

A diy tutorial

Version 1.0

January 17, 2010

diy is a tool suite for testing shared memory models. We provide two tools, *litmus* (Part I) for running tests, and *diy proper* (Part II) for producing tests from concise specifications. The software is written in Objective Caml¹, and released as sources. The web site of *diy* is <http://diy.inria.fr/>, authors can be contacted at diy-devel@inria.fr

This software is released under the terms of the Lesser GNU Public License.

Contents

I	Running tests with <i>litmus</i>	3
1	A tour of <i>litmus</i>	3
1.1	A simple run	3
1.2	Cross compilation	4
1.3	Running several tests at once	5
2	Controlling test parameters	6
2.1	Architecture of tests	6
2.2	Controlling executable files	7
2.3	Controlling <i>litmus</i>	7
II	Generating tests	10
3	Preamble	10
3.1	Relaxation of Sequential Consistency	10
3.2	Introduction to relaxations	11
3.3	More relaxations	12
3.4	Summary of simple relaxations	15
3.4.1	Communication relaxations	15
3.4.2	Program order relaxations	15
3.4.3	Barrier relaxations	16
4	Testing relaxations with <i>diy</i>	16
4.1	Principle	16
4.2	Testing x86	17
5	Additional relaxations	18
5.1	Intra-processor dependencies	18
5.2	Composite relaxations and cumulativity	19

¹<http://caml.inria.fr/ocaml/>

6	Command usage	20
6.1	Usage of diyone	20
6.2	Usage of diy	20
6.3	Usage of readRelax	21

Part I

Running tests with litmus

Traditionally, a *litmus test* is a small parallel program designed to exercise the memory model of a parallel, shared-memory, computer. Given a litmus test in assembler (X86 or Power) `litmus` runs the test.

Using `litmus` thus requires a parallel machine, which must additionally feature `gcc` and the `pthread`s library. At the moment, `litmus` is a prototype and has numerous limitations (recognised instructions, limited porting). Nevertheless, `litmus` should accept all tests produced by the companion `diy` tool and has been successfully used on Linux, Mac OS, and on two versions of AIX.

The authors of `litmus` are Luc Maranget and Susmit Sarkar. The present `litmus` is inspired from a prototype by Thomas Braibant (INRIA Rhône-Alpes) and Francesco Zappa Nardelli (INRIA Paris-Rocquencourt).

1 A tour of litmus

1.1 A simple run

Consider the following (rather classical) `classic.litmus` litmus test for X86:

```
X86 classic
"Fre PodWR Fre PodWR"
{ x=0; y=0; }
  P0          | P1          ;
  MOV [y],$1  | MOV [x],$1  ;
  MOV EAX,[x] | MOV EAX,[y] ;
exists (0:EAX=0 /\ 1:EAX=0)
```

A litmus test source has three main sections:

1. The initial state defines the initial values of registers and memory locations. Initialisation to zero may be omitted.
2. The code section defines the code to be run concurrently — above there are two threads. Yes we know, our X86 assembler syntax is a mistake.
3. The final condition applies to the final values of registers and memory locations.

Run the test by:

```
$ litmus classic.litmus
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
% Results for classic.litmus %
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
X86 classic
"Fre PodWR Fre PodWR"

{ x=0; y=0; }

  P0          | P1          ;
  MOV [y],$1  | MOV [x],$1  ;
  MOV EAX,[x] | MOV EAX,[y] ;

exists (0:EAX=0 /\ 1:EAX=0)
Generated assembler
```

```

_litmus_P0_0_: movl $1, (%rcx)
_litmus_P0_1_: movl (%rsi), %eax
_litmus_P1_0_: movl $1, (%rsi)
_litmus_P1_0_: movl $1, (%rsi)
_litmus_P1_1_: movl (%rcx), %eax

```

```

Test classic Allowed
Histogram (4 states)
34      :>0:EAX=0; 1:EAX=0;
499911:>0:EAX=1; 1:EAX=0;
499805:>0:EAX=0; 1:EAX=1;
250     :>0:EAX=1; 1:EAX=1;
Ok

```

```

Witnesses
Positive: 34, Negative: 999966
Condition exists (0:EAX=0 /\ 1:EAX=0) is validated
Hash=eb447b2ffe44de821f49c40caa8e9757
Time classic 0.60
...

```

The litmus test is first reminded, followed by actual assembler — the machine is an AMD64, in-line address references disappeared, registers may change, and assembler syntax is now more familiar. The test has run one million times, producing one million final states, or *outcomes* for the registers `EAX` of threads P_0 and P_1 . The test run validates the condition, with 34 positive witnesses.

1.2 Cross compilation

With option `-o <name.tar>`, `litmus` does not run the test. Instead, it produces a tar archive that contains the C sources for the test.

