

---

# **ATS Knowledge Documentation**

*Release 0.1*

**Zhiqiang Ren**

June 01, 2015



<b>1</b>	<b>Annotated ATS Programs</b>	<b>3</b>
1.1	Goal . . . . .	3
1.2	Examples . . . . .	3
<b>2</b>	<b>ATS Tools' documentation</b>	<b>9</b>
2.1	Tag Generator For ATS-Anairiats . . . . .	9
2.2	Vim plug-in for ATS-Postiats . . . . .	10
<b>3</b>	<b>ATS-Postiats' Blog</b>	<b>11</b>
3.1	Syntax for <i>staload</i> . . . . .	11
3.2	Various <i>print</i> functions in ATS-Postiats . . . . .	12
3.3	Miscellaneous topics related to ATS-Postiats . . . . .	12
3.4	Targeting C# from program of ATS-Postiats . . . . .	13
3.5	Get value in the statics of ATS . . . . .	17
3.6	Template for Makefile of simple ATS project . . . . .	18
3.7	Global value with linear type . . . . .	19
<b>4</b>	<b>ATS-Postiats' syntax</b>	<b>21</b>
4.1	Mapping from source code to data structure . . . . .	21
<b>5</b>	<b>Model Checking ATS</b>	<b>23</b>
5.1	Tutorial . . . . .	23
5.2	Table of Contents . . . . .	27
5.3	Bibliography . . . . .	34
<b>6</b>	<b>Indices and tables</b>	<b>35</b>



Contents:



---

## Annotated ATS Programs

---

### 1.1 Goal

The normal way to start up a programming project using ATS is via imitating an existing example. The website serves as a collection of such examples along with intensive comments about the usage of ATS in these examples.

These examples come primarily from the [Google group for ATS](#).

xxx

### 1.2 Examples

WireShark has been installed on all the Windows machines in the Instructional Lab (EMA 302). It is only usable during the period of the lab session.

#### 1.2.1 Dining Philosopher with Animation (by Cairo)

##### Source File Download

##### Code

- `dp_gui.dats`

```
1 Copy from
2 Assignment 3:
3 Class: BU CAS CS520, Fall, 2013
4 Due: Thursday, the 26th of September, 2013
5 *)
6
7 (* ***** ***** *)
8 //
9 #include
10 "share/atspre_define.hats"
11 #include
```

- `dp_observer.dats`

```

1  (*
2  **
3  **
4  **
5  *)
6
7  (* ***** *)
8
9  (*
10 Copy from
11 Assignment 3:
12 Class: BU CAS CS520, Fall, 2013
13 Due: Thursday, the 26th of September, 2013
14 *)
15
16 (* ***** *)
17 //
18 #include
19 "share/atspre_define.hats"
20 #include
21 "share/atspre_staload.hats"
22 //
23 (* ***** *)
24
25 staload
26 UN = "prelude/SATS/unsafe.sats"
27
28 (* ***** *)
29
30 staload "libc/SATS/time.sats"
31
32 (* ***** *)
33
34 staload "${SCAIRO}/SATS/cairo.sats"
35
36 (* ***** *)
37
38
39 staload "${GTK}/SATS/gdk.sats"
40 staload "${GTK}/SATS/gtk.sats"
41 staload "${GLIB}/SATS/glib.sats"
42 staload "${GLIB}/SATS/glib-object.sats"
43
44 (* ***** *)
45
46 staload "mythread.sats"
47
48 // staload "libc/SATS/stdlib.sats"
49 staload "libc/SATS/unistd.sats"
50
51
52 staload "dp_observer.sats"
53 dynload "dp_observer.dats"
54
55 staload "DiningPhil.sats"
56 dynload "DiningPhil.dats"
57 (* ***** *)
58

