

MILEPOST GCC: machine learning based research compiler

Grigori Fursin
Cupertino Miranda
Olivier Temam
INRIA Saclay, France

Mircea Namolaru
Elad Yom-Tov
Ayal Zaks
Bilha Mendelson
IBM Haifa, Israel

Phil Barnard
Elton Ashton
ARC International
UK

Eric Courtois
Francois Bodin
CAPS Enterprise
France

Edwin Bonilla
John Thomson
Hugh Leather
Chris Williams
Michael O'Boyle

University of Edinburgh, UK

Contact: grigori.fursin@inria.fr

Motivation

Tuning hardwired compiler optimizations for rapidly evolving hardware makes porting an optimizing compiler for each new platform extremely challenging. Our radical approach is to develop a modular, extensible, self-optimizing compiler that automatically learns the best optimization heuristics based on the behavior of the platform. In this poster we describe MILEPOST[™] GCC, a machine-learning-based compiler that automatically adjusts its optimization heuristics to improve the execution time, code size, or compilation time of specific programs on different architectures. Our preliminary experimental results show that it is possible to considerably reduce execution time of the MiBench benchmark suite on a range of platforms entirely automatically.

* MILEPOST Project - Machine Learning for Embedded Program Optimization
<http://www.milepost.eu>

Challenges


Current state-of-the-art compilers often fail to deliver best performance due to:

- *hardwired optimization heuristics (cost models) for rapidly evolving hardware (often impossible to fine-tune programs externally)*
- *interaction between optimizations*
- *large irregular optimization spaces*
- *difficult to add new transformations to already tuned optimization heuristics*
- *inability to reuse optimization knowledge among different programs and architectures*
- *lack of run-time information and inability to adapt to varying program and system behavior at run-time with low overhead*

Need modular self-tuning compilers that can continuously and automatically learn how to optimize programs, and have an ability to make program adaptable at run-time for different behavior and constraints

Iterative compilation

Optimization spaces (set of all possible program transformations) are large, non-linear with many local minima



Finding a good solution may be long and non-trivial

matmul, 2 transformations, search space = 2000
swim, 3 transformations, search space = 10⁵²

Iterative compilation:

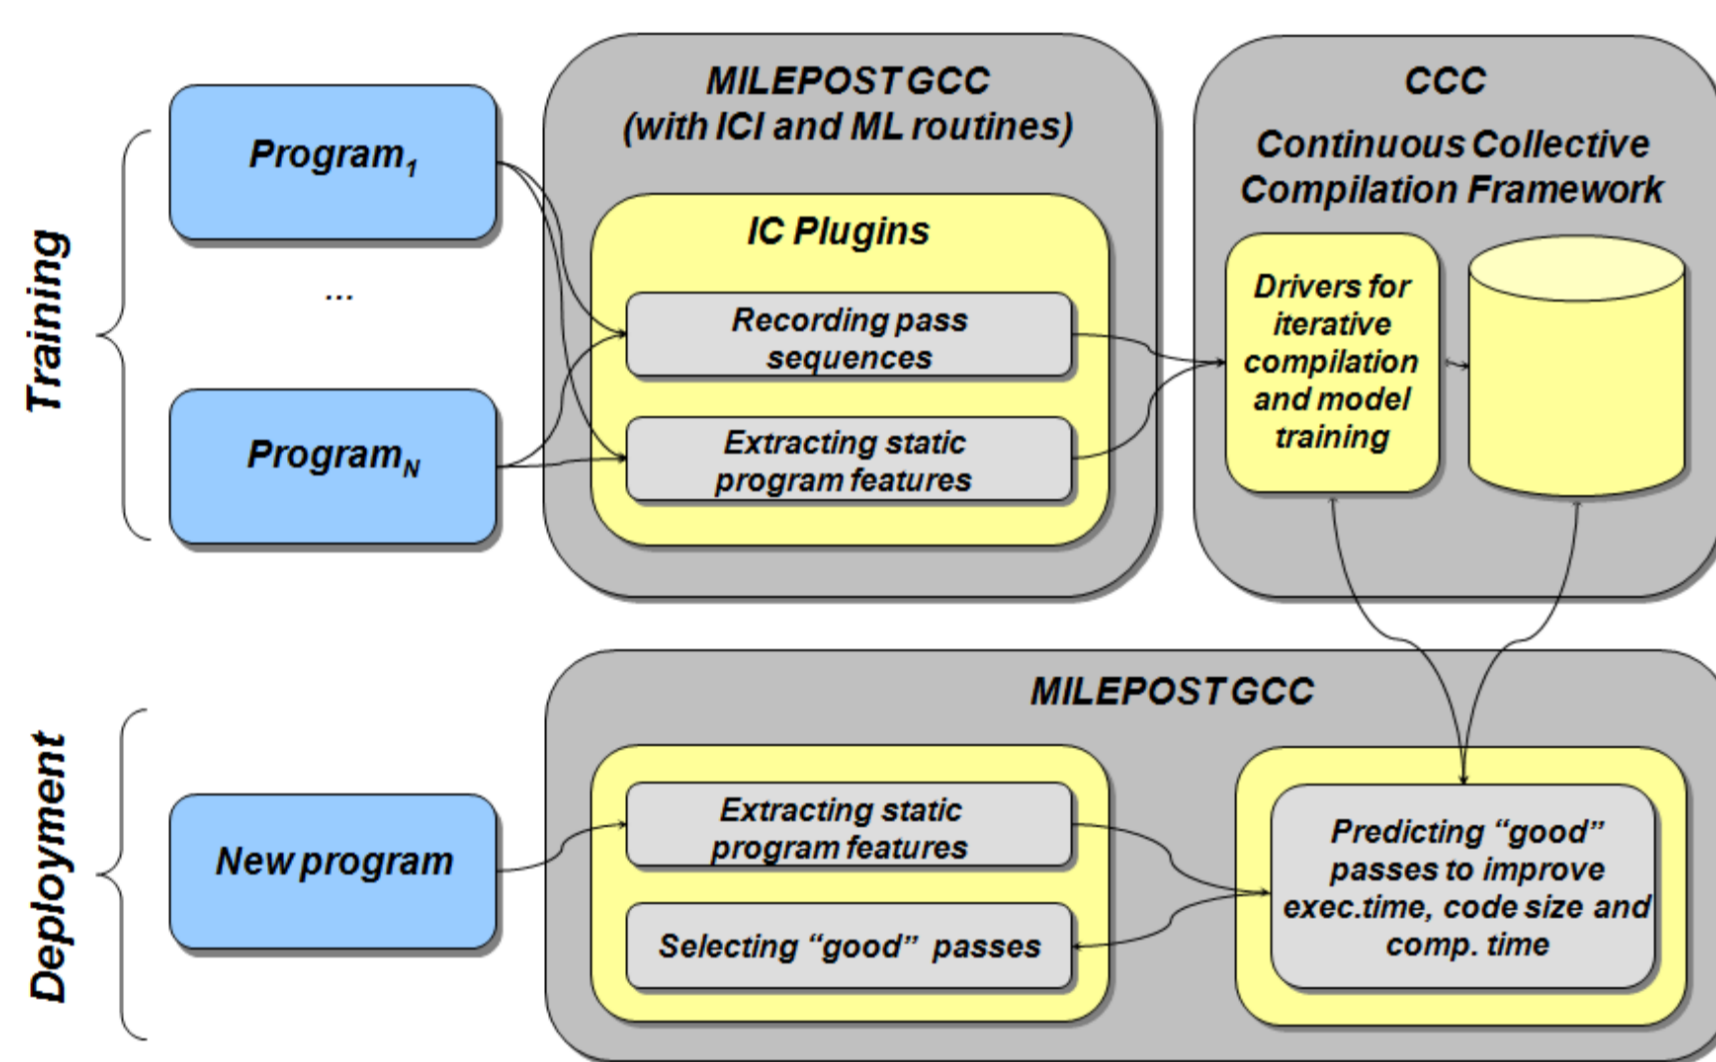
learn program behavior across executions

High potential (O'Boyle, Cooper since 1998), but:

- slow
- the same dataset is used
- no run-time adaptation
- no optimization knowledge reuse