Consider `ppc-classic.litmus`, a Power version of the previous test:

```

PPC ppc-classic
"Fre PodWR Fre PodWR"
{
0:r2=y; 0:r4=x;
1:r2=x; 1:r4=y;
}
P0      | P1      ;
li r1,1  | li r1,1    ;
stw r1,0(r2) | stw r1,0(r2) ;
lwz r3,0(r4) | lwz r3,0(r4) ;
exists (0:r3=0 /\ 1:r3=0)

```

Our target machine (ppc) runs Mac OS, wich we specify with the `-os` option:

```

$ litmus -o /tmp/a.tar -os mac ppc-classic.litmus
$ scp /tmp/a.tar ppc:/tmp

```

Then, on the remote machine ppc:

```

ppc$ mkdir classic && cd classic
ppc$ tar xf /tmp/a.tar
ppc$ ls
comp.sh run.sh ppc-classic.c outs.c utils.c

```

Test is compiled by the shell script `comp.sh` and run by the shell script `run.sh`:

```
$ sh comp.sh
$ sh run.sh
...
Test ppc-classic Allowed
Histogram (3 states)
3947  :>0:r3=0; 1:r3=0;
499357:>0:r3=1; 1:r3=0;
496696:>0:r3=0; 1:r3=1;
Ok

Witnesses
Positive: 3947, Negative: 996053
Condition exists (0:r3=0 /\ 1:r3=0) is validated
...
```

As we see, the condition validates also on Power.

The compilation script `comp.sh` produces an executable: `ppc-classic.exe`. Notice that `ppc-classic.exe` can be run directly, for a less verbose output.

1.3 Running several tests at once

Consider the additional test `ppc-storefwd.litmus`:

```
PPC ppc-storefwd
"DpdR Fre Rfi DpdR Fre Rfi"
{
0:r2=x; 0:r6=y;
1:r2=y; 1:r6=x;
}
P0          | P1          ;
li r1,1     | li r1,1     ;
stw r1,0(r2) | stw r1,0(r2) ;
lwz r3,0(r2) | lwz r3,0(r2) ;
xor r4,r3,r3 | xor r4,r3,r3 ;
lwzx r5,r4,r6 | lwzx r5,r4,r6 ;
exists (0:r3=1 /\ 0:r5=0 /\ 1:r3=1 /\ 1:r5=0)
```

To compile the two tests together, we can give two file names as arguments to `litmus`:

```
$ litmus -o /tmp/a.tar -os mac ppc-classic.litmus ppc-storefwd.litmus
```

Or, more conveniently, list the litmus sources in a file whose name starts with `@`:

```
$ cat @ppc
ppc-classic.litmus
ppc-storefwd.litmus
$ litmus -o /tmp/a.tar -os mac @ppc
```

To run the test on the remote ppc machine, the same sequence of commands as in the one test case applies:

```
ppc$ tar xf /tmp/a.tar && sh comp.sh && sh run.sh
...
Test ppc-classic Allowed
```

```

Histogram (3 states)
4167  :>0:r3=0; 1:r3=0;
499399:>0:r3=1; 1:r3=0;
496434:>0:r3=0; 1:r3=1;
Ok

Witnesses
Positive: 4167, Negative: 995833
Condition exists (0:r3=0 /\ 1:r3=0) is validated
...
Test ppc-storefwd Allowed
Histogram (4 states)
37    :>0:r3=1; 0:r5=0; 1:r3=1; 1:r5=0;
499837:>0:r3=1; 0:r5=1; 1:r3=1; 1:r5=0;
499912:>0:r3=1; 0:r5=0; 1:r3=1; 1:r5=1;
214   :>0:r3=1; 0:r5=1; 1:r3=1; 1:r5=1;
Ok

Witnesses
Positive: 37, Negative: 999963
Condition exists (0:r3=1 /\ 0:r5=0 /\ 1:r3=1 /\ 1:r5=0) is validated
...

```

Now, the output of `run.sh` shows the result of two tests.

2 Controlling test parameters

Users can control some of testing conditions. Those impact efficiency and outcome variability.

Sometimes one looks for a particular outcome — for instance, one may seek to get the outcome `0:r3=1; 1:r3=1;` that is missing in the previous experiment for test `ppc-classical`. To that aim, varying test conditions may help.

2.1 Architecture of tests

Consider a test `a.litmus` designed to run on t threads P_0, \dots, P_{t-1} . The structure of the executable `a.exe` that performs the experiment is as follows:

- So as to benefit from parallelism, we run $n = \max(1, a/t)$ (euclidean division) tests concurrently on a machine where a cores are available.
- Each of these (identical) tests consists in repeating r times the following sequence:
 - Fork t (POSIX) threads T_0, \dots, T_{t-1} for executing P_0, \dots, P_{t-1} . Which thread executes which code is either fixed, or changing, controlled by the *launch mode*. In our experience, the launch mode has marginal impact.

- Each thread T_k executes a loop of size s . Loop iteration number i executes the code of P_k (in fixed mode) and saves the final contents of its observed registers in some arrays indexed by i . Furthermore, still for iteration i , memory location x is in fact an array cell.

How this array cell is accessed depends upon the *memory mode*. In *direct mode* the array cell is accessed directly as $x[i]$; as a result, cells are accessed sequentially and false sharing effects are likely. In *indirect mode* the array cell is accessed by the means of a shuffled array of pointers; as a result we observed a much greater variability of outcomes.

If the *preload mode* is enabled, a preliminary loop of size s reads a random subset of the memory locations accessed by P_k . Preload have a noticeable effect.

The iterations performed by the different threads T_k may be unsynchronised, exactly synchronised by a pthread based barrier, or approximately synchronised by specific code. Absence of synchronisation may be interesting when t exceeds a . As a matter of fact, in this situation, any kind of synchronisation leads to prohibitive running times. However, for a large value of parameter s and small t we have observed spontaneous concurrent execution of some iterations amongst many. Pthread based barriers are exact but they are slow and in fact offers poor synchronisation for short code sequences. The approximate synchronisation is thus the preferred technique.