```

```

59 %{^
60 typedef
61 struct { char buf[32] ; } bytes32 ;
62 %} // end of [%{^]
63 abst@type bytes32 = $extype"bytes32"
64
65 (* ***** ***** *)
66
67 %{^
68 #define mystrftime(bufp, m, fmt, ptm) strftime((char*)bufp, m, fmt, ptm)
69 %} // end of [%{^]
70
71 (* ***** ***** *)
72
73 extern fun worker (darea: !GtkDrawingAreal): void
74
75 (* ***** ***** *)
76 val x: ref int = ref<int>(0)
77 val cg: ref double = ref<double>(0.0)
78 val cb: ref double = ref<double>(1.0)
79
80
81 (* ***** ***** *)
82
83 %{^
84 typedef char **charptrptr ;
85 %} ;
86 abstype charptrptr = $extype"charptrptr"
87
88
89 (* ***** ***** *)
90 (* ***** ***** *)
91 extern
92 fun{} fexpose (!GtkDrawingAreal): gboolean
93
94 implement{
95 } fexpose (darea) = let
96   val () = draw_drawingarea (darea) in GFALSE
97 end // end of [fexpose]
98
99 extern fun{
100 } dp_anime_main (): void
101
102 (* ***** ***** *)
103
104 #define nullp the_null_ptr
105
106 #define W 400
107 #define H 400
108
109 implement{}
110 dp_anime_main
111   ((*void*)) = () where
112   {
113   //
114   val win0 =
115     gtk_window_new (GTK_WINDOW_TOPLEVEL)
116   val win0 = win0

```

```

117 val () = assertloc (ptrcast(win0) > 0)
118 val () = gtk_window_set_default_size (win0, (gint)W, (gint)H)
119 //
120 val opt = stropt_some"dining philosopher"
121 val issome = stropt_is_some(opt)
122 //
123 val () =
124 if issome then let
125   val title = stropt_unsome (opt)
126 in
127   gtk_window_set_title (win0, gstring(title))
128 end // end of [if] // end of [val]
129 //
130 val darea =
131   gtk_drawing_area_new ()
132 val p_darea = gobjref2ptr (darea)
133 val () = assertloc (p_darea > 0)
134 val () = gtk_container_add (win0, darea)
135 //
136 val _sid = g_signal_connect
137 (
138   darea, (gsignal)"expose-event", G_CALLBACK(fexpose), (gpointer)nullp
139 )
140 //
141 val _sid = g_signal_connect
142 (
143   win0, (gsignal)"delete-event", G_CALLBACK(gtk_main_quit), (gpointer)nullp
144 )
145 val _sid = g_signal_connect
146 (
147   win0, (gsignal)"destroy-event", G_CALLBACK(gtk_widget_destroy), (gpointer)nullp
148 )
149 //
150 val () = gtk_widget_show_all (win0)
151 //
152 val () = g_object_unref (win0) // HX: refcount of [win0] decreases from 2 to 1
153 //
154 // todo: Why does this pass the typechecking?
155 val () = mythread_create_cloptr (llam () => worker (darea))
156 val () = dp_init ()
157
158 //
159 val ((*void*)) = gtk_main ((*void*))
160
161 val () = g_object_unref (darea)
162 //
163 } // end of [dp_anime_main]
164
165 (* ***** ***** *)
166
167 implement
168 main0 (argc, argv) =
169 {
170 //
171 var argc: int = argc
172 var argv: charptrptr = $UN.castvwtpl{charptrptr}(argv)
173 //
174 val () = $extfcall (void, "gtk_init", addr@(argc), addr@(argv))

```

```
175
176 val ((*void*)) = dp_anime_main ((*void*))
177 //
178 } (* end of [main0] *)
179
180 implement worker (darea) = let
181   val () = ignoret (sleep(1))
182   //
183   val (fpf_win | win) = gtk_widget_get_window (darea)
184   //
185   val isnot = g_object_isnot_null (win)
186   //
187   prval () = minus_addback (fpf_win, win | darea)
188   //
189   in
190   //
191   if isnot then let
192     var alloc: GtkAllocation
193     val () = gtk_widget_get_allocation (darea, alloc)
194     val () = gtk_widget_queue_draw_area (darea, (gint)0, (gint)0, alloc.width, alloc.height)
195     in
196     worker(darea)
197   end else
198     ()
199   // end of [if]
200   //
201 end // end of [worker]
202
203 (* ***** *)
204
205 (* end of [dp_gui.dats] *)
```

## 1.2.2 Operations pertaining to array0

### Source File Download

### Code



---

## ATS Tools' documentation

---

Contents:

### 2.1 Tag Generator For ATS-Anairiats

#### 2.1.1 Description

For users of Emacs and Vim, it is common to browse source code spanning over multiple files with the support of *tag*. Here, we provide a tool for generating *tag* files accepted by these two editors.

