

ROSE Compiler Framework/Print version

Contents

ROSE Compiler Framework/Print version	1
Contents	1
About the Book	5
Mission and Milestones	5
Mission Statement.....	5
Use cases	6
ROSE on ROSE	6
ROSE on DOE applications.....	6
ROSE for average programmers	6
Core Technologies	6
AST	7
bin(tools) directory.....	7
Compilation.....	8
Translation	8
Analysis.....	8
Optimization	9
Usability	9
Installation.....	9
Documentation	10
API	10
Portability.....	10
Workflow	11
ROSE's Documentations	11
Obtaining ROSE	12
git 1.7.10 or later for github.com	12
Installation.....	13

Platform Requirement	13
Software Requirement	13
Installing boost.....	14
Installing Java JDK.....	14
./build	15
configure	15
make.....	15
make check.....	16
make install	16
set environment variables	16
try out a rose translator	16
ROSE tools.....	17
identityTranslator	17
Supported Programming Languages.....	17
OpenMP support	17
CUDA support	18
Abstract Syntax Tree (Intermediate Representation).....	18
visualization of AST	18
text output of AST	19
preprocessing info.....	19
Program Translation.....	19
Expected behavior of a ROSE Translator.....	19
SageBuilder and SageInterface	20
workflow for writing translators	20
Order to traverse AST.....	20
example translators	21
Program Analysis.....	21
control flow graph.....	21
virtual control flow graph	21
static control flow graph	22
static and interprocedural CFGs.....	22
Virtual function analysis	22
def-use analysis.....	22
pointer analysis	23

SSA	24
Generic dataflow framework	24
Dependence analysis	25
Program Optimizations	25
Developer's Guide	25
Basic skills for ROSE developers	25
Expected Deliverables for Interns	26
Don't use branches, use separate repositories	26
Summary of development tools	27
code review	27
issue tracking	27
How-tos	28
Add a new project	28
sys admin's notes	28
Code Review Process	29
Motivation	29
Goals	29
Software	30
Github	30
Developer Checklist	30
Coding Standards	30
Workflow	31
Reviewer Checklist	32
what to avoid	33
criticism	33
references	33
Testing	34
Workflow	34
Motivation and Goals	34
Review of workflow changes	34
Current workflow	35
Frequently Asked Questions (FAQ)	36
How to search rose-public mailinglist for previously asked questions?	36
Compilation	36

How to speedup compiling ROSE?	36
Can ROSE accept incomplete code?.....	36
Can ROSE analyze Linux Kernel sources?	37
Can ROSE compile C++ Boost library?	37
AST	37
How to find XYZ in AST?.....	37
How does the AST merge work?	38
How to filter out header files from AST traversals?.....	38
Should SgIfStmt::get_true_body() return SgBasicBlock?.....	38
How to handle #include "header.h", #if, #define etc. ?	39
SgClassDeclaration::get_definition() returns NULL?	39
Translation	40
Can ROSE identityTranslator generate 100% identical output file?	40
How to build a tool inserting function calls?	40
How to copy/clone a function?	41
Can I transform code within a header file?.....	42
how to work with formal and actual arguments of functions?.....	42
how to clone a function?.....	43
Daily work	43
git clone returns error: SSL certificate problem?.....	44
What is the best IDE for ROSE developers?	44
Portability.....	44
What is the status for supporting Windows?	45
Who is using ROSE	45
Universities	45
DOE national laboratories.....	45
TODO List	45

About the Book

The goal of this book is to have a **community documentation** providing extensive and up-to-date instructional information about how to use the ROSE compiler framework, developed at Lawrence Livermore National Laboratory.

While the ROSE project website (<http://www.rosecompiler.org>) already has a variety of official documentations, having a wikibook for ROSE **allows anybody to contribute** to gathering information about this software.

Again, please note that this wikibook is not the official documentation of ROSE. It is the community efforts contributed by anyone just like you.

Mission and Milestones

Some brainstorming for milestones of ROSE. The order of the sections is used to reflect the priorities. The most important ones come first.

Mission Statement

Our primary mission is to optimize applications within the [U.S. Department of Energy \(DOE\)](#).

Goa: To develop the **BEST** open-source, `[[w:Source-to-source_compiler|source-to-source compiler]]` infrastructure in the world.

We achieve this goal through:

- **cutting-edge research** on source- and high-level compiler analysis and optimization algorithms
- **best-practice software engineering** to encapsulate existing and newly-developed compiler techniques into easy-to-use APIs
- **pre-built ROSE tools** to perform program transformation, analysis and optimization of your code
- **an easy-to-use API** to help you to build your own customized, domain-specific compiler tools

Focus: To make advanced compiler techniques accessible to non-expert compiler or tool writers.

Benefit to You: Improved programmer productivity, source code correctness, performance, and efficiency.

ROSE Is...

- A **research platform**
- A **library** (and set of associated tools) to quickly and easily apply compiler techniques to your code
- A **compiler infrastructure** for you to write your own custom source-to-source translators to perform source code transformations, analyses, and optimizations.

Use cases

ROSE on ROSE

We use ROSE in our own daily work. In other words, we [eat our own dog food](#).

ROSE uses ROSE to:

- Build ROSE translators to improve the code quality of ROSE
 - e.g., refactor code
- Run static analysis tools built on ROSE to improve the code quality of ROSE
 - enforce coding conventions
 - catch errors not found by compilers and other commercial tools
- Run dynamic analysis tools built on ROSE to catch run-time errors in ROSE

ROSE on DOE applications

ROSE is a [DOE](#)-fund software project within [LLNL](#). A priority of ROSE is to serve the mission of LLNL and build up the software capability of DOE.

ROSE is being used to analyze, translate, and optimize DOE applications to make them run faster and more efficiently on high performance computing platforms.

ROSE for average programmers

ROSE is funded by the Office of Science of DOE and is released under a BSD-like license. So everybody can obtain it to use the pre-built tools shipped with ROSE and/or build them custom tools using ROSE.

Core Technologies

The uniqueness of ROSE lies on its high level program representation and its associated compiler analysis, translation, and optimizations.

AST

The Abstract Syntax Tree (AST) is the intermediate representation (IR) of programs in ROSE. It is important to have an intuitive program representation.

TODO:

- clean up name conventions of AST nodes, member functions, member data
- Using Doxygen to document each node for its purpose: like which languages are using it, the corresponding code constructs, etc
- document important data member and member functions
- use major versions to separate out big changes to names

Doxygen supports using a separated file to document classes/functions. For example, for SgProject node, we can add documents for it in docs/testDoxygen/SgProject.docs

bin(tools) directory

A place to provide popular and hardened tools built using ROSE. Users can use them directly and immediately get real sense about how powerful/useful ROSE can be. We also use these tools in our daily work.

tool list

- dot graph generator
- identityTranslator
- call graph generator
- control flow graph generator
- outliner
- inliner
- loop transformation tools
- constant folding tool
- static analysis tool (compass)
- runtime error checking tool(RTED)
- automatic parallelization tool (autoPar)

So sample tools from tutorials and immature projects will not conflict with real hardened tools.

By hardening these tools, we can really improve

- intra- and inter-procedural analysis
- global analysis using AST merge

Previous efforts in this direction:

- rose/tools: to be removed or merged into the new bin directory

Compilation

This is really no point to talk about a compiler if it cannot compile the code you are interested in.

Enrich the supported benchmarks for C, C++, Fortran, etc.

- ROSE itself: so we can eat our own dog food
- DOE/LLNL applications
- spec cpu benchmark: many conferences in compiler need proof of impact using industrial benchmarks
- Linux kernel
- boost C++ library
- Plum Hall test suites and others

Fix blocking bugs

Translation

With the AST representing an input program, the immediately possible thing is to do program translation by restructuring the AST.

List

- refine SageBuilder so no low level AST constructors are needed to create AST pieces
- refine SageInterface so AST insertion, deletion, copying is simpler

Sample translators built using ROSE

- AST outliner
- OpenMP implementation
- Cross language translation:
 - Fortran 77 to C
 - C++ to C

Analysis

The core of compiler technology is the set of sophisticated compiler analyses. These analyses are the basis for advanced compiler optimizations.

Implement and harden common analysis

- program representation for analysis: often raw AST may not be sufficient or efficient
 - control flow graph: intra-procedural and inter-procedural
 - dominance analysis
 - call graph
 - SSA, including array SSA
- data flow framework: enabling writing a set of data flow analyses, including
 - def-use
 - liveness analysis
 - constant propagation
 - side-effect analysis
 - alias, points-to analysis
 - dependence analysis

Extend the analysis support for multiple languages

- C/C++
- Fortran
- OpenMP

Optimization

Implementing an optimization often involving calling some relevant analyses and then doing some translations.