- Wait for the t threads to terminate and collect outcomes in some histogram like structure.

- Wait for the n tests to terminate and sum their histograms.

Hence, running `a.exe` produces $n \times r \times s$ outcomes. Parameters n , a , r and s can first be set directly while invoking `a.exe`, using the appropriate command line options. For instance, assuming $t = 2$, `./a.exe -a 201 -r 10000 -s 1` and `./a.exe -n 1 -r 1 -s 1000000` will both produce one million outcomes, but the latter is probably more efficient. If our machine has 8 cores, `./a.exe -a 8 -r 1 -s 1000000` will yield 4 millions outcomes, in a time that we hope not to exceed too much the one experienced with `./a.exe -n 1`. Also observe that the memory allocated is roughly proportional to $n \times s$, while the number of T_k threads created will be $t \times n \times r$. The `run.sh` shell script transmits its command line to all the executable (`.exe`) files it invokes, thereby providing a convenient means to control testing condition for several tests. Satisfactory test parameters are found by experimenting and the control of executable files by command line options is designed for that purpose.

Once satisfactory parameters are found, it is a nuisance to repeat them for every experiment. Thus, parameters a , r and s can also be set while invoking `litmus`, with the same command line options. In fact those settings command the default values of `.exe` files controls. Additionally, the synchronisation technique for iterations, the memory mode, and several others compile time parameters can be selected by appropriate `litmus` command line options. Finally, users can record frequently used parameters in configuration files.

2.2 Controlling executable files

Any executable file produced by `litmus` accepts the following command line options.

`-v` Be verbose, can be repeated to increase verbosity.

`-q` Be quiet.

`-a <n>` Run maximal number of tests concurrently for n available cores — parameter a in Section 2.1.

`-n <n>` Run n tests concurrently — parameter n in Section 2.1.

`-r <n>` Perform n runs — parameter r in Section 2.1.

`-fr <f>` Multiply r by f (f is a floating point number).

`-s <n>` Size of a run — parameter s in Section 2.1.

`-fs <f>` Multiply s by f .

`-f <f>` Multiply s by f and divide r by f .

2.3 Controlling litmus

Arguments

`litmus` takes file names as command line arguments. Those files are either a single litmus test, when having extension `.litmus`, or a list of file names, when prefixed by `@`. Of course, the file names in `@files` can themselves be `@files`.

Options

There are many command line options. We describe the more useful ones:

General behaviour

- `-version` Show version number and exit.
- `-libdir` Show installation directory and exit.
- `-v` Be verbose, can be repeated to increase verbosity.
- `-mach <name>` Read configuration file `name.cfg`. See the next section for the syntax of configuration files.
- `-o <name.tar>` Cross compile tests into tar file `name.tar`.

Test conditions The following options set the default values of the options of the executable files produced:

- `-a <n>` Run maximal number of tests concurrently for n available cores — set default value for `-a` of Section 2.2. Default is 0 (run one test).
- `-limit <bool>` Do not process tests with more than n threads, where n is the number of available cores defined above. Default is `false`.
- `-r <n>` Perform n runs — set default value for option `-r` of Section 2.2. Default is 10.
- `-s <n>` Size of a run — set default value for option `-s` of Section 2.2. Default is 100000.

The following additional options control the various modes described in Section 2.1. Those cannot be changed without running `litmus` again:

- `-barrier (user|pthread|none)` Set synchronisation mode, default `user`.
- `-launch (changing|fixed)` Set launch mode, default `changing`.
- `-mem (indirect|direct)` Set memory mode, default `indirect`.
- `-preload <bool>` Enable or disable preload, default enabled.
- `-safer <bool>` Enable or disable safer mode, default enabled. In safer mode, executable files performs some consistency checks. Those are intended both for debugging and for dynamically checking some assumption on POSIX threads that we rely upon.
- `-para (self|shell)` Perform several tests concurrently, either by forking POSIX threads (as described in Section 2.1), or by forking Unix processes. Only applies for cross compilation. Default is `self`.
- `-speedcheck <bool>` Enable or disable quick condition check mode, default enabled. When enabled, stop test as soon as condition is settled.
- `-ccopts <flags>` Set additional `gcc` compilation flags (defaults: X86="`-fomit-frame-pointer -O2`", PPC="`-O`").

Target architecture description Litmus compilation chain may slightly vary depending on the following parameters:

- `-os (linux|mac|aix|aix5)` Set target operating system. This parameter mostly impacts some of `gcc` options. Default `linux`.
- `-ws (w32|w64)` Set word size. This option only applies to Power, selecting `gcc` 32 or 64 bits mode. It also slightly impacts code generation in the corner case where memory locations hold other memory locations. Default `w32`.

Configuration files

The syntax of configuration files is minimal: lines “*key = arg*” are interpreted as setting the value of parameter *key* to *arg*. Each parameter has a corresponding option, usually *-key*, except for single-letter options:

<i>option</i>	<i>key</i>	<i>arg</i>
-a	avail	integer
-s	size_of_test	integer
-r	number_of_run	integer

As command line option are processed left-to-right, settings from a configuration file (option **-mach**) can be overridden by a later command line option. Some configuration files for the machines we have tested are present in the distribution. As an example here is the configuration file **hpcx.cfg**.

```
size_of_test = 2000
number_of_run = 20000
os = AIX
ws = W32
# A node has 16 cores X2 (SMT)
avail = 32
```

Lines introduced by # are comments and are thus ignored.