#### 2.1.2 Download

Please download the file `ats-lang-tools.jar`.

#### 2.1.3 Usage

The generation of *tag* file includes two steps as follows.

1. Use *atsopt* to collect information from the source files you have.

```
atsopt -o MYTAGS --taggen -s fact.sats -d fact.dats
```

2. Use *ats-lang-tools.jar* to generate the tag file from *MYTAGS*.

```
java -jar ats-lang-tools.jar --input MYTAGS -c --output tags # "c" for ctags with vim
java -jar ats-lang-tools.jar --input MYTAGS -e --output TAGS # "e" for etags with emacs
```

The next example shows how to generate the tag file for the source files of ATS-Postiats. Using option *-output-a*, *atsopt* would output to file *MYTAGS* accumulatively, so we can combine *find* and *atsopt* to generate a large *MYTAGS*.

```
# Make sure we start from a clean slate.
rm -f MYTAGS

find ${PATSHOME}/src -name "*.sats" -exec atsopt --output-a MYTAGS --taggen -s {} \;
find ${PATSHOME}/src -name "*.dats" -exec atsopt --output-a MYTAGS --taggen -d {} \;
java -jar ats-lang-tools.jar --input MYTAGS -c --output tags
```

I use `${PATSHOME}/src` in the *find* command so that the generated *tags* would use absolute path for each file. In this way, we can open vim at any location.

The aforementioned method can be applied to ATS-Postiats as well. The following example shows how to generate the tag file for source files in the *prelude* of ATS-Postiats, which are written in ATS-Postiats.

```
PATH_PRELUDE=${PATSHOME}/prelude
MYTAGS_PATS_PRELUDE_PATS=${PATH_PRELUDE}/MYTAGS_PATS_PRELUDE_PATS

rm -rf ${MYTAGS_PATS_PRELUDE_PATS}
# Exclude two subdirectories "CODEGEN" and "DOCUGEN"
find ${PATH_PRELUDE}/SATS \( -name "CODEGEN" -o -name "DOCUGEN" \) -prune -o -name "*.sats" \
  -exec patsopt --output-a ${MYTAGS_PATS_PRELUDE_PATS} --taggen -s {} \;
find ${PATH_PRELUDE}/DATS \( -name "CODEGEN" -o -name "DOCUGEN" \) -prune -o -name "*.dats" \
  -exec patsopt --output-a ${MYTAGS_PATS_PRELUDE_PATS} --taggen -d {} \;

java -jar ${PATSHOME}/ats-lang-tools.jar -c --input ${MYTAGS_PATS_PRELUDE_PATS} --output ${PATH_
```

### 2.1.4 Development

The project is held on Github with the following address <https://github.com/alex-ren/org.ats-lang.toolats>.

## 2.2 Vim plug-in for ATS-Postiats

Please refer to this [post](#) for detailed information.

---

## ATS-Postiats' Blog

---

Some other sources of ATS: [ATS Wiki](#).

Contents:

### 3.1 Syntax for *staload*

#### 3.1.1 Macros in the path

In **ATS2**, keyword *staload* is followed by a literal string, which specifies the location of the file to be loaded. Besides those relative path and absolute path commonly seen, we can use *macro* inside the literal string. One example goes as follows:

```
staload "${GTK}/SATS/gdk.sats"
```

The *GTK* is not a system environment variable. It is actually a macro defined in the ATS code. For example, we can use the following code to set the *GTK* to the directory containing **ATS2** code for gdk library.

```
#define GTK_targetloc "~/codebase/contrib/GTK"
staload "${GTK}/SATS/gdk.sats"
```

**Note:** When we use *define* to define the macro, we are actually using the name *GTK\_targetloc*, not *GTK*.

There are some special macros we can use, which are provided by the **ATS2** compiler, e.g. *PATSHOME* and *PATSHOMERELOC*. The following example shows their usage:

```
#define GTK_targetloc "$PATSHOMERELOC/contrib/GTK"
```

The macro *PATSHOMERELOC* in **ATS2** compiler is from the environment variable *PATSHOMERELOC*, which is set in the environment before executing *patscc* or *patsopt*.