List

- constant folding
- partial redundancy elimination
- common subexpression elimination
- loop optimizations
- parallelization
- vectorization

Usability

Installation

The first step every ROSE user must encounter is to install ROSE. We should try our best to simplify, speedup, and automate this step as much as possible.

List

- ultimately: one step installation like `apt-get install rose`
- more realistic: two step process
 - `apt-get install rose-prerequisite`

- ./build, configure, make, make install
- milestones, support two-step process on
 - ubuntu
 - Centos
- speedup the installation: goal is 30 minutes or less
- release a VM version of ROSE (use VeeWee, Vagrant, etc.)

Documentation

Types of documentation

- Doxygen documents
 - we promote source comment style so they can be automatically processed to generate class/function reference web pages.
- Design and implementation docs
 - written in LaTeX, reusable for writing academic papers, grant proposals.
 - separate them out from the rose source base
- Developer/User guide
 - we are experimenting with Markdown using github
 - this wikibook is another experiment
 - How-tos: like how to document an AST node
 - coding standard/convention: directory layout
- Tutorials
 - with embedded source code examples which must be tested with the latest version of ROSE
 - choice 1: keep adding the existing tutorial written in LaTeX
 - choice 2: experimenting .md

API

ROSE essentially is a library encapsulating compiler techniques. We try to design intuitive, organized, and well documented API so programmers can easily leverage sophisticated compiler technology.

Checklist

- define top level namespaces to organize functionalities
- encapsulate core technologies of ROSE into simple functions

Portability

The goal here is to make ROSE widely available on mainstream platforms.

Support

- popular Linux distributions and their major versions

- RedHat Enterprise Workstation or its open-source version, Centos
 - Ubuntu LTS
 - Fedora
- Mac OS X
- Windows: native support, not through Cygwin

Often, the key is to support

- more recent versions of GCC
- Boost C++ library

Workflow

Quality comes from a good process.

See details at [Workflow](#)

Streamlined, simplified, and automated workflow involving both users and developers to improve the quality of ROSE and simplify our daily work.

- review of workflow changes

Current components and their roles:

- wikibook:
 - big pictures and milestones
 - instructional tutorials, how-tos, FAQ etc.
 - coding standard/convention: file names, directory layout
- mailing list: interaction with users, feel users' need
- redmine: create projects based on milestones and user input, create and track tasks
- Jenkins: continuous integration of new features, bugfixes
- github:
 - internal:code review only,
 - external: hosting code, issue tracking
 - "rosebot" to automate Github workflow: preliminary testing, policies (git-hooks), automatically add reviewers, etc.
- website: content management system hooked up with all other components

ROSE's Documentations

ROSE uses a range of materials to document the project.

- ROSE manual: the design, algorithm, and implementation details. Written in LaTeX, the content of the manual can come from published papers. It may contain intense academic citations and math formula.

- ROSE tutorial: code examples for tools built on top of ROSE, step-by-step instructions for doing things
- Doxygen web reference: class/namespace references of source code
- this wikibook: non-official, community documentation. Editable by anyone, aimed to supplement official documents and to collect tutorials, FAQ and quick pointers to important topics.

Obtaining ROSE

ROSE's source files are managed by git, a distributed revision control and source code management system. There are several ways to download the source tree:

- Private Git repos within LLNL
 - Private Git repository hosted within Lawrence Livermore National Laboratory: the internal file path is `/usr/casc/overture/ROSE/git/ROSE.git`: central repo of ROSE, mostly automatically updated by Jenkins only after incoming commits pass all regression tests
 - Private Git repository hosted by `github.llnl.gov`: used for daily pushes and code review
- Public repositories
 - Public Git repository hosted at <https://github.com/rose-compiler/rose>: the content is identical to the private Git repository's master branch at LLNL, except that the proprietary EDG submodule is not released.
 - Downloadable packages and a subversion repository (synchronized with stable snapshots of ROSE's git repository): <https://outreach.scidac.gov/projects/rose/>

git 1.7.10 or later for github.com

github requires git 1.7.10 or later to avoid HTTPS cloning errors, as mentioned at <https://help.github.com/articles/https-cloning-errors>

Ubuntu 10.04's package repository has git 1.7.0.4. So building later version of git is needed. But you still need an older version of git to get the latest version of git.

```
apt-get install git-core
```

Now you can clone the latest git

```
git clone https://github.com/git/git.git
```

Install all prerequisite packages needed to build git from source files (assuming you already installed GNU tool chain with GCC compiler, make, etc.)

```
sudo apt-get install gettext zlib1g-dev asciidoc libcurl4-openssl-dev
$ cd git # enter the cloned git directory
```

```
$ make configure ;# as yourself
$ ./configure --prefix=/usr ;# as yourself
$ make all doc ;# as yourself
# make install install-doc install-html;# as root
```

Installation

ROSE is released as an open source software package. Users are expected to compile and install the software.

Platform Requirement

ROSE is portable to Linux and Mac OS X on IA-32 and x86-64 platforms. In particular, ROSE developers often use the following development environments:

- Red Hat Enterprise Linux 5.6 or its open source equivalent [Centos 5.6](#)
- Ubuntu 10.04.4 LTS. Higher versions of Ubuntu are NOT supported due to the GCC versions supported by ROSE.
- Mac OS X 10.5 and 10.6

Software Requirement

Here is a list for prerequisite software packages for installing ROSE

- GCC 4.0.x to 4.4.x , the range of supported GCC versions is checked by [support-rose.m4](#) during configuration
 - gcc
 - g++
 - gfortran (optional for Fortran support)
- GNU autoconf >=2.6 and automake >= 1.9.5, GNU m4 >=1.4.5
- libtool
- bison (byacc),
- flex
- glibc-devel
- Sun Java JDK
- git
- boost library: version 1.36 to 1.48. Again the range of supported Boost versions is checked by [support-rose.m4](#) during configuration
- ZGRViewer, a GraphViz/DOT Viewer: essential to view dot graphs of ROSE AST
 - install Graphviz first - Graph Visualization Software

Optional packages for additional features or advanced users

- libxml2-devel

- sqlite
- texlive-full, need for building LaTeX docs

Installing boost

The installation of Boost may need some special attention.

Download a supported boost version from
<http://sourceforge.net/projects/boost/files/boost/>

For version 1.36 to 1.38

```
./configure --prefix=/home/usera/opt/boost-1.35.0
make
make install
```

Ignore the warning like : Unicode/ICU support for Boost.Regex?... not found.

For version 1.39 and 1.48: create the boost installation directory first

In boost source tree

- ./bootstrap.sh --prefix=your_boost_install_path
- ./bjam install --prefix=your_boost_install_path --libdir=your_boost_install_path/lib

Remember to export LD_LIBRARY_PATH for the installed boost library, for example

```
D_LIBRARY_PATH=$LD_LIBRARY_PATH:/home/leo/opt/boost_1.45.0_inst/lib
export PATH LD_LIBRARY_PATH
```

Installing Java JDK

Download Java SE JDK from
<http://www.oracle.com/technetwork/java/javase/downloads/index.html>

For example, you can download jdk-7u5-linux-i586.tar.gz for your Linux 32-bit system.

After untar it to your installation path, remember to set environment variables for Java JDK

```
# jdk path should be search first before other paths
PATH=/home/leo/opt/jdk1.7.0_05/bin:$PATH

# lib path for libjvm.so
D_LIBRARY_PATH=$LD_LIBRARY_PATH:/home/leo/opt/jdk1.7.0_05/jre/lib/i386/server
```

```
# Don't forget to export both variables!!
export PATH LD_LIBRARY_PATH
```

./build

In general, it is better to rebuild the configure file in the top level source directory of ROSE. Just type:

```
rose_sourcetree>./build
```

configure

The next step is to run configure in a separated build tree. ROSE will complain if you try to build it within its source directory.

There are many configuration options. You can see the full list of options by typing `./sourcetree/configure --help` . But only `--prefix` and `--with-boost` are required as the minimum options.

```
mkdir buildrose
cd buildrose
../rose_sourcetree/configure --prefix=/home/user/opt/rose_tux284 --
with-boost=/home/user/opt/boost-1.36.0/
```

ROSE's configure turns on debugging option by default. The generated object files should already have debugging information.

Additional useful configure options

- Specify where a gcc's OpenMP runtime library libgomp.a is located. Only GCC 4.4's gomp lib should be used to have OpenMP 3.0 support
 - `--with-gomp_omp_runtime_library=/usr/apps/gcc/4.4.1/lib/`

make

In ROSE's build tree, type

```
cd buildrose
make -j4
```

will build the entire ROSE, including librose.so, tutorials, projects, tests, and so on. `-j4` means to use four processes to perform the build. You can have bigger numbers if your machine supports more concurrent processes. Still, the entire process will take hours to finish.