Configuration files are searched first in the current directory; then in any directory specified by setting the shell environment variable LITMUSDIR; and then in litmus installation directory, which is defined while compiling litmus.

Part II

Generating tests

The authors of diy are Jade Alglave and Luc Maranget (INRIA Paris–Rocquencourt).

3 Preamble

We wrote diy as part of our empirical approach to studying relaxed memory models: developping in tandem testing tools and models of multiprocessor behaviour. In this tutorial, we attempt an independant tool presentation. Readers interested by the companion formalism are invited to refer to our publications (submission...).

3.1 Relaxation of Sequential Consistency

Relaxation is one of the key concepts behind simple analysis of weak memory models. We define a relaxation by reference to the most natural model of parallel execution in shared memory: Sequential Consistency (SC), as defined by L. Lamport [2]. A parallel program running on a sequentially consistent machine behaves as an interleaving of its sequential threads.

Consider once more the example `classic.litmus`:

```
X86 classic
"Fre PodWR Fre PodWR"
{ x=0; y=0; }
  P0      | P1      ;
  MOV [y],$1 | MOV [x],$1 ; #(a)Wy1 | (c)Wx1
  MOV EAX,[x] | MOV EAX,[y] ; #(b)Rx0 | (d)Ry0
exists (0:EAX=0 /\ 1:EAX=0)
```

To focus on interaction through shared memory, let us consider memory accesses, or *memory events*. A memory event will hold a direction (write, written W, or read, written R), a memory location (written x, y) a value and a unique label. In any run of the simple example above, four memory events occur: two writes (c) Wx1 and (a) Wy1 and two reads (b) Rxv₁ with a certain value v₁ and (d) Ryv₂ with a certain value v₂.

If the program's behaviour is modelled by the interleaving of its events, the first event must be a write of value 1 to location x or y and at least one of the loads must see a 1. Thus, a SC machine would exhibit only three possible outcomes for this test:

Allowed: 0:EAX = 0 ∧ 1:EAX = 1
Allowed: 0:EAX = 1 ∧ 1:EAX = 0
Allowed: 0:EAX = 1 ∧ 1:EAX = 1

However, running (see Sec. 1.1) this test on a x86 machine yields an additional result:


Allowed: 0:EAX = 0 ∧ 1:EAX = 0

And indeed, x86 allows each write-read pair on both processors to be reordered [1]: thus the write-read pair in program order is relaxed on each of these architectures. We cannot use SC as an accurate memory model for modern architectures. Instead we analyse memory models as *relaxing* the ordering constraints of the SC memory model.

3.2 Introduction to relaxations

Consider again our classical example, from a SC perspective. We briefly argued that the outcome “0:EAX = 0 \wedge 1:EAX = 0” is forbidden by SC. We now present a more complete reasoning:

- From the condition on outcome, we get the values in read events: (b) Rx0 and (d) Ry0.
- Because of these values, (b) Rx0 must precede the write (c) Wx1 in the final interleaving of SC. Similarly, (d) Ry0 must precede the write (a) Wy1. This we note (b) \xrightarrow{fr} (c) and (d) \xrightarrow{fr} (a).
- Because of sequential execution order on one single processor (a.k.a. *program order*), (a) Wy1 must precede (b) Rx0 (first processor); while (c) Wx1 must precede (d) Ry0 (second processor). This we note (a) \xrightarrow{po} (b) and (c) \xrightarrow{po} (d).
- We synthesise the four constraints above as the following graph:




Constraint arrows or *global* arrows are shown in brown colour. As the graph contains a cycle of brown arrows, the events cannot be ordered. Hence the execution presented is not allowed by SC.

The key idea of diy resides in producing programs from similar cycles. To that aim, the edges in cycles must convey additional information:

- For \xrightarrow{po} edges, we consider whether the locations of the events on both sides of the edge are the same or not ('s' or 'd'); and the direction of these events (W or R). For instance the two \xrightarrow{po} edges in the example are PodWR. (program order edge between a write and a read whose locations are different).
- For \xrightarrow{fr} edges, we consider whether the processor of the events on both sides of the edge are the same or not ('i' for internal, or 'e' for external). For instance the two \xrightarrow{fr} edges in the example are Fre.

So far so good, but our x86 machine produced the outcome 0:EAX=0 \wedge 1:EAX=0. The Intel Memory Ordering White Paper [1] specifies: “Loads may be reordered with older stores to different locations”, which we rephrase as: PodWR is relaxed. Considering Fre to be safe, we have the graph:



And the brown sub-graph becomes acyclic.

We shall see later why we choose to relax PodWR and not Fre. At the moment, we observe that we can assume PodWR to be relaxed and Fre not to be (*i.e.* to be *safe*) and test our assumptions, by producing and running more litmus tests. The diy suite precisely provides tools for this approach.