**Note:** The macro *PATSHOMERELOC* is normally set to the *contrib* package released along with ATS.

Going further, let's have a look of the file `$PATSHOME/share/HATS/atspre_define_pdgreloc.hats`, which is included by the file `$PATSHOME/share/atspre_define.hats`, which is commonly included in almost all *dats* files.

```
//
#define
ZLOG_targetloc "$PATSHOMERELOC/contrib/zlog"
//
```

```
#define
JSONC_sourceloc "$PATSLIB_URL/contrib/json-c"
#define
JSONC_targetloc "$PATSHOMERELOC/contrib/json-c"
```

An ATS program using libraries of zlog and json-c would need to include the following snippet of code:

```
#include "share/atspre_define.hats"

staload "${JSONC}/SATS/json_ML.sats"
```

## 3.2 Various *print* functions in ATS-Postiats

### 3.2.1 dddd

---

#### Todo

- `fprint_val<a>(stdout_ref, x)`
  - `fprint_newline`
- 

## 3.3 Miscellaneous topics related to ATS-Postiats

### 3.3.1 Setting of Environment

Please follow the instructions on [ATS' website](#) to install *ATS-Postiats* as well as *ATS2-contrib*. Here I just want to emphasize the setting of environment variables *PATSHOME* as well as *PATSHOMERELOC*. These two variables should be set in the environment before invoking the ATS compiler (*patscc* or *patsopt*). Also they are commonly referred in *Makefile* used in ATS projects. The first one (*PATSHOME*) should be set to the directory where ATS is installed. Normally I just download tarball for the source of ATS, decompress it, build ATS, and use built executables. (Simply put, I don't do `make install`.) Therefore I just set *PATSHOME* to the folder resulting from decompressing the tarball. *PATSHOMERELOC* should be set to the folder resulting for decompress the tarball for *ATS2-contrib*. Also the *PATH* variable should be set accordingly so that system can locate the compiler of ATS. Normally, I put these settings into one script file, say *pats.xxx.sh*. Then I do `source pats.xxx.sh` when I open a terminal for the first time. My *pats.xxx.sh* looks like the following:

```
# Script for setting environment for ATS-Postiats

PATSHOME=${HOME}/programs/ats2/ATS2-Postiats-0.0.8; export PATSHOME
PATSHOMERELOC=${HOME}/programs/ats2/ATS2-Postiats-contrib-0.0.7; export PATSHOMERELOC

PATH=${PATSHOME}/bin:${PATH}; export PATH
```

If you use `make install` to install ATS at the system level, your script would probably be like the following:

```
# Script for setting environment for ATS-Postiats

PATSHOME=/usr/local/lib/ATS2-Postiats-0.0.8; export PATSHOME
PATSHOMERELOC=${HOME}/programs/ats2/ATS2-Postiats-contrib-0.0.7; export PATSHOMERELOC
```

---

**Note:** Since we choose to install ATS at the system level, there's no need to set *PATH*.

---

### 3.3.2 Usage of *staload*

Some useful information about *staload* can be found in [Wiki](#) and my previous [blog](#).

### 3.3.3 Standard *header* files

To use the *prelude* library of ATS, we would include the following code:

```
#include "share/at spre_define.hats"
#include "share/at spre_staload.hats"
```

To use the *ML* library of ATS, we would include the following code:

```
#include "share/HATS/at spre_staload_libats_ML.hats"
```

If we want to generate C code used on lower level systems, such as embedded system, we can replace these *header* files with appropriate ones to fit the targeting platform.

### 3.3.4 Type conversion

#### Todo

- `$UNSAFE.cast {natLt (n) }`

## 3.4 Targeting C# from program of ATS-Postiats

I am working on generating C# code from the second layer of the syntax tree of ATS-Postiats. The following includes some lessons I've learned and corresponding design decision I've made.