For most users, building librose.so should be enough for most of their work. In this case, just type

```
make -C src/ -j4
```

make check

Optionally, you can type `make check` to make sure the compiled rose pass all its shipped tests. This takes hours again to go through all make check rules within projects, tutorial, and tests directories.

To save time, you can just run partial tests under a selected directory, like the `buildrose/tests`

```
make -C tests/ check -j4
```

make install

After "make", it is recommended to run "make install" so rose's library (librose.so), headers (rose.h) and some prebuilt rose-based tools can be installed under the specified installation path using `--prefix`.

set environment variables

After the installation, you should set up some standard environment variables so you can use rose. For bash, the following is an example:

```
ROSE_INS=/home/userx/opt/rose_installation_tree
PATH=$PATH:$ROSE_INS/bin
LD_LIBRARY_PATH=$LD_LIBRARY_PATH:$ROSE_INS/lib
# Don't forget to export variables !!!
export PATH LD_LIBRARY_PATH
```

try out a rose translator

There are quite some pre-built rose translators installed under `$ROSE_INS/bin`.

You can try `identityTranslator`, which just parses input code, generates AST, and unparses it back to original code:

```
identityTranslator -c helloWorld.c
```

It should generate an output file named `rose_helloWorld.c`, which should just look like your input code.

ROSE tools

ROSE is a compiler framework to build customized compiler-based tools. A set of example tools are provided as part of the ROSE release to demonstrate the use of ROSE. Some of them are also useful for daily work of ROSE developers.

We list and briefly explain some tools built using ROSE. They are installed under `ROSE_INSTALLATION_TREE/bin`.

identityTranslator

This is the simplest tool built using ROSE. It takes input source files, builds AST, and then unparses the AST back to compilable source code. It tries its best to preserve everything from the input file. But due to limitations of the frontends and the internal processing, it cannot generate 100% identical output compared to the input file.

Some notable changes it may introduce include:

- "int a, b, c;" are transformed to three `SgVariableDeclaration` statements,
- macros are expanded.
- extra brackets are added around constants of typedef types (e.g. `c=Typedef_Example(12);` is translated in the output to `c = Typedef_Example((12));`)
- Converting `NULL` to `0`.

Supported Programming Languages

ROSE supports a wide range of main stream programming languages, with different degrees of maturity. The list of supported languages includes:

- C and C++: based on the [EDG C++ frontend](#) version 3.3.
 - An ongoing effort is to upgrade the EDG frontend to its recent 4.4 version.
 - Another ongoing effort is to use clang as an alternative, open-source C/C++ frontend
- Fortran 77/95/2003: based on the [Open Fortran Parser](#)
- OpenMP 3.0: based on ROSE's own parsing and translation support for both C/C++ and Fortran OpenMP programs.
- UPC 1.1: this is also based on the EDG 3.3 frontend

OpenMP support

ROSE supports OpenMP 3.0 for C/C++ (and limited Fortran support).

- The ROSE manual has a chapter (Chapter 12 OpenMP Support) explaining the details. [pdf](#)
- A paper was published for the uniqueness of the ROSE OpenMP Implementation [pdf](#)
- Frontend parsing source files (ompparser.yy and ompFortranParser.C) are located under <https://github.com/rose-compiler/rose/tree/master/src/frontend/SageIII>
- The transformation of OpenMP into threaded code is located in omp_lowering.cpp, under <https://github.com/rose-compiler/rose/blob/master/src/midend/programTransformation/ompLowering>
- The OpenMP runtime interface is defined in libxomp.h and xomp.c under the same ompLowering directory mentioned above

CUDA support

ROSE has an experimental connection to EDG 4.0, which helps us support CUDA.

To enable parsing CUDA codes, please use the following configuration options:

```
--enable-edg-version=4.0 --enable-cuda --enable-edg-cuda
```

Chapter 16 of ROSE User Manual has more details about this.

Abstract Syntax Tree (Intermediate Representation)

The main intermediate representation of ROSE is its abstract syntax tree (AST).

visualization of AST

We provide ROSE_INSTALLATION_TREE/bin/dotGeneratorWholeASTGraph to generate a dot graph of the detailed AST of input code.

To visualize the generated dot graph, you have to install

- ZGRViewer here: <http://zytm.sourceforge.net/zgrviewer.html#download>.
- Graphviz: <http://www.graphviz.org/Download.php>.

A complete example

```
# make sure the environment variables(PATH, LD_LIBRARY_PATH) for the
installed rose are correctly set
which dotGeneratorWholeASTGraph
~/workspace/masterClean/build64/install/bin/dotGeneratorWholeASTGraph

# run the dot graph generator
```

```
dotGeneratorWholeASTGraph -c ttt.c

#see it
which run.sh
~/64home/opt/zgrviewer-0.8.2/run.sh

run.sh ttt.c_WholeAST.dot
```

text output of AST

just call: `SgNode::unparseToString()`. You can call it from any `SgLocatedNode` within the AST to dump partial AST's text format.

preprocessing info.

In addition to nodes and edges, ROSE AST may have some extra attributes attached for preprocessing information like `#include`, `#if .. #else`. They are attached before, after, or within a nearby IAST node (only the one with source location information.)

An example translator will traverse the input code's AST and dump information about the found preprocessing information,

```
exampleTranslators/defaultTranslator/preprocessingInfoDumper -c
main.cxx
-----
Found an IR node with preprocessing Info attached:
(memory address: 0x2b7e1852c7d0 Sage type: SgFunctionDeclaration) in
file
/export/tmp.liao6/workspace/userSupport/main.cxx (line 3 column 1)
-----PreprocessingInfo #0 ----- :
classification = CpreprocessorIncludeDeclaration:
  String format = #include "all_headers.h"

relative position is = before
```

Please read more about this topic from ROSE tutorial: "Chapter 29 Handling Comments, Preprocessor Directives, And Adding Arbitrary Text to Generated Code" . You can download it from http://www.rosecompiler.org/ROSE_Tutorial/ROSE-Tutorial.pdf

Program Translation

With its high level intermediate representation, ROSE is suitable for building source-to-source translators. This is achieved by re-structuring the AST of the input source code, then unparsing the transformed AST to the output source code.

Expected behavior of a ROSE Translator

A translator built using ROSE is designed to act like a compiler (gcc, g++,gfortran ,etc depending on the input file types).

So users of the translator only need to change the build system for the input files to use the translator instead of the original compiler.

SageBuilder and SageInterface

The official guide for restructuring/constructing AST **highly recommends** using helper functions from SageBuilder and SageInterfaces namespaces to create AST pieces and moving them around. These helper functions try to be stable across low-level changes and be smart enough to transparently set many edges and maintain symbol tables.

Users who want to have lower level control may want to directly invoke the member functions of AST nodes and symbol tables to explicitly manipulate edges and symbols in the AST. But this process is very tedious and error-prone.

workflow for writing translators

Generic steps:

- prepare a simplest source file (a.c) as an example input of your translator
 - avoid including any system headers so you can visualize the whole AST
 - use ROSE_INSTALLATION_TREE/bin/dotGeneratorWholeASTGraph to generate a whole AST for a.c
- prepare another simplest source file (b.c) as an example output of your translator
 - again, avoid including any system headers
 - use ROSE_INSTALLATION_TREE/bin/dotGeneratorWholeASTGraph to generate a whole AST for b.c
- compare the two dot graphs side by side
- restructuring the source AST graph to be the AST graph you want to generate

Order to traverse AST

Naive pre-order traversal is not suitable for building a translator since the translator may change the nodes the traversal is expected to visit later on. Conceptually, this is essentially the same problem with C++ iterator invalidation.

To safely transform AST, It is recommended to use a reverse iterator of the statement list generated by a preorder traversal. This is different from a list generated from a post order traversal.

For example, assuming we have a subtree of : parent <child 1, child 2>,

- Pre order traversal will generate a list: parent, child 1, child2

- Post order traversal will generate a list: child 1, child2, parent.
- Reverse iterator of the pre order will give you : child2, child 1, and parent.
Transforming using this order is the safest based on our experiences.

example translators

split one complex statement into multiple simpler statements

- ROSE/projects/backstroke/ExtractFunctionArguments.C

Program Analysis

ROSE have implemented the following compiler analysis

- call graph analysis
- control flow graph
- data flow analysis: including liveness analysis, def-use analysis, etc.
- dependence analysis
- side effect analysis

control flow graph

ROSE provides several variants of control flow graphs

virtual control flow graph

The virtual control flow graph (vcfg) is dynamically generated on the fly when needed. So there is no mismatch between the ROSE AST and its corresponding control flow graph. The downside is that the same vcfg will be re-generated each time it is needed. This can be a potentially a performance bottleneck.