As a first example, `classic.litmus` can be created as follows:

```
% diyone -arch X86 -name classic Fre PodWR Fre PodWR
```

As a second example, we can produce several similar tests as follows:

```
% diy -arch X86 -safe Fre -relax PodWR -name classic
Generator produced 2 tests
Relaxations tested: {PodWR}
```

diy produces two litmus tests, `classical000.litmus` and `classical001.litmus`, plus one index file `@all`. One of the litmus tests generated is the same as above, while the new test is:

```
% cat classic001.litmus
X86 classic001
"Fre PodWR Fre PodWR Fre PodWR"
Cycle=Fre PodWR Fre PodWR Fre PodWR
Relax=PodWR
Safe=Fre
{ }
  P0          | P1          | P2          ;
  MOV [z],$1  | MOV [x],$1  | MOV [y],$1  ;
  MOV EAX,[x] | MOV EAX,[y] | MOV EAX,[z] ;
exists (0:EAX=0 /\ 1:EAX=0 /\ 2:EAX=0)
% cat @all
# diy -arch X86 -safe Fre -relax PodWR -name classic
# Revision: 3333
classic000.litmus
classic001.litmus
```

diy first generates cycles from the relaxations given as arguments, up to a limited size, and then generates litmus tests from these cycles.


3.3 More relaxations

We assume the memory to be *coherent*. Coherence implies that, in a given execution, the writes to a given location are performed by following a sequence, or *coherence order*, and that all processors see the same sequence.

In diy, the coherence orders are specified indirectly. For instance, the relaxation Wse (resp. Wsi) specifies two writes, performed by different processors (resp. the same processor), to the same location ℓ , the first write preceding the second in the coherence order of ℓ . The condition of the produced test then selects the specified coherence orders. Consider for instance:

```
% diyone -arch X86 -name ws Wse PodWW Wse PodWW
```

The cycle that reveals a violation of the SC memory model is:



So the coherence order is 0 (initial store, not depicted), 1, 2 for both locations x and y. While the produced test is:

```
X86 ws "Wse PodWW Wse PodWW"
{ }
  P0      | P1      ;
  MOV [y],$2 | MOV [x],$2 ;
  MOV [x],$1 | MOV [y],$1 ;
exists (x=2 /\ y=2)
```

By the coherence hypothesis, checking the final value of locations suffices to check the coherence orders, as expressed by the final condition of `ws.litmus`:

```
exists (x=2 /\ y=2)
```

Relaxations Rfe and Rfi relate writes to reads that load their value. We are now equipped to generate the famous iriw test (independent reads of independent writes):

```
% diyone -arch X86 Rfe PodRR Fre Rfe PodRR Fre -name iriw
```

We generate its internal variation (*i.e.* where all Rfe are replaced by Rfi) as easily:

```
% diyone -arch X86 Rfi PodRR Fre Rfi PodRR Fre -name iriw-internal
```

We get the cycles of Fig. 1, and the litmus tests of Fig. 2.

Relaxations given as arguments really are a “concise specification”. As an example, we get iriw for Power, simply by changing `-arch X86` into `-arch PPC`.

```
% diyone -arch PPC Rfe PodRR Fre Rfe PodRR Fre
PPC a
"Rfe PodRR Fre Rfe PodRR Fre"
{
0:r2=y; 0:r4=x;
```

Figure 1: Cycles for iriw and iriw-internal




Figure 2: Litmus tests iriw and iriw-internal

```
X86 iriw
"Rfe PodRR Fre Rfe PodRR Fre"
{ }
P0          | P1          | P2          | P3          ;
MOV EAX,[y] | MOV [x],$1 | MOV EAX,[x] | MOV [y],$1 ;
MOV EBX,[x] |              | MOV EBX,[y] |              ;
exists (0:EAX=1 /\ 0:EBX=0 /\ 2:EAX=1 /\ 2:EBX=0)
```

```
X86 iriw-internal
"Rfi PodRR Fre Rfi PodRR Fre"
{ }
P0          | P1          ;
MOV [x],$1 | MOV [y],$1 ;
MOV EAX,[x] | MOV EAX,[y] ;
MOV EBX,[y] | MOV EBX,[x] ;
exists
(0:EAX=1 /\ 0:EBX=0 /\
 1:EAX=1 /\ 1:EBX=0)
```

```

1:r2=x;
2:r2=x; 2:r4=y;
3:r2=y;
}
P0          | P1          | P2          | P3          ;
lwz r1,0(r2) | li r1,1          | lwz r1,0(r2) | li r1,1          ;
lwz r3,0(r4) | stw r1,0(r2) | lwz r3,0(r4) | stw r1,0(r2) ;
exists (0:r1=1 /\ 0:r3=0 /\ 2:r1=1 /\ 2:r3=0)

```

Also notice that without the `-name` option, `diyone` writes its result to standard output.

3.4 Summary of simple relaxations

We summarise relaxations available on all architectures.

3.4.1 Communication relaxations

We call communication relaxations the relations between two events communicating through memory, though they could belong to the same processor. Thus, these events operate on the same memory location.

diy relaxation	Source	Target	Processor	Additional property
Rfi	W	R	Same	Target reads its value from source
Rfe	W	R	Different	Target reads its value from source
Wsi	W	W	Same	Source precedes target in coherence order
Wse	W	W	Different	Source precedes target in coherence order
Fri	R	W	Same	Source reads a value from a write that precedes target in coherence order
Fre	R	W	Different	Source reads a value from a write that precedes target in coherence order

3.4.2 Program order relaxations

We call program order relaxations each relation between two events in the program order. These events are on the same processor, since they are in program order. As regards code output, `diy` interprets a program order relaxation by generating two memory instructions (load or store) following one another.