### 3.4.1 Conversion of types

Since C# supports powerful generic types, it feels so natural to try to translate the polymorphic types in ATS into generic types in C#. The following is the ATS code from which I want to generate C# code.

```
extern val p: ptr

fun foo1 {a:type} (x: (int, a)): int = x.0

fun test_foo1 (): void = let
  val x = foo1 ((0, p))
in
end

fun foo2 {a,b:type} (f: a -> b): int = 42

fun foo3 (): ptr -> ptr = let
  fun foo4 (x: ptr): ptr = x
in
  foo4
end

fun test_foo2 (): void = let
```

```

    val f = foo3 ()
    val x = foo2 (f)
in
end

/* ***** */

fun foo5 (f: {a:type} a -> a): ptr = let
  val r = f (p)
in
  r
end

```

In ATS, tuple, record and function types have no names, and their equality is based on their content. In C#, all types must have names including structure, class, and delegate. Therefore I have to translate ATS code into C# code in the way shown in the sample code below. That is to create very general generic types in C# for tuple and functions types in ATS. It seems that there's no better way around this except the usage of *Object* type every where. Since I want to avoid *ugly* type casting, I plot to generate the code as follows

```

public class Ptr {};

/*
 * Tuple type must have name.
 */
public class Tuple2<T1, T2> {
    public T1 m1;
    public T2 m2;
    public Tuple2(T1 t1, T2 t2) {
        m1 = t1;
        m2 = t2;
    }
}

/*
 * Function type must have name.
 */
public delegate T2 Fool<T1, T2>(T1 x);

public class Code {

    static public Ptr p = new Ptr();

    static public int foo1<T1>(Tuple2<int, T1> x) {
        return x.m1;
    }

    static public void test_fool() {
        var x = foo1(new Tuple2<int, Ptr>(0, p));
    }

    static public int foo2<T1, T2>(Fool<T1, T2> f) {
        return 42;
    }

    static public Fool<Ptr, Ptr> foo3() {
        return foo4;
    }
}

```

```

static public Ptr foo4(Ptr x) {
    return x;
}

static public void test_foo2() {
    var f = foo3();
    int x = foo2(f);
}

/* ***** */

// C# doesn't allow this.
// static public Ptr foo5(Fool f) {
//     var r = f(p);
//     return r;
// }

static public void Main() {
    return;
}
}

```

The code above compiles when function *foo5* commented out. The reason I have to comment out *foo5* goes as follows. In ATS, polymorphic function is of first class. (Object of polymorphic function type can be passed around. E.g. *foo5* takes one as input argument.) In C#, no object can be of open generic type (including generic delegate). Therefore the function “foo5” in ATS cannot be simply translated into a generic delegate of C#.

Therefore I choose to use *Object* type in C# to represent type parameter of polymorphic type in ATS. But this is still not good. Such candidate in C# is shown below.

```

using System;

public class Ptr {};

/*
 * Tuple type must have name.
 */
public class Tuple2<T1, T2> {
    public T1 m1;
    public T2 m2;
    public Tuple2(T1 t1, T2 t2) {
        m1 = t1;
        m2 = t2;
    }
}

/*
 * Function type must have name.
 */
public delegate T Fool<T>(T x);

public class Code {

    static public Ptr p = new Ptr();

    static public int fool(Tuple2<int, Object> x) {
        return x.m1;
    }
}

```

```

static public void test_fool() {
    var x = fool(new Tuple2<int, Object>(0, p));
}

static public int foo2(Fool<Object> f) {
    return 42;
}

static public Fool<Ptr> foo3() {
    return foo4;
}

static public Ptr foo4(Ptr x) {
    return x;
}

static public void test_foo2() {
    var f = foo3();
    int x = foo2((Fool<Object>)(Object)f);
}

static public Ptr foo5(Fool<Object> f) {
    var r = f(p);
    return (Ptr)r;
}

static public void Main() {
    return;
}
}

```

Due to aforementioned decision, I have to give *foo2* the type *Fool<Object>* as shown above. Then to make *test\_foo2* compilable, I have to cast *f* to *Object*, then to *FOOL<Object>*. Also I have to use cast again in *foo5*. Such heavy usage of casting contradicts my original idea of relying on the generic type system of C#.