Facts

- documentation: virtual CFG is documented in **Chapter 19 Virtual CFG** of ROSE tutorial [pdf](#)
- source files:
 - src/frontend/SageIII/virtualCFG/virtualCFG.h
 - src/ROSETTA/Grammar/Statement.code // prototypes of member functions for located nodes, etc.
 - src/frontend/SageIII/virtualCFG/memberFunctions.C // implementation of virtual CFG related member functions for each AST node
 - this file will help the generation of buildTree/src/frontend/SageIII/Cxx_Grammar.h
- test directory: tests/CompileTests/virtualCFG_tests

- a dot graph generator: generator a dot graph for either the raw or interesting virtual CFG.
 - source: tests/CompileTests/virtualCFG_tests/generateVirtualCFG.C
 - Installed under rose_ins/bin

static control flow graph

Due to the performance concern of virtual control flow graph, we developed another static version which persistently exists in memory like a regular graph.

Facts:

- documentation: **19.7 Static CFG** of ROSE tutorial [pdf](#)
- test directory: rose/tests/CompileTests/staticCFG_tests

static and interprocedural CFGs

Facts:

- documentation: **19.8 Static, Interprocedural CFGs** of ROSE tutorial [pdf](#)
- test directory: rose/tests/CompileTests/staticCFG_tests

Virtual function analysis

Facts

- Original contributor: Faizur from UTSA, done in Summer 2011
- Code: at src/midend/programAnalysis/VirtualFunctionAnalysis.
- implemented with the techniques used in the following paper: "Interprocedural Pointer Alias Analysis - <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.42.2382>". The paper boils down the virtual function resolution to pointer aliasing problem. The paper employs flow sensitive inter procedural data flow analysis to solve aliasing problem, using compact representation graphs to represent the alias relations.
- some test files in the roseTests folder of the ROSE repository and he told me that the implementation supports function pointers as well as code which is written across different files (header files etc).
- documentation: Chapter 24 Dataflow Analysis based Virtual Function Analysis, of ROSE tutorial pdf

def-use analysis

If you want a def-use analysis, try this

http://www.rosecompiler.org/ROSE_HTML_Reference/classVariableRenaming.html

```
VariableRenaming v(project);  
v.run();  
v.getReachingDefsAtNode(...);
```

pointer analysis

<https://mailman.nersc.gov/pipermail/rose-public/2010-September/000390.html>

On 9/1/10 11:49 AM, Fredrik Kjolstad wrote: > Hi all, > > I am trying to use Rose as the analysis backend for a refactoring > engine and for one of the refactorings I am implementing I need > whole-program pointer analysis. Rose has an implementation of > steensgard's algorithm and I have some questions regarding how to use > this. > > I looked at the file steensgaardTest2.C to figure out how to invoke > this analysis and I am a bit perplexed: > > 1. The file SteensgaardPtrAnal.h that is included by the test is not > present in the include directory of my installed version of Rose. > Does this mean that the Steensgaard implementation is not a part of > the shipped compiler, or does it mean that I have to retrieve an > instance of it through some factory method whose static return type is > PtrAnal? I believe it is in the shipped compiler. And you're using the correct file to figure out how to use it. It should be in the installed include directory --- if it is not, it's probably something that needs to be fixed. But you can copy the include file from ROSE/src/midend/programAnalysis/pointerAnal/ as a temporary fix

> > 2. How do I initialize the alias analysis for a given SgProject? Is > this done through the overloaded ()?

The steensgaardTest2.C file shows how to set up everything to invoke the analysis. Right now you need to go over each function definition and invoke the analysis explicitly, as illustrated by the main function in the file. > > 3. Say I want to query whether two pointer variables alias and I have > SgNodes to their declarations. How do I get the AstNodePtr needed to > invoke the may_alias(AstInterface&, const AstNodePtr&, const > AstNodePtr&) function? Or maybe I should rather invoke the version of > may_alias that takes two strings (varnames)? > To convert a SgNode* x to AstNodePtr, wrap it inside an AstNodePtrImpl object, i.e., do AstNodePtrImpl(x), as illustrated inside the () operator of TestPtrAnal in steensgaardTest2.C.

> 4. How do I query whether two parameters alias? > The PtrAnal class has the following interface method

```
may_alias(AstInterface& fa, const AstNodePtr& r1, const AstNodePtr&
```

r2); It is implemented in SteensgaardPtrAnal class, which inherit PtrAnal class. To build AstInterface and AstNodePtr, you simply need to wrap SgNode* with some wrapper classes, illustrated by steensgaardTest2.C

-Qing Yi

```

void func(void) {
int* pointer;
int* aliasPointer;

pointer = malloc(sizeof(int));
aliasPointer = pointer;
*aliasPointer = 42;

printf("%d\n", *pointer);
}

```

The SteensgaardPtrAnal::output function returns:

```

c:(sizeof(int )) LOC1=>LOC2
c:42 LOC3=>LOC4
v:func LOC5=>LOC6 (inparams: ) ->(outparams: LOC7)
v:func-0 LOC8=>LOC7
v:func-2-1 LOC9=>LOC10
v:func-2-3 LOC11=>LOC12 (pending LOC10 LOC13=>LOC14 =>LOC4 )
v:func-2-4 LOC15=>LOC16 =>LOC17
v:func-2-5 LOC18=>LOC14 =>LOC4
v:func-2-aliasPointer LOC19=>LOC14 =>LOC4
v:func-2-pointer LOC20=>LOC13 =>LOC14 =>LOC4
v:malloc LOC21=>LOC22 (inparams: LOC2) ->(outparams: LOC12)
v:printf LOC23=>LOC24 (inparams: LOC16=>LOC17 LOC14=>LOC4 ) -
>(outparams:
LOC25)

```

SSA

ROSE has implemented an SSA form. Some discussions on the mailing list: [link](#).

Rice branch has an implementation of array SSA. We are waiting for their commits to be pushed into Jenkins. --[Liao](#) ([discuss](#) • [contribs](#)) 18:17, 19 June 2012 (UTC)

Generic dataflow framework

As the ROSE project goes on, we have collected quite some versions of dataflow analysis. It is painful to maintain and use them as they

- duplicate the iterative fixed-point algorithm
- scatter in different directories and
- use different representations for results.

An ongoing effort is to consolidate all dataflow analysis work within a single framework.

Quick facts

- original author: Greg Bronevetsky
- code reviewer: Chunhua Liao
- Documentation:

- source codes: files under ./src/midend/programAnalysis/genericDataflow
- tests: tests/roseTests/programAnalysisTests/generalDataFlowAnalysisTests
- currently implemented analysis
 - dominator analysis: dominatorAnalysis.h dominatorAnalysis.C
 - liveness variable analysis, or liveness analysis: liveDeadVarAnalysis.h liveDeadVarAnalysis.C
 - constant propagation: constantPropagation.h constantPropagation.C:
TODO need to move the files into src/ from /tests

Dependence analysis

The interface for dependence graph could be found in DependencyGraph.h. The underlying representation is n DepGraph.h. BGL is required to access the graph.

[Here](#) are 6 examples attached with this email. In deptest.C, there are also some macros to enable more accurate analysis.

If USE_IVS is defined, the induction variable substitution will be performed. if USE_FUNCTION is defined, the dependency could take a user-specified function side-effect interface. Otherwise, if non of them are defined, it will perform a normal dependence analysis and build the graph.

Program Optimizations

ROSE provides the following program optimizations and transformations:

- loop transformation, including loop fusion, fission, unrolling, blocking, loop interchange, etc.
- inlining
- outlining
- constant folding
- partial redundancy elimination

Developer's Guide

We briefly describe the workflow of ROSE developers.