Solving these problems is non-trivial

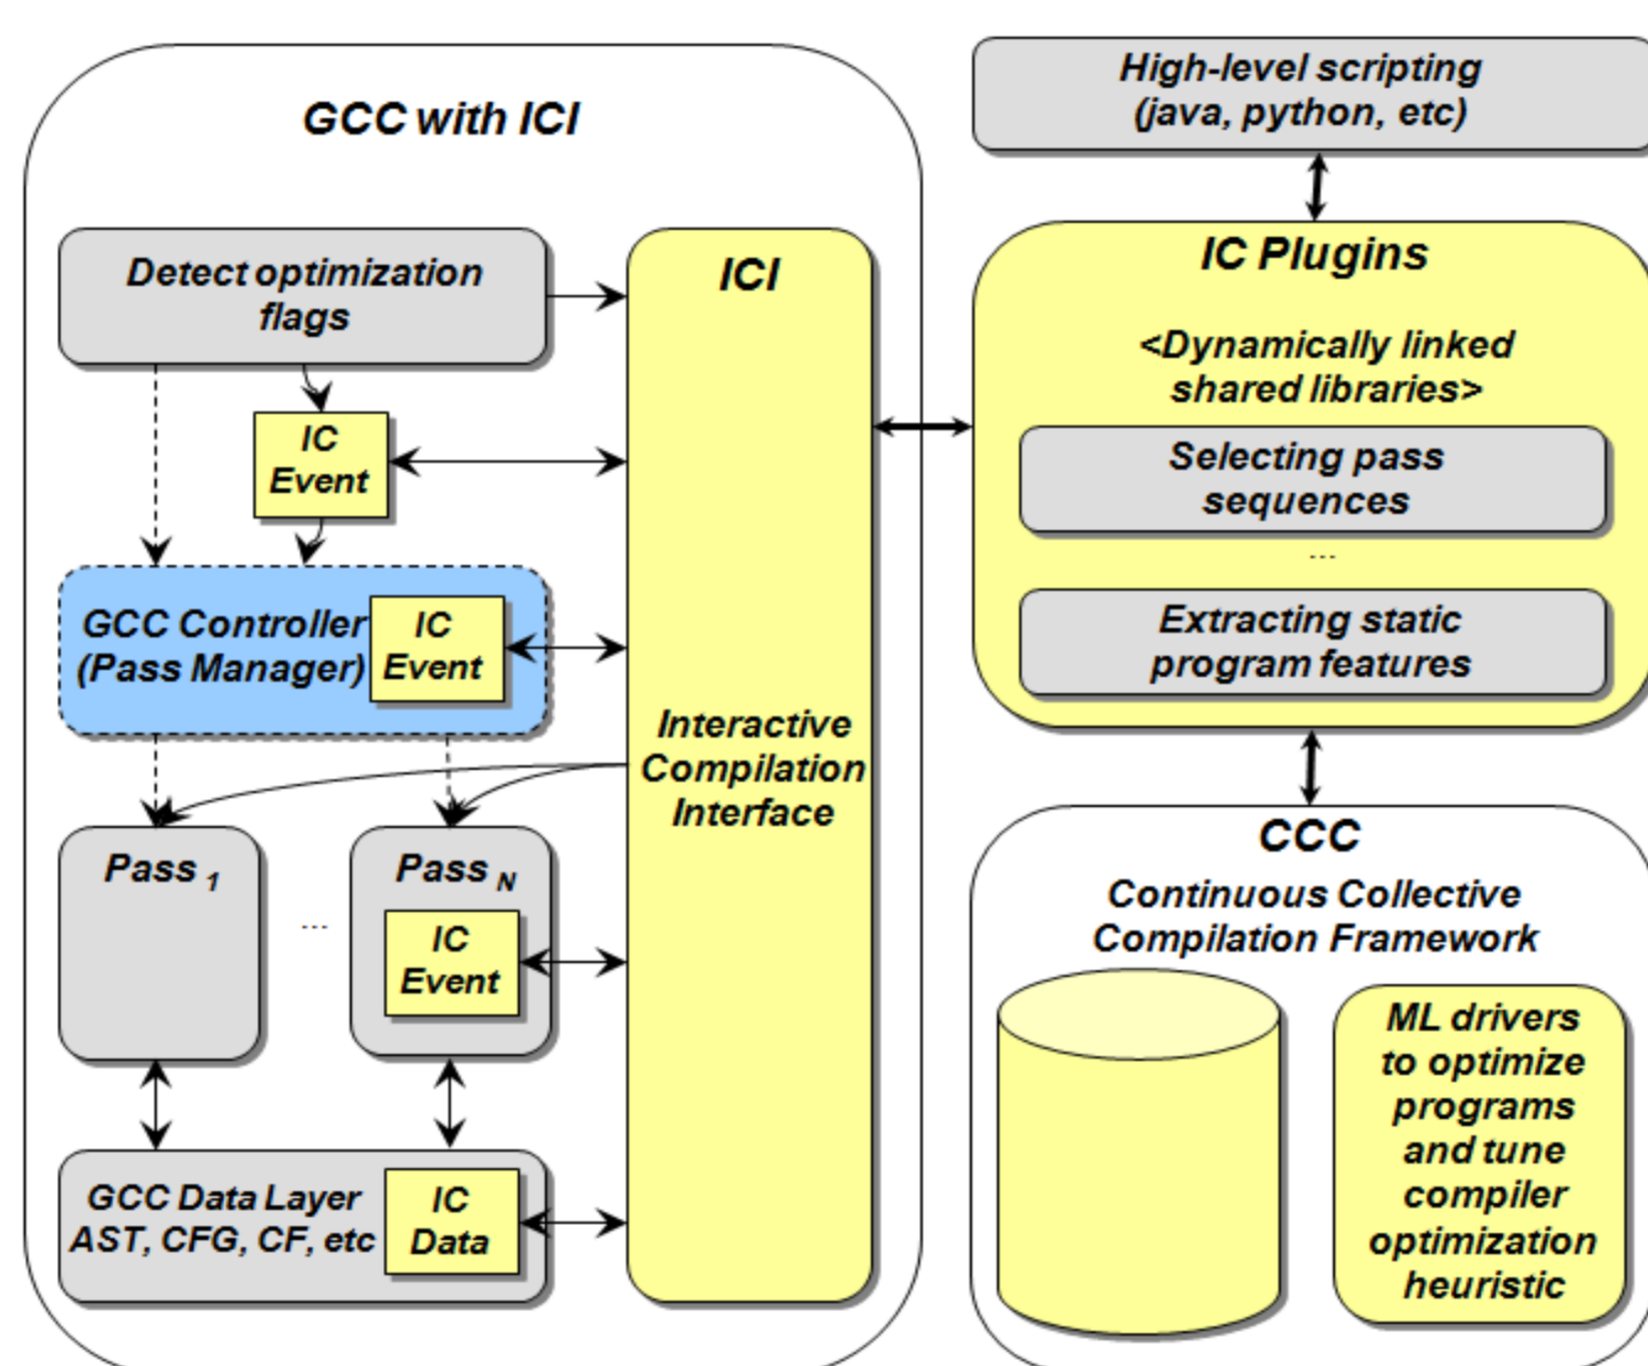
MILEPOST Framework



Training: Gathering information about the structure of programs and record how they behave when compiled under different optimization settings to build machine learning models.

Deployment: ML model is able to predict good optimization strategies for a given set of program features and is built as a plugin within MILEPOST GCC.

Interactive Compilation Interface



The ICI provides opportunities for external control and examination of the compiler. Optimization settings at a fine-grained level, beyond the capabilities of command line options or pragmas, can be managed through IC plugins.

GCC IC Plugins

```
file: ic-framework.c
static int
load_ici(char *dynlib_file)
{
    void *ICLib;
    bool error = 0;
    ICLib = dlopen(dynlib_file, RTLD_LAZY);
    error = !check_for_derror();
    ICI_start = (func) dlsym(ICLib, "start");
    ICI_stop = (func) dlsym(ICLib, "stop");
    error = !check_for_derror();
}

file: save-executed-passes.c (IC Plugin)
#include "ic-controller.h"
#include "gcc-ici-interface.h"
void executed_pass(void)
{
    char *pass_name;
    char *func_name;
    * to save original GCC pass order */
    ICI_start = (func) dlsym(ICLib, "start");
    ICI_stop = (func) dlsym(ICLib, "stop");
    error = !check_for_derror();
}

file: passes.c
bool
execute_one_pass(struct tree_opt_pass *pass)
{
    bool initializing_dump;
    unsigned int todo_after = 0;
    static bool gate_status;
    gate_status = (pass->gate == NULL) ? true :
    pass->gate;
    ICI_register_parameter("gate_status", &gate_status);
    ICI_call_event("avoid_gate");
    ICI_unregister_parameter("gate_status");
    if (!gate_status)
    return false;
    ICI_register_parameter("pass_name", (void *)
    (pass->name));
    ICI_call_event("pass_execution");
    ICI_unregister_parameter("pass_name");
}

char start(void)
{
    ICI_register_event("pass_execution",
    &executed_pass);
}
```

Modifications needed to enable GCC ICI:

ic-framework.c: GCC plugin (dynamic library) invocation

save-executed-passes.c: Plugin to monitor executed passes. It registers an event handler function executed_pass on an IC-Event called pass_execution.

passes.c: Modifications in GCC Controller (Pass Manager) to enable manipulation with optimization passes.