Therefore I simply choose to turn all the boxed types into *Object* and add proper type conversion whenever deemed necessary. (E.g. getting member of a tuple, invoking via a function pointer.) This is also the convention when generating C code from ATS program. The difference is that in C we rely on *void \** instead of *Object*. A hand written candidate is shown below.

```

using System;

/*
 * Tuple type must have name.
 */
public class Tuple2<T1, T2> {
    public T1 m1;
    public T2 m2;
    private Tuple2(T1 t1, T2 t2) {
        m1 = t1;
        m2 = t2;
    }
    static public Object create(T1 t1, T2 t2) {
        return new Tuple2<T1, T2>(t1, t2);
    }
}

/*

```

```

* Function type must have name.
*/
public delegate Object Foo1(Object x);

public class Code {

    static public Object p = new Object();

    static public int foo1(Object x) {
        return ((Tuple2<int, Object>)x).m1;
    }

    static public void test_foo1() {
        var x = foo1(Tuple2<int, Object>.create(0, p));
    }

    static public int foo2(Foo1 f) {
        return 42;
    }

    static public Foo1 foo3() {
        return foo4;
    }

    static public Object foo4(Object x) {
        return x;
    }

    static public void test_foo2() {
        var f = foo3();
        int x = foo2(f);
    }

    static public void Main() {
        return;
    }
}

```

In my implementation of C# code generator, I track the usage of all the tuples and records, define corresponding generic types for them. And I track all the function definitions, define corresponding delegate types for them.

## 3.5 Get value in the statics of ATS

### Todo

organize

ATS-Postiats/prelude/basics\_dyn.sats

```

dataprop
EQINT (int, int) = {x: int} EQINT (x, x)
//
extern prfun eqint_make {x,y:int | x == y} (): EQINT (x, y)
//
extern prfun
eqint_make_gint

```

```
{tk:tk}{x:int} (x: glint (tk, x)): [y: int] EQINT (x, y)

fun goo {x:int | x == 1} (): void = ()

fun foo (): void = let
  val x = 1
  val [y:int] EQINT () = eqint_make_gint (x)
in
  goo {y} ()
end
```

## 3.6 Template for Makefile of simple ATS project

An old version of Makefile template I wrote for ATS can be found here (todo).

The newer version of Makefile, which relies on *Makefile* provided by ATS, is shown below.

Makefile

```
#Makefile for untyped lambda calculus
#####

include $(PATSHOME)/share/atsmake-pre.mk

#####
#
CFLAGS += -O2
#
#####
#
# TODO
#LDFLAGS := # uncomment this for a clean start of LDFLAGS
# LDFLAGS += -lm # uncomment this for math library
#
# TODO
# Uncomment the following if you want the Boehm GC.
# By default, MALLOCFLAG is -DATS_MEMALLOC_LIBC
# LDFLAGS += -lgc
# MALLOCFLAG := -DATS_MEMALLOC_GCBDW
#
#####
# TODO
# Add source files here.
SOURCES_DATS += UTLC.dats
SOURCES_SATS +=
SOURCES_C +=

# TODO
# Specify the name of the target.
TARGET=UTLC

MYTARGET=$(TARGET)

include $(PATSHOME)/share/atsmake-post.mk
```

```
#####

#####

cleanats:: ; $(RMF) *_?ats.c

cleanall:: ; $(RMF) $(TARGET)

#####

#####

# For the future when tag generator is provided.
# .PHONY: tags
# tags:
#     $(RMF) tags MYTAGS
#     find ./ -name "*.sats" -exec $(PATSOPT) --output-a MYTAGS --taggen -s {} \;
#     find ./ -name "*.dats" -exec $(PATSOPT) --output-a MYTAGS --taggen -d {} \;
#     java -jar ../../../../tool/ats-lang-tools.jar --input MYTAGS -c --output tags

##### end of [Makefile] #####
```

### 3