Basic skills for ROSE developers

These are some basic skills that you, as a ROSE developer, should have or acquire as you work on your project(s):

- Linux shell: bash is the default shell for ROSE. know common commands (grep, find, ...) and basic scripting (bash, ...)
- C++ programming - be conscious of applying consistent coding-style conventions and writing code that will be maintainable when you leave -
- Debugging: GDB will be indispensable to make sure your code works as expected
- Git - source code management. Get familiar with the basics of Git: <http://git-scm.com/>
- Build systems - GNU Autotools (autoconf, automake), GNU Make, GNU libtool, and CMake (primarily so you won't break our existing Windows port)
- LaTeX - Document your work in ROSE/docs!! -
- Be familiar with ROSE documents (tutorials, installation, and developer guides): <http://rosecompiler.org/documents.html>. This also includes the project's Doxygen documentation.
- Compiler - ROSE is a compiler project after all. Take some compiler courses
 - Read online free course materials related to compilers
 - Keep learning topics related to your work

Expected Deliverables for Interns

Projects should result in some concrete deliverables, including but not limited to

- source codes: must pass code review and Jenkins so the codes can be merged into rose's master branch. Unfortunately, we have no resources to fix/merge dangling branches after students leave.
- bug fixes: fix issues/bugs assigned to you. The goal is to close these issues after your work.
- documentation: write new chapters, improve existing ones, mostly in LaTeX format.
- presentation: give a talk to share what you have learn and/or what you have done. It is a group-learning process. We will also use some of your slides to report to our sponsors.
- publication: try to publish a paper during your stay with us.
- weekly status reports: please update your redmine project's tasks to reflect what you have done each week

Of course, not every project should produce all of the deliverables.

Don't use branches, use separate repositories

A developer often works on several tasks, a project, a bug fix, a documentation task, etc.

Using branches can cause interferences among your tasks.

It is recommended to have an isolated local git repository for each task you are working on. So the progress/status of one task won't interfere with other tasks.

Summary of development tools

We leverage a range of software tools to facilitate the workflow of our project. Here is a brief list and the corresponding contact persons (ideally at least two to avoid "single-point failure" unless security rules say otherwise):

- rosecompiler.org: official website, admin1: Justin, admin2: Liao
- rose-public mailing list: admin 1: Justin, admin2: Liao
- Jenkins (<http://hudson-rose-30:8080/>) : continuous integration server, admin 1: Justin, admin2:???
- Redmine (<http://hudson-rose-30:3000/>) : project management, admin1: Justin, admin2:??
- github public: rose-compiler account, admin1: Justin, admin2: ???
- github enterprise (<http://github.llnl.gov/>) : rose-compiler organization account, admin1: Justin, admin2: ???

code review

see [Code Review](#) for details

issue tracking

We actually have more issue trackers than we would like to have right now. This is caused by the different purposes of issue tracking and the balancing between transparency and privacy/security. Maintainability is also a big factor since internal issue trackers may have quite some down-time compared to commercial websites.

- SciDAC outreach center:
 - external issue tracker: to collect user-submitted issues (ROSE bugs and feature requests). A nice thing about this tracker is that it allows submitting issues anonymously, very handy for busy and/or shy users. We should try to keep this quick channel open until we find alternative
 - private issue tracker: internal-submitted issues, which are not suitable for public view for any reasons. All issues should be moved to other internal issue tracker we decide to use.
- Internal redmine: the nice thing here is that it is a dedicated and intuitive project management web application. Tasks (like papers, presentations) can be not related to source code at all. Also important is that it is internal. We can have very secure internal project information. The downside is that it is hard to involve external collaborators.
 - each research project (and summer student project) has its own issue trackers
- github.com: potentially handy to connect with issue trackers of github. But not full-featured as a project management software. Everything has to be tied to a git

- repo. It is not a general assumption. Another downside is that it is public, not suitable for our internal projects.
- rosecompiler/rose 's issue tracking. A very good way to involving collaborators without remote access to LLNL
 - rosecompiler/rose-docs 's issue tracking. Being phased out.
 - github enterprise within LLNL: having benefits of being internal. But still not a general project management software.
 - We only use it for code review. Issue tracking there is turned off.

<https://github.com/blog/831-issues-2-0-the-next-generation>

<http://stackoverflow.com/questions/8888675/in-github-issue-tracker-can-non-admin-users-assign-users-and-labels>

<http://programmers.stackexchange.com/questions/129714/how-to-manage-github-issues-for-priority-etc>

How-tos

Add a new project

[How to add a new project directory](#)

sys admin's notes

The Apache Tomcat servlet is run locally on `*@hudson-rose-30@*`. Startup

- `@hudson-rose-30 $ /export/tmp.hudson-rose/applications/apache/tomcat/bin/startup.sh@`

On `@reboot`

```
@hudson-rose-30 $ crontab -e@  
<pre>@reboot /export/tmp.hudson-rose/applications/apache/tomcat/bin/startup.sh
```

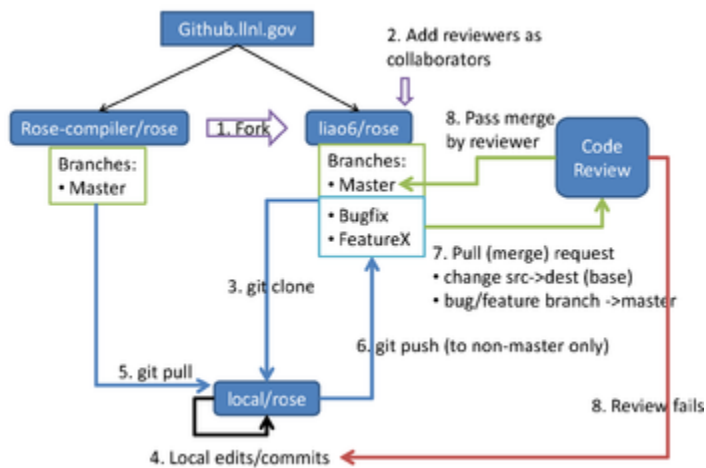
`@startup.sh@` calls `@catalina.sh start@` which sources our custom `@jenkins-env.sh@`.

Shutdown

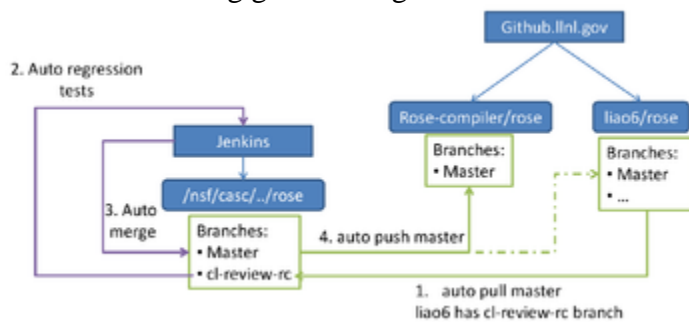
- `@hudson-rose-30 $ /export/tmp.hudson-rose/applications/apache/tomcat/bin/shutdown.sh@`

Backup: The local installation is synced by Jenkins to the `@NFS@` at `@/home/hudson-rose/Applications/apache/tomcat/@`.

Code Review Process



Code review using github.lnl.gov



Connection between github and Jenkins

Motivation

Without code review, developers have:

- added files into wrong directories, with improper names
- committed hundreds of reformatted files
- disabled tests to subvert our stringent Jenkins CI regression tests
- re-invented the wheel by implementing features that already exist

Goals

Our primary goals for code reviewing ROSE are to:

- share knowledge about the code: coder + reviewer will know the code, instead of just the coder

- group-study: learn through studying other peoples' code
- enforce policies for consistent usability and maintainability of ROSE code
- avoid reinventing the wheel and eliminating unnecessary redundancy
- safe-guarding the code: disallowing subversive attempts to disable or remove regression tests

Software

We are currently testing [Github Enterprise](#) and looking into the possibility of leveraging [Redmine](#) for internal code review.

In the past, we have looked at [Google's Gerrit code review system](#).

Github

Releases: <https://enterprise.github.com/releases>

Support: <https://support.enterprise.github.com>

Developer Checklist

Read these tips and guidelines before sending a request for code review.

Coding Standards

Your code should be written in a way that makes it easily maintainable and reviewable:

- write easy to understand code; avoid using exotic techniques which nobody can easily understand.
- add sufficient documentation (source-code comments, README, etc.) to aid the understandability of your code, your documentation should cover
 - why do you do this (motivation)
 - how do you do it (design and/or algorithm)
 - where are the associated tests (works as expected)
- before submission of your code for review, make sure
 - you have merged with the latest central repository's master branch without conflicts
 - your working copy can pass local tests via: make, make check, and make distcheck
 - you have fixed all compiler warnings of your code whenever possible
- submit a logical unit of work (one or more commits); something coherent like a bug fix, an improvement of documentation, an intermediate stage for reaching a big new feature.
- balance code submissions with a good ratio of [lines of code] and [complexity of code]. A good balance needs to be achieved to make the reviewer's life easier.