Optimization selection

Using Interactive Compilation Interface we can conduct research on optimization pass selection and reordering

```
fixupcfg_inil_datastructures,all_optimizations,...,retslot,copyrename,cop,...,dce,dm,phicprop,phiopt,alias,tail,
profile,ch,...,alias,copyrename,dm,phicprop,reassoc,dce,dse,alias,...,copyprop,lit,unswitch,...,ctrlroll,livopts,
loopdone,reassoc,vrp,dm,phicprop,cddce,dse,forprop,phiopt,tail,copyrename,uncprop,optimized,...,cse1,
gcse1,bypass,cse1,...,final_clean_state
```

Sequence of compiler passes for -O3

```
fixupcfg_inil_datastructures,all_optimizations,...,retslot,copyrename,dce,phiopt,alias,profile,ch,...,alias,reassoc,dce,alias,...,
ctrlroll,livopts,loopdone,reassoc,vrp,dm,phicprop,cddce,dse,forprop,phiopt,tail,copyrename,uncprop,optimized,...,cse1,
gcse1,combine,cse2,regmove,split1,mode-sw,life2,sched1,...,final_clean_state
```

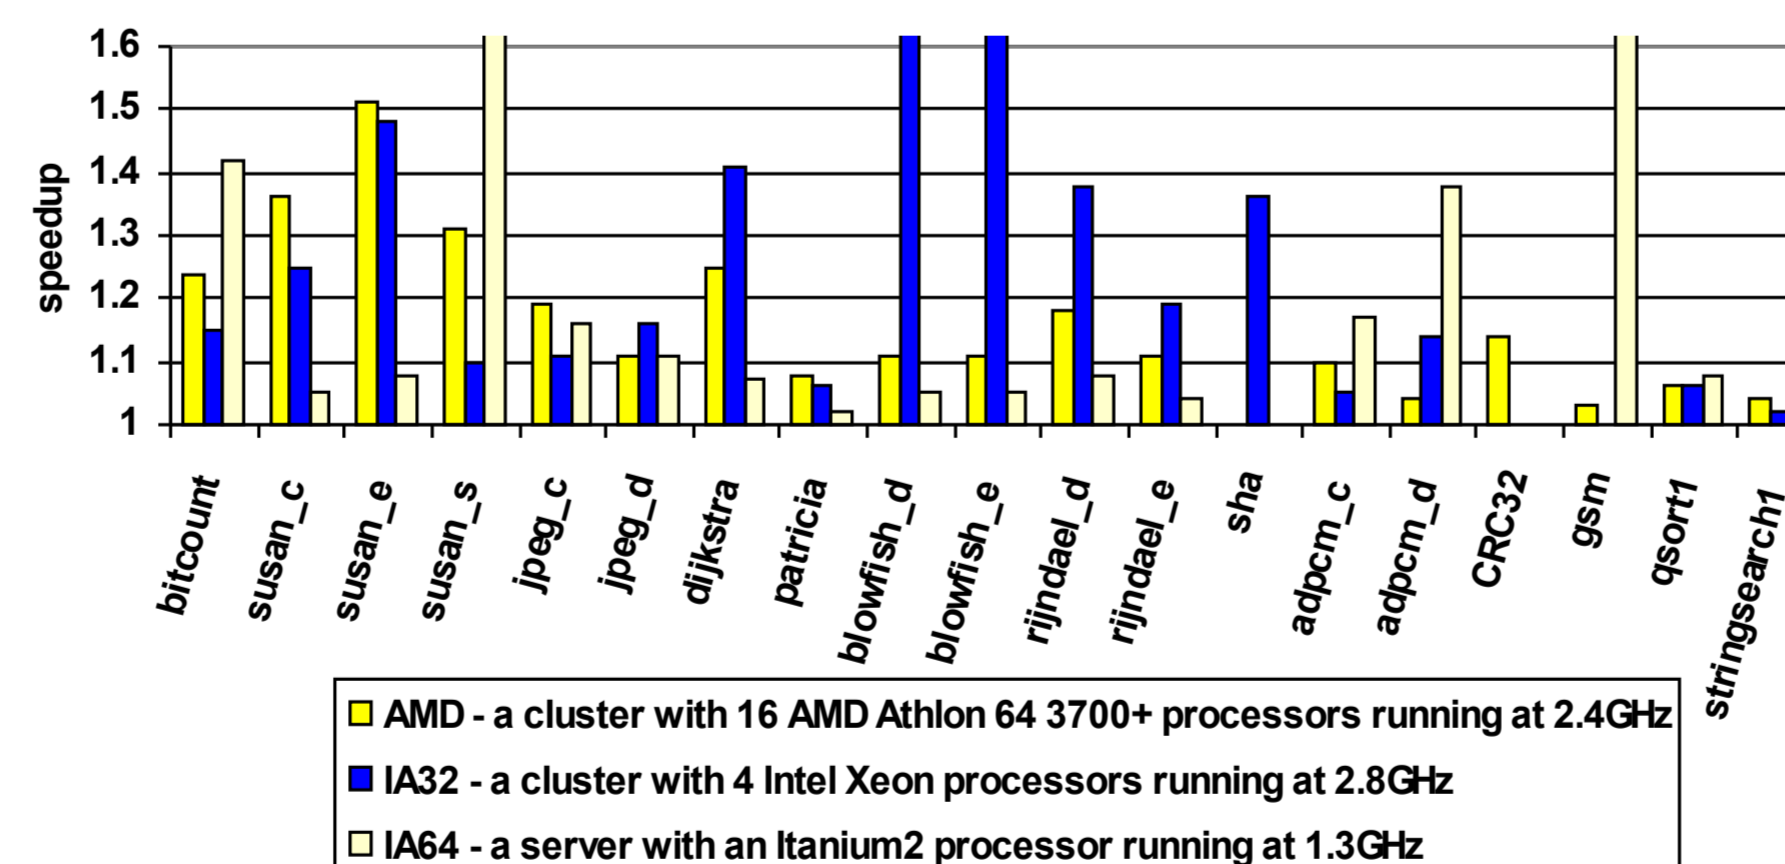
Sequence of GCC passes for the "good" set of compiler flags (ISSUE with UNROLLING, PEELING)