Program order relaxations have the following syntax:

$$\text{Po}(s|d)(R|W)(R|W)$$

where:

- s (resp. d) indicates that the two events are to the same (resp. different) location(s);
- R (resp. W) indicates an event to be a read (resp. a write);

In practice, we have:

diy relaxation	Source	Target	Location
PosRR	R	R	Same
PodRR	R	R	Diff
PosRW	R	W	Same
PodRW	R	W	Diff
PosWW	W	W	Same
PodWW	W	W	Diff
PosWR	W	R	Same
PodWR	W	R	Diff

It is to be noticed that PosWR, PosWW and PosRW are similar to Rfi, Wsi and Fri, respectively. More precisely, diy is unable to consider a PosWR (or PosWW, or PosRW) relaxation as not being also a Rfi (or Wsi, or Fri) relaxation. However, litmus tests conditions may be more informative in the case of Rfi and Fri.

3.4.3 Barrier relaxations

Relaxed architectures provide specific instructions, namely *barriers* or *fences*, to enforce order of memory accesses. In diy the presence of a fence instruction is specified with Fenced relaxations, similar to Po relaxations, except that a fence instruction is inserted. Hence we have FencedsRR, FenceddRR. etc. The inserted fence is the stronger fence provided by the architecture — that is, **mfence** for x86 and **sync** for Power.

Barriers can also be specified by using specific names. More precisely, we have MFence for x86; while on Power we have Sync and LwSync. Hence, to yield two reads to different locations and separated by the lightweight Power barrier **lwsync**, we specify LwSyncdRR.

4 Testing relaxations with diy

diy can probably be used in various, creative, ways; but the tool first stems from our technique for testing relaxed memory models. The **-safe** and **-relax** options are crucial here. We describe our technique by the means of an example: X86-TSO.

4.1 Principle

Before engaging in testing it is important to categorise relaxations as safe or relaxed.

This can done by interpretation of vendor’s documentation. For instance, the iriw test of Sec. 3.3 is the example 7.7 of [1] “Stores Are Seen in a Consistent Order by Other Processors”, with a Forbid specification. Hence we deduce that Fre, Rfe and PodRR are safe. Then, from test iriw-internal of Sec. 3.3, which is Intel’s test 7.5 “Intra-Processor Forwarding Is Allowed” with an allow specification, we deduce that Rfi is relaxed. Namely, the cycle of iriw-internal is “Fre Rfi PodRR Fre Rfi PodRR”. Therefore, the only possibility is for Rfi to be relaxed.

Overall, we deduce:

- Relaxations PosWR (Rfi) and PodWR are relaxed
- The remaining relaxations PosRR, PowRR, PosWW (Wsi), PodWW, PosRW (Fri), Fre and Wse are safe. Barrier relaxations FencedsWR and FenceddWR are also safe and worth testing.

Of course these remain assumptions to be tested. To do so, we perform one series of tests per relaxed relaxation, and one series of tests for confirming safe relaxations as much as possible. Let S be all safe relaxations.

- Let r be a relaxed relaxation. We produce tests for confirming r being relaxed by **diy -relax r -safe S** . We run these tests with **litmus**. If one of the tests yields **Ok**, then r is confirmed to be relaxed (provided the experiments on S below do not fail...)
- For confirming the safe set, we produce tests by **diy -safe S** . We run these tests as much as possible and expect never to see **Ok**.

Namely, diy builds cycles as follows:

- **diy -relax r -safe S** build cycles with at least one r taking other relaxations from S .
- **diy -safe S** build cycles from the relaxations in S .

For the purpose of confirming relaxed relaxations, S can be replaced by a subset.

4.2 Testing x86

Repeating command line options is painful and error prone. Besides, configuration parameters may get lost. Thus, we regroup those in configuration files that simply list the options to be passed to diy, one option per line. For instance here is the configuration file for testing the safe relaxations of x86, `x86-safe.conf`.

```
#safe x86 conf file
-arch X86
#Generate tests on four processors or less
-nprocs 4
#From cycles of size at most six
-size 6
#With names safe000, safe0001,...
-name safe
#List of safe relaxations
-safe PosR* PodR* PodWW PosWW Rfe Wse Fre FencedsWR FenceddWR
```

Observe that the syntax of relaxations allows one shortcut: the wildcard `*` stands for `W` and `R`. Thus `PodR*` gets expanded to `PodRR PodRW`.