- the time needed to review your code should not exceed 1 hour

Workflow

Steps for initializing code review:

One time setup

- log into <http://github.llnl.gov> using your OUN and PAC
- fork your own rose repo from rose-compiler/rose
 - Goto <http://github.llnl.gov/rose-compiler/rose>, click Fork on the right upper corner of the webpage
- add potential reviewers (liao6, too1, vanderbrugge1, aananthakris1, vanka1) into your forked repo as collaborators, so they can review and merge your branches into your master branch later on
- add hudson-rose as your collaborator also so it can automatically push latest commits to your master branch to keep it up-to-date
 - Goto youraccount/rose of github.llnl.gov, then click Admin -> Collaborators
- create your public-private key pair using ssh-keygen, and copy your public key into your profile with github.llnl.gov
 - this is necessary since only ssh is supported by github.llnl.gov for now. https is not yet supported.

Daily work process

- have a local git repo to do your work and submit local commits, you have two choices:
 - clone it from /nfs/casc/overture/rose/rose.git as we usually do before
 - clone your fork on github.llnl.gov to a local repo: use the ssh URL option for now since the https option won't work.
 - don't use branches, use separated git repositories for each of your tasks. So status/progress of one task won't interfere with other tasks.
- When ready to push your commits, synchronize with the latest rose-compiler/master to resolve merge conflicts, etc.
 - type: git pull origin master # this should always work since master branches on github.llnl.gov are automatically kept up-to-date
 - make sure your local changes can pass 1)make -j8, 2)make check -j8, and 3)make distcheck -j8
- push your commits to your fork's non-master branch, like bugfix-rc , featurex-rc. You have total freedom in creating any branches in your forked repo, with any names you like

```
# If your local repository was cloned from
/nfs/casc/overture/ROSE/rose.git.
```

```
# There is no need to discard it. You can just add the github.llnl's
repo as an additional remote repository and push things there:
```

```
git remote add github-llnl-youraccount-rose
http://github.llnl.gov/youraccount/rose.git
git push github-llnl-youraccount-rose HEAD:refs/heads/bugfix-rc
```

- add a pull(merge) request to merge bugfix-rc into your own fork's master,
 - please note that the default pull request will use rose-compiler/rose's master as the base branch (destination of the merge). Please change it to be your own fork's master branch instead.
 - Also make sure the source (head) branch of the pull (merge) request is the one you want (bugfix-rc in this example)
- notify a reviewer that you have a pull request (requesting to merge your bugfix-rc into your master branch)
 - You can assign the pull request to the reviewer so an email notification will be automatically sent to the reviewer
 - Or you can add discussion for the pull request using @revieweraccount
 - Or you can just email the reviewer
- waiting for reviewer's feedback:
 - if passes, reviewer should have merged your bugfix-rc into your master. Jenkins will automatically poll your master and do the testing/merging
 - if reviewer wants additional changes such as better naming, better places to put files, more source comments, accompanying regression tests, etc. Just repeat the process: do local edits, local commits, push to your remote branch, send merge request again
 - A third possible outcome is that reviewers may accept the commits. But some additional tasks are needed in the future to improve the code.

Reviewer Checklist

What to look as a code reviewers?

- allocate up to 1 hour each time to review 500 to 1000 lines of code: longer time may not pay off due to the attention span limit of human brains
- directory paths and file names: are files in the intuitive paths or conform to our convention? are the names readable?
 - source codes, test input, documentation files are added into the right directories
- clarify of the code: can somebody who did not write the code easily understand what the code does?
 - the reason/motivation for writing the code
 - name convention: variable, function, class names should be intuitive
 - source comments: sufficient explanations for what each function, class does, what is the algorithm used, what is the paper/book chapter the implementation is based on.
- no duplication or reinvent of the wheel: similar code already exists or can be extended
- refactored: can part of the code be refactored to be reusable by others?

- no big functions: a function with hundreds line of code
- make check rules are associated with each new feature to ensure the new feature will be tested and verified for expected behaviors
- No turning off/relaxing other make check rules to make developers' commits pass Jenkins
- make a decision for the code review
 - pass. The code does what it is supposed to do with clear documentation and test cases. Merge the pull request and close the review.
 - fail. Additional changes are needed, such as better naming, better places to put files, more source comments, accompanying regression tests, etc. Notify the developers the issues and ask for a new push with suggested improvements.
 - pass but with future tasks. The commits are accepted. But some additional tasks are needed in the future to improve the code. They can be put into a separate push later on.

what to avoid

- Judging code by whether it's what the reviewer would have written
 - Given a problem, there are usually a dozen different ways to solve it. And given a solution, there's a million ways to render it as code.
- degenerating into nitpicks:
 - perfectionism may hurt the progress. we should allow some non-critical improvements to be done in the next version/commits.
- feel obligated to say something critical: it is perfectly fine to say "looks good, pass"
- delay in review: we should not rush it but we should keep in mind that somebody is waiting for the review to be done to move forward

criticism

Code reviews often degenerate into nitpicks. Brainstorming and design reviews to be more productive.

- I think this makes sense, the early we catch the problems, the better. Design happens earlier. Design should be reviewed. The same idea applies to requirement analysis also. --[Liao](#) ([discuss](#) • [contribs](#)) 18:18, 22 June 2012 (UTC)

references

- <http://www.possibility.com/wiki/index.php?title=CodeReviews>
- <http://scientopia.org/blogs/goodmath/2011/07/06/things-everyone-should-do-code-review/>
- <http://stackoverflow.com/questions/3730527/workflow-for-github-based-code-review>

- <http://stackoverflow.com/questions/4262693/what-to-look-for-in-a-code-review>
- LLNL Internal URL: <http://github.llnl.gov/>

Testing

ROSE uses [Jenkins](#) to implement a contiguous integration software development process. It leverages a range of software packages to test its correctness, robustness, and performance. The software used by the ROSE's Jenkins include:

- SPEC CPU 2006 benchmark: a subset is supported for now
- SPEC OMP benchmark: a subset is supported for now
- [NAS parallel benchmark](#): developed by NASA Ames Research Center. Both C (customize version) and OpenMP versions are used
- Plum Hall C and C++ Validation Test Suites: a subset is supported for now

Workflow

Motivation and Goals

Quality comes from a good process.

The goal is to have a streamlined, simplified, and automated workflow involving both users and developers to

- improve the quality of ROSE: source codes and documentations
- improve our productivity: optimize and simplify our daily work process so we can do more quality work using less time and other resources

Review of workflow changes

Major workflow improvements and changes should be thoroughly tested and reviewed by staff members before deployment since they may have profound impact on the project

review criteria:

- improve our productivity: optimize our workflow to allow us to do more quality and use less time and other resources
 - address what is slowing us down or distracting us
 - Simplify daily life (less hoops, the better). Side-by-side comparison of how many hoops we eliminated/automated by the workflow improvements
 - It is counterproductive to improve workflow by adding more hoops/steps/clicks into daily work.
- improve quality of work incrementally:

- accepting incremental improvements is more realistic than asking for perfection for the first attempt.
- workflow should allow quick new contributions and fast revision of existing contributions
- automation: new steps of workflow should be automated as much as possible.
- backwards compatibility:
 - preserve existing work, not create anything from scratch
 - Or having a way to converting existing code/documents into the new form
- preference of major contributors: whoever contributes the most should has a little bit more weight to say

Current workflow

Requirement Analysis

- wikibook:
 - draft big pictures and milestones
- mailing list: interaction with users, feel users' need

Design

- wikibook: quick draft design documents and provoke discussion
- powerpoint slides: more formal communication about what is the design

Implementation

- redmine: create projects based on milestones and user input, create and track tasks
- github:
 - internal:code review only,
 - external: hosting code, issue tracking
 - "rosebot" to automate Github workflow: preliminary testing, policies (git-hooks), automatically add reviewers, etc.

Testing

- Jenkins: continuous integration of new features, bugfixes

Documentation

- See more at [ROSE Compiler Framework/Documentation](#)

Publicity

- website: content management system hooked up with all other components

Frequently Asked Questions (FAQ)

We collect a list of frequently asked questions about ROSE, mostly from the rose-public mailing list [link](#)

How to search rose-public mailinglist for previously asked questions?

google.com supports search things within the scope of a URL. For example, if you have a problem with a keyword MY PROBLEM, you can try to search the mailing list by using the following keyword in google.com:

```
"MY PROBLEM site:https://mailman.nersc.gov/pipermail/rose-public/"
```

Compilation

How to speedup compiling ROSE?

Question It takes hours to compile ROSE, how can I speed up this process?

Answer:

- if you have multi-core processors, try to use `make -j4` (make by using four processes).
- also try to only build `librose.so` under `src/` by typing `make -C src/ -j4`
- Just to show some extreme case: using an old ThinkPad T60 laptop, Intel(R) Core(TM)2 CPU T7200 @ 2.00GHz, 3G Memory, it took 95 minutes for "`make -C src/ -j4`" to finish.

