenges
  - Compile time
  - Memory consumption
  - Combining CUs from different languages

- Extensibility mechanism to allow 3<sup>rd</sup> party tools
- Wrap some internal APIs for external use
- Allow loading of external shared modules
  - Loaded module becomes another pass
  - Compiler flag determines location
- Versioning scheme prevents mismatching
- Useful for
  - Static analysis
  - Experimenting with new transformations

- Several concurrent efforts targetting 4.3 and 4.4
  - Schedule over larger regions for increased parallelism
  - Most target IA64, but benefit all architectures
- Enhanced selective scheduling
- Treeregion scheduling
- Superblock scheduling
- Improvements to swing modulo scheduling

- Several efforts over the years
  - Complex problem
  - Many different targets to handle
  - Interactions with reload and scheduling
- YARA (Yet Another Register Allocator)
  - Experimented with several algorithms
- IRA (Integrated Register Allocator)
  - Priority coloring, Chaitin-Briggs and region based
  - Expected in 4.4
  - Currently works on x86, x86-64, ppc, IA64, sparc, s390

- SSA may cause excessive register pressure
  - Pathological cases → ~800 live registers
  - RA battle lost before it begins
- Short term project to cope with RA deficiencies
- Implement register pressure reduction in GIMPLE before going to RTL
  - Pre-spilling combined with live range splitting
  - Load rematerialization
  - Tie RTL generation into out-of-ssa to allow better instruction selection for spills and rematerialization

- Delay compilation until runtime (JIT)
  - Emit bytecodes
  - Implement virtual machine with optimizing transformations
- Leverage on existing infrastructure (LLVM, LTO)
- Not appropriate for every case
- Challenges
  - Still active research
  - Different models/costs for static and dynamic compilers

- Speed up edit-compile-debug cycle
- Speeds up ordinary compiles by compiling a given header file “once”
- Incremental changes fed to compiler daemon
- Incremental linking as well
- Side effects
  - Refactoring
  - Cross-referencing
  - Compile-while-you-type (e.g., Eclipse)

- Phase ordering not optimal for every case
- Current static ordering difficult to change
- Allow external re-ordering
  - Ultimate control
  - Allow experimenting with different orderings
  - Define `-On` based on common orderings
- Problems
  - Probability of finding bugs increases
  - Enormous search space

- Home page <http://gcc.gnu.org/>
- Wiki <http://gcc.gnu.org/wiki>
- Mailing lists  
[gcc@gcc.gnu.org](mailto:gcc@gcc.gnu.org)  
[gcc-patches@gcc.gnu.org](mailto:gcc-patches@gcc.gnu.org)
- IRC <irc.oft.net/#gcc>