We get safe tests by issuing the following command, preferably in a specific directory, say `safe`.

```
% diy -conf x86-safe.conf
Generator produced 38 tests
Relaxations tested: {}
```

Here are the configuration files for confirming that `Rfi` and `PodWR` are relaxed, `x86-rfi.conf` and `x86-podwr.conf`.

```
#rfi x86 conf file
-arch X86
-nprocs 4
-size 6
-name rfi
-safe PosR* PodR* PodWW PosWW Rfe Wse Fre FencedsWR FenceddWR
-relax Rfi
```

```
#podwr x86 conf file
-arch X86
-nprocs 4
-size 6
-name podwr
-safe Fre
-relax PodWR
```

Notice that we used the complete safe relaxation list in `x86-rfi.conf` and a reduced list in `x86-podwr.conf`. Tests are to be generated in specific directories.

```
% cd rfi
% diy -conf x86-rfi.conf
Generator produced 11 tests
Relaxations tested: {Rfi}
% cd ../podwr
% diy -conf x86-podwr.conf
Generator produced 2 tests
Relaxations tested: {PodWR}
% cd ..
```

Now, let us run all tests at once, with the parameters of machine `saumur` (4 physical cores with hyper-threading):

```
% litmus -mach saumur rfi/@all > rfi/saumur.rfi.00
% litmus -mach saumur podwr/@all > podwr/saumur.podwr.00
% litmus -mach saumur safe/@all > safe/saumur.safe.00
```

If your machine has 2 cores only, try `litmus -a 2 -limit true...`

We now look for the tests that have validated their condition in the result files of `litmus`. A simple tool, `readRelax`, does the job:

```
% readRelax -arch X86 rfi/saumur.rfi.00 podwr/saumur.podwr.00 safe/saumur.safe.00
.
.
.
** Relaxation summary **
{Rfi} With {Rfe, Fre, Wse, PodRW, PodRR} {Rfe, Fre, PodRR}\
{Fre, Wse, PodWW, PodRR} {Fre, PosWW, PodRR, MFencedWR}\
{Fre, PodWW, PodRR, MFencedWR} {Fre, PodRR} {Fre, PodRR, MFencedWR}
{PodWR} With {Fre}
```

The tool `readRelax` first lists the result of all tests (which is omitted above), and then dumps a summary of the relaxations it found. The sets of the relaxations that need to be safe for the tests to indeed reveal a relaxed relaxation are also given. Here, `Rfi` and `PodWR` are confirmed to be relaxed, while no relaxation in the `safe` set is found to be relaxed. Had it been the case, a line `{ } With {...}` would have occurred in the relaxation summary. The `safe` tests need to be run a lot of times, to increase our confidence in the `safe` set.

5 Additional relaxations

We introduce some additional relaxations that are specific to the Power architecture. We shall not detail here our experiments on Power machines. See our experience report <http://diy.inria.fr/phat/> for more details.

5.1 Intra-processor dependencies

In a very relaxed architecture such as Power, *intra-processor dependencies* becomes significant. Roughly, intra-processor dependencies fall into two categories:

Data dependencies occur when a memory access instruction reads a register whose contents depends upon a previous (in program order) load. In `diy` we specify such a dependency as:

$$Dp(s|d)(R|W)$$

where, as usual, `s` (resp. `d`) indicates that the source and target events are to the same (resp. different) location(`s`); and `R` (resp. `W`) indicates that the target event is a read (resp. a write). As a matter of fact, we do not need to specify the direction of the source event, since it always is a read.

Control dependencies occur when the execution of a memory access is conditioned by the contents of a previous load. Their syntax is similar to the one of `Dp` relaxations, with a `Ctrl` tag:

$$Ctrl(s|d)(R|W)$$

In the produced code, `diy` expresses a data dependency by a *false dependency* (or *dummy dependency*) that operates on the address of the target memory access. For instance:

```
% diyone DpdW Rfe DpdW Rfe
PPC a "DpdW Rfe DpdW Rfe"
{ 0:r2=x; 0:r5=y; 1:r2=y; 1:r5=x; }
P0      | P1      ;
lwz r1,0(r2) | lwz r1,0(r2) ;
```

```

xor r3,r1,r1 | xor r3,r1,r1 ;
li r4,1      | li r4,1      ;
stwx r4,r3,r5 | stwx r4,r3,r5 ;
exists (0:r1=1 /\ 1:r1=1)

```

On P_0 , the effective address of the indexed store `stwx r4,r3,r5` depends on the contents of the index register `r3`, which itself depends on the contents of `r1`. The dependency is a “false” one, since the contents of `r3` always is zero, regardless of the contents of `r1`.

A control dependency is implemented by the means of an useless compare and branch sequence, plus the `isync` instruction when the target event is a load. For instance

```

% diyone CtrlldR Fre SyncdWW Rfe
PPC a
"CtrlldR Fre SyncdWW Rfe"
{ 0:r2=x; 0:r4=y; 1:r2=y; 1:r4=x; }
P0          | P1          ;
li r1,1      | lwz r1,0(r2) ;
stw r1,0(r2) | cmpw r1,r1   ;
sync        | beq LC00   ;
li r3,1      | LC00:         ;
stw r3,0(r4) | isync        ;
              | lwz r3,0(r4) ;
exists (1:r1=1 /\ 1:r3=0)

```

Of course, in both cases, we assume that dependencies are not “optimised out” by the assembler or the hardware.

5.2 Composite relaxations and cumulativity

Users may specify a small sequence of single relaxations as behaving as a single relaxation to `diy`. The syntax is:

$$[r1, r2, \dots]$$

The main usage of the feature is to specify *cumulativity relaxations*, that is, the sequence of `Rfe` and of a barrier relaxation (A-cumulativity), or the sequence of a barrier relaxation and of `Rfe` (B-cumulativity). Cumulativity relaxations are best expressed by the following syntactical shortcuts: let r be a barrier relaxation, then `ACr` stands for `[Rfe,r]`, while `BCr` stands for `[r,Rfe]`.