Can ROSE accept incomplete code?

<https://mailman.nersc.gov/pipermail/rose-public/2011-July/001015.html>

ROSE does not handle incomplete code. Though this might be possible in the future. It would be language dependent and likely depend heavily on some of the language specific tools that we use internally. This is however, not really a priority for our work. If you want to for example demonstrate how some of the internal tools we are using or alternative tools that we could use might handle incomplete code, this might be interesting and we could discuss it.

For example, we are not presently using Clang, but if it handled incomplete code that might be interesting for the future. I recall that some of the latest EDG work might handle some incomplete code, and if that is true then that might be interesting as well. I have not attempted to handle incomplete code with OFP, so I am not sure how well that could be

expected to work. Similarly, I don't know what the incomplete code handling capabilities of ECJ Java support is either. If you know any of these questions we could discuss this further.

I have some doubts about how much meaningful information can come from incomplete code analysis and so that would worry me a bit. I expect it is very language dependent and there would be likely some constraints on the incomplete code. So understanding the subject better would be an additional requirement for me.

Can ROSE analyze Linux Kernel sources?

<https://mailman.nersc.gov/pipermail/rose-public/2011-April/000856.html>

Question: I'm trying to analyze the Linux kernel. I was not sure of the size of the code-base that can be handled by ROSE, and could not find references as to whether it has been tried on the Linux kernel source. As of now I'm trying to run the identity translator on the source, and would like to know if it can be done using ROSE, and if it has been successfully tested before.

Short answer: Not for now

Long answer: We are using EDG 3.3 internally by default and this version of EDG does not handle the GNU specific register modifiers used in the `asm()` statements of the Linux Kernel code. There might be other problems, but that was at least the one that we noticed in previous work on this some time ago. But we are working on upgrading the EDG frontend to be a more recent version 4.4.

Can ROSE compile C++ Boost library?

<https://mailman.nersc.gov/pipermail/rose-public/2010-November/000544.html>

not yet.

I know of a few cases where ROSE can't handle parts of Boost. In each case it is an EDG problem where we are using an older version of EDG. We are trying to upgrade to a newer version of EDG (4.x), but that version's use within ROSE does not include enough C++ support, so it is not ready. The C support is internally tested, but we need more time to work on this.

AST

How to find XYZ in AST?

The usually steps to retrieve information from AST are:

- prepare a simplest (preferably 5-10 lines only), compilable sample code with the code feature (array[i+1]), avoid including any headers (#include file.h) to keep the code small.
 - Please note: don't include any headers in the sample code. A header (#include <stdio.h> for example) can bring in thousands of nodes into AST.
- use dotGeneratorWholeASTGraph to generate a detailed AST dot graph of the input code
- use zgrviewer-0.8.2's run.sh to visualize the dot graph
- visually/manually locate the information you want in the dot graph, understand what to look and where to look
- use code (AST member functions, traversal, SageInterface functions, etc) to retrieve the information.

How does the AST merge work?

tests that demonstrate the AST Merge are in the directory:

```
tests/CompileTests/mergeAST_tests
```

(run "make check" to see hundreds of tests go by).

How to filter out header files from AST traversals?

<https://mailman.nersc.gov/pipermail/rose-public/2010-April/000144.html> Question: I want to exclude functions in #include files from my analysis/transformations during my processing.

By default, AST traversal may visit all AST nodes, including the ones come from headers.

So AST processing classes provide three functions :

- T traverse (SgNode * node, ..): traverse full AST , nodes which represent code from include files
- T traverseInputFiles(SgProject* projectNode,..) traverse the subtree of AST which represents the files specified on the command line
- T traverseWithinFile(SgNode* node,..): only the nodes which represent code of the same file as the start node

Should SgIfStmt::get_true_body() return SgBasicBlock?

<https://mailman.nersc.gov/pipermail/rose-public/2011-April/000930.html>

Both true/false bodies were SgBasicBlock before.

Later, we decided to have more faithful representation of both blocked (with {...}) and single-statement (without { ..}) bodies. So they are SgStatement (SgBasicBlock is a subclass of SgStatement) now.

But it seems like the document has not been updated to be consistent with the change.

You have to check if the body is a block or a single statement in your code. Or you can use the following function to ensure all bodies must be SgBasicBlock.

```
//A wrapper of all ensureBasicBlockAs*() above to ensure the parent of s is a scope statement with list of statements as children, otherwise generate a SgBasicBlock in between.
```

```
SgLocatedNode * SageInterface::ensureBasicBlockAsParent (SgStatement *s)
```

How to handle #include "header.h", #if, #define etc. ?

It is called preprocessing info. within ROSE's AST. They are attached before, after, or within a nearby AST node (only the one with source location information.)

An example translator is provided to traverse the input code's AST and dump information about the found preprocessing information,

```
exampleTranslators/defaultTranslator/preprocessingInfoDumper -c
main.cxx
-----
Found an IR node with preprocessing Info attached:
(memory address: 0x2b7e1852c7d0 Sage type: SgFunctionDeclaration) in
file
/export/tmp.liao6/workspace/userSupport/main.cxx (line 3 column 1)
-----PreprocessingInfo #0 ----- :
classification = CpreprocessorIncludeDeclaration:
  String format = #include "all_headers.h"

relative position is = before
```

SgClassDeclaration::get_definition() returns NULL?

If you look at the whole AST graph carefully, you can find defining and non-defining declarations for the same class.

A symbol is usually associated with a non-defining declaration. A class definition is associated with a defining declaration.

You may want to get the defining declaration from the non-defining declaration before you try to grab the definition.

Translation

Can ROSE identityTranslator generate 100% identical output file?

<https://mailman.nersc.gov/pipermail/rose-public/2011-January/000604.html>

Questions: Rose identityTranslator performs some modifications, "automatically".

These modifications are:

- Expanding the assert macro.
- Adding extra brackets around constants of typedef types (e.g. `c=Typedef_Example(12);` is translated in the output to `c = Typedef_Example((12));`)
- Converting NULL to 0.

How can I avoid these modifications?

Answer: No.

There is no easy way to avoid these changes currently. Some of them are introduced by the cpp preprocessor. Others are introduced by the EDG front end ROSE uses. 100% faithful source-to-source translation may require significant changes to preprocessing directive handling and the EDG internals.

We have had some internal discussion to save raw token strings into AST and use them to get faithful unparsed code. But this effort is still at its initial stage as far as I know.

How to build a tool inserting function calls?

<https://mailman.nersc.gov/pipermail/rose-public/2010-July/000319.html>

Question: I am trying to build a tool which insert one or more function calls whenever in the source code there is a function belonging to a certain group (e.g. all functions beginning with `foo_*`). During the ast traversal, how can I find the right place, i.e., there is a function in ROSE that searches for a string pattern or something similar?

Answers:

- In Chapter 28 AST Construction of the ROSE tutorial, there are examples to instrument function calls into the AST using traversals or a queryTree. I would approach this by checking the node for the specific `SgFunctionDefinition` (or whatever you need) and then check the name of the node to find its location.
- You can

- use the AST query mechanism to find all functions and store them in a container. e.g `Rose_STL_Container<SgNode*> nodeList = NodeQuery::querySubTree(root_node, V_Sg????);`
- Then iterate the container to check each function to see if the function name matches what you want.
- use SageBuilder namespace's `buildFunctionCallStmt()` to create a function call statement.
- use SageInterface namespace's `insertStatement ()` to do the insertion.

How to copy/clone a function?

<https://mailman.nersc.gov/pipermail/rose-public/2011-April/000919.html>

We need to be more specific about the function you want to copy. Is it just a prototype function declaration (non-defining declaration in ROSE's term) or a function with a definition (defining declaration in ROSE's term)?