Feature extraction

We can now add new passes that are not included into default optimization heuristic but called through ICI and plugins:

- ft1** - Number of basic blocks in the method
- ...
- ft20** - Number of conditional branches in the method
- ft21** - Number of assignment instructions in the method
- ...
- ft25** - Average of number of instructions in basic blocks
- ...
- ft55** - Number of static/extern variables that are pointers in the method


Training



■ AMD - a cluster with 16 AMD Athlon 64 3700+ processors running at 2.4GHz
■ IA32 - a cluster with 4 Intel Xeon processors running at 2.8GHz
■ IA64 - a server with an Itanium2 processor running at 1.3GHz

We generate training data using Continuous Collective Compilation Framework. We select 500 random sequences of flags (or associated passes) either turned on or off. We can already achieve considerable speedups across all platforms, however it is very time-consuming process motivating the use of machine learning to automatically build specialized compilers and predict the best optimization flags or sequences of passes for different architectures.

Predicting



Once ML model is built, we can evaluate it by introducing a new program to a system and measuring how well the prediction performs. This graphs shows the speedups obtained on ARC725D after iterative compilation (500 iterations) and after 1 prediction. It demonstrates that except a few pathological cases using CCC Framework, MILEPOST GCC and Machine Learning Models we can improve original ARC GCC by around 11%.

Future work

- continue research on pass selection and reordering for reconfigurable processors and design space exploration
- improve and automate selection of static and dynamic (hardware counters) program features
- use ICI for adaptive libraries
- enable run-time adaptation and parallelization for static programs



More information

Project news: <http://www.milepost.eu>

GCC-ICI: <http://gcc-ici.sourceforge.net>

UNIDAPT: <http://unidapt.org>

Extended version (GCC Summit'08): <http://unidapt.org/papers/fmtpp2008.pdf>