Hence, a simple way to generate iriw-like (see Sec. 3.3) litmus tests with `lwsync` is as follows:

```

% diy -name iriw-lwsync -nprocs 8 -size 8 -relax ACLwSyncdRR -safe Fre
Generator produced 3 tests
Relaxations tested: {ACLwSyncdRR}

```

where we have for instance:

```

% cat iriw-lwsync001.litmus
PPC iriw-lwsync001
"Fre Rfe LwSyncdRR Fre Rfe LwSyncdRR Fre Rfe LwSyncdRR"
Cycle=Fre Rfe LwSyncdRR Fre Rfe LwSyncdRR Fre Rfe LwSyncdRR
Relax=ACLwSyncdRR
Safe=Fre
{
0:r2=z; 0:r4=x; 1:r2=x;
2:r2=x; 2:r4=y; 3:r2=y;

```

```

4:r2=y; 4:r4=z; 5:r2=z;
}
P0          | P1          | P2          | P3          | P4          | P5          ;
lwz r1,0(r2) | li r1,1      | lwz r1,0(r2) | li r1,1      | lwz r1,0(r2) | li r1,1      ;
lwsync      | stw r1,0(r2) | lwsync      | stw r1,0(r2) | lwsync      | stw r1,0(r2) ;
lwz r3,0(r4) |              | lwz r3,0(r4) |              | lwz r3,0(r4) |              ;
exists (0:r1=1 /\ 0:r3=0 /\ 2:r1=1 /\ 2:r3=0 /\ 4:r1=1 /\ 4:r3=0)

```

6 Command usage

The diy suite consists in three tools:

diyone generates one litmus test from the specification of a violation of the sequential consistency memory model as a cycle—see Sec. 3.2.

diy generates several tests, aimed confirming that relaxations are relaxed or safe—see Sec. 4.

readRelax Extract relevant information from the results of tests—see Sec. 4.2.

6.1 Usage of diyone

diyone takes a list of relaxations as arguments and outputs a litmus test. Note that **diyone** may fail to produce the test, with a message that briefly details the failure.

```

% diyone Rfe Rfe PodRR
Test a [Rfe Rfe PodRR] failed:
Impossible direction PodRR Rfe

```

diyone accepts the following documented options.

-v Be verbose, repeat to increase verbosity.

-version Show version number and exit.

-obs Enable observers. Observers are additional threads that check the coherence order of locations that are written into three times or more. Without observers, **diyone** fails if three writes to the same location are specified.

-optcond Optimise conditions by disregarding the values of loads that are neither the target of Rf, nor the source of Fr. This is the default.

-nooptcond Do not optimise conditions.

-name <name> Set the name of the test to **<name>** and output it into file **<name>.litmus**. By default, the test name is **a** and output goes to standard output.

-arch (X86|PPC) Set architecture. Default is PPC.

6.2 Usage of diy

The tool **diy** accepts the same options as **diyone**, option **-name <name>** being mandatory and setting the base name of generated litmus tests: *i.e.* **diy** produces tests **<name>000**, **<name>001**, etc., in files with extension **.litmus**. Moreover, **diy** produces an index file **@all** that lists file names **<name>000.litmus**, **<name>001.litmus** etc.

diy also accepts the following, additional, documented options.

-conf <file> Read configuration file <file>. A configuration file consists in a list of options, one option per line. Lines introduced by # are comments and are thus ignored.

-o <name.tar> Output litmus tests as an archive <name.tar>. The default is to output them in the current directory.

-size <n> Set the maximal size of cycles. Default is 6.

-exact Produce cycles of size exactly <n>, in place of size up to <n>.

-nprocs <n> Reject tests with more than <n> threads. Default is 4.

-c <bool> Avoid equivalent cycles. Default is **true**.

-relax <relax-list> Set relax list. Default is empty. The syntax of <relax-list> is a comma (of space) separated list of relaxations.

-safe <relax-list> Set safe list. Default is empty.

The relax and safe lists command the internal generation of cycles as follows:

1. When the relax list is empty, cycles are built from the relaxations of the safe list.
2. When the relax list is of size 1, cycles are built from its single element r and from the elements of the safe list. Additionally, the cycle produced contains r at least once.
3. When the relax list is of size n , with $n > 1$, diy generates n independent sets of cycles, each set being built with one relaxation from the relax list and all the relaxations in the safe list. In other words, diy on a relax list of size n behaves similarly to n runs of diy on each relaxation in the list.

6.3 Usage of readRelax

readRelax is a simple tool to extract relevant information out of litmus run logs. For a given run of a given litmus test, the relevant information is:

- Whether the test yielded Ok or not,
- An optional relaxation, which is the one given as argument to diy option **-relax** at test build time, or none.
- The safe list relevant to the given test, *i.e.* the safe relaxations that appear in the tested cycle.

See Sec. 4.2 for an example.

The tool readRelax takes file names as arguments. If no argument is present, it reads a list of file names on standard input, one name per line. Only the **-arch** option is accepted, with default PPC. This is to be noticed, since forgetting the option **-arch X86** for reading x86 logs will make readRelax fail.

References

- [1] Intel 64 Architecture Memory Ordering White Paper, August 2007.
- [2] L. Lamport. How to make a correct multiprocess program execute correctly on a multiprocessor. *IEEE Trans. Comput.*, 46(7):779–782, 1979.