- Copying a non-defining function declaration can be achieved by using the following function instead:

```
// Build a prototype for an existing function declaration (defining or
nondefining is fine).
SgFunctionDeclaration* SageBuilder::buildNondefiningFunctionDeclaration
(const SgFunctionDeclaration *funcdecl, SgScopeStatement *scope=NULL)
```

Copying a defining function declaration is semantically a problem since it introduces redefinition of the same function. It is at least a hack to first introduce something wrong and later correct it. Here is an example translator to do the hack (copy a defining function, rename it, fix its symbol):

```
#include <rose.h>
#include <stdio.h>
using namespace SageInterface;

int main(int argc, char** argv)
{
    SgProject* project = frontend(argc, argv);
    AstTests::runAllTests(project);

    // Find a defining function named "bar" under project

    SgFunctionDeclaration* func=
findDeclarationStatement<SgFunctionDeclaration> (project, "bar", NULL,
true);
    ROSE_ASSERT (func != NULL);

    // Make a copy and set it to a new name
    SgFunctionDeclaration* func_copy =
isSgFunctionDeclaration(copyStatement (func));
    func_copy->set_name("bar_copy");
```

```

// Insert it to a scope
SgGlobal * glb = getFirstGlobalScope(project);
appendStatement (func_copy,glb);

#if 1 // fix up the missing symbol, this should be optional now since
SageInterface::appendStatement() should handle it transparently.
SgFunctionSymbol *func_symbol = glb->lookup_function_symbol
("bar_copy", func_copy->get_type());
if (func_symbol == NULL);
{
    func_symbol = new SgFunctionSymbol (func_copy);
    glb ->insert_symbol("bar_copy", func_symbol);
}
#endif
AstTests::runAllTests(project);
backend(project);
return 0;
}

```

Can I transform code within a header file?

<https://mailman.nersc.gov/pipermail/rose-public/2011-May/000971.html>

No. ROSE does not unparse AST from headers right now. A summer project tried to do this. But it did not finish.

<https://mailman.nersc.gov/pipermail/rose-public/2010-August/000344.html>

I guess ROSE does not support writing out changed headers for safety/practical reasons. A changed header has to be saved to another file since writing to the original header is very dangerous (imaging debugging a header translator which corrupts input headers). Then all other files/headers using the changed header have to be updated to use the new header file.

Also all files involved have to be writable by user's translators.

As a result, the current unparser skips subtrees of AST from headers by checking file flags (compiler_generated and/or output_in_code_generation etc.) stored in Sg_File_Info objects.

how to work with formal and actual arguments of functions?

<https://mailman.nersc.gov/pipermail/rose-public/2011-June/001008.html>

```

//Get the actual arguments
SgExprListExp* actualArguments = NULL;
if (isSgFunctionCallExp(callSite))
    actualArguments = isSgFunctionCallExp(callSite)->get_args();
else if (isSgConstructorInitializer(callSite))

```

```

        actualArguments = isSgConstructorInitializer(callSite)-
>get_args();
        ROSE_ASSERT(actualArguments != NULL);

        const SgExpressionPtrList& actualArgList =
actualArguments->get_expressions();

        //Get the formal arguments.
        SgInitializedNamePtrList formalArgList;
        if (calleeDef != NULL)
            formalArgList = calleeDef->get_declaration()->get_args();

        //The number of actual arguments can be less than the number of
formal arguments (with implicit arguments) or greater
        //than the number of formal arguments (with varargs)

```

how to clone a function?

<https://mailman.nersc.gov/pipermail/rose-public/2011-April/000919.html>

Here is an example translator to do the hack (copy a defining function, rename it, fix its symbol):

```

-----
#include <rose.h>
#include <stdio.h>
using namespace SageInterface;

int main(int argc, char** argv)
{
    SgProject* project = frontend(argc, argv);
    AstTests::runAllTests(project);

    // Find a defining function named "bar" under project
    SgFunctionDeclaration* func=
findDeclarationStatement<SgFunctionDeclaration> (project, "bar", NULL,
true);
    ROSE_ASSERT (func != NULL);

    // Make a copy and set it to a new name
    SgFunctionDeclaration* func_copy =
isSgFunctionDeclaration(copyStatement (func));
    func_copy->set_name("bar_copy");

    // Insert it to a scope
    SgGlobal * glb = getFirstGlobalScope(project);
    appendStatement (func_copy, glb);

    AstTests::runAllTests(project);
    backend(project);
    return 0;
}

```

Daily work

git clone returns error: SSL certificate problem?

Symptom:

```
git clone https://github.com/rose-compiler/rose.git
Cloning into rose...
error: SSL certificate problem, verify that the CA cert is OK. Details:
error:14090086:SSL routines:SSL3_GET_SERVER_CERTIFICATE:certificate
verify failed while accessing https://github.com/rose-
compiler/rose.git/info/refs

fatal: HTTP request failed
```

The reason may be that you are behind a firewall which tweaks the original SSL certification.

Solutions: Tell cURL to not check for SSL certificates:

```
#Solution 1: Environment variable (temporary)
$ env GIT_SSL_NO_VERIFY=true git pull

# Solution 2: git-config (permanent)
# set local configuration
$ git config --local http.sslVerify false

# Solution 2: set global configuration
$ git config --global http.sslVerify false
```

What is the best IDE for ROSE developers?

<https://mailman.nersc.gov/pipermail/rose-public/2010-April/000115.html>

There may not be a widely recognized best integrated development environment. But developers have reported that they are using

- vim
- emacs
- KDevelop
- Source Navigator
- Eclipse
- Netbeans

The thing is that ROSE is huge and has some ridiculously large generated source file (CxxGrammar.h and CxxGrammar.C are generated in the build tree for example). So many code browsers may have trouble in handling ROSE.

Portability

What is the status for supporting Windows?

<https://mailman.nersc.gov/pipermail/rose-public/2011-December/001349.html>

We have not finished the Windows work yet. IT is on our list of things to do. It was started and ROSE internally compiles using MS Visual Studio (using project files generated from the Cmake build that we maintain and test within our release process for ROSE) but does not pass our tests. So it is not ready. The distribution of the EDG binaries for Windows is another step that would come after that. We don't know at present when this will be done, it is important, but not a high priority for our DOE specific work, but important for other work. The effort required is something that we could discuss. If you want to call me that would be the best way to proceed. Send me email off of the main list and we can set that up.

<https://mailman.nersc.gov/pipermail/rose-public/2011-March/000798.html>

Under Windows ROSE uses CMake. This is a project that is currently under development. As of November 2010 we are able to compile and link the src directory. We are also able to run example programs that link against librose and execute the frontend and backend. However, this is an internal capability and not available externally yet since we don't distribute the Windows generated EDG binaries that would be required. Also the current support for Windows is still incomplete, ROSE does not yet pass its internal tests under Windows. }

Who is using ROSE

We are aware of the following ROSE users (people who write their own ROSE-based tools). They are the reason of the ROSE's existence. Feel free to add your name if you are using ROSE.

Universities

- University of California, San Diego, CUDA code generator [link](#)
- University of Utah, compiler-based parameterized code transformation for autotuning
- University of Oregon, performance tools
- University of Wyoming, OpenMP error checking

DOE national laboratories

- Argonne National Laboratory, performance modeling

TODO List

What is missing (so you can help if you want)

- Generate a pdf output of this book
 - printable, pdf version: http://en.wikibooks.org/wiki/Help:Print_versions
- how to download a backup of this wiki book?

Retrieved from

["http://en.wikibooks.org/w/index.php?title=ROSE_Compiler_Framework/Print_version&oldid=2373014"](http://en.wikibooks.org/w/index.php?title=ROSE_Compiler_Framework/Print_version&oldid=2373014)

[Category:](#)

- [ROSE Compiler Framework](#)

What do you think of this page? Reliability:

Completeness: Neutrality:

Presentation:

Re-review this revision

Quality: poor/unrated minimal average good

Comment:

Personal tools

- [Liao](#)
- [My discussion](#)
- [My preferences](#)
- [My watchlist](#)
- [My contributions](#)
- [Log out](#)

Namespaces

- [Book](#)
- [Discussion](#)

Variants

Views

- [Read](#)
- [Edit](#)
- [View history](#)
- [Unwatch](#)

Actions

- [Move](#)

Search

Navigation

- [Main Page](#)
- [Help](#)
- [Browse](#)
- [Cookbook](#)
- [Wikijunior](#)
- [Featured books](#)
- [Recent changes](#)
- [Donations](#)
- [Random book](#)

Community

- [Reading room](#)
- [Community portal](#)
- [Bulletin Board](#)
- [Help out!](#)
- [Policies and guidelines](#)
- [Contact us](#)

Toolbox

- [What links here](#)
- [Related changes](#)
- [Upload file](#)
- [Special pages](#)
- [Permanent link](#)
- [Cite this page](#)
- [Page rating](#)

Sister projects

- [Wikipedia](#)
- [Wikiversity](#)
- [Wiktionary](#)
- [Wikiquote](#)
- [Wikisource](#)
- [Wikinews](#)
- [Commons](#)

[Print/export](#)

- [Create a collection](#)
- [Download as PDF](#)
- [Printable version](#)

- This page was last modified on 4 July 2012, at 17:51.
- Text is available under the [Creative Commons Attribution-ShareAlike License](#); additional terms may apply. See [Terms of Use](#) for details.

- [Privacy policy](#)
- [About Wikibooks](#)
- [Disclaimers](#)
- [Mobile view](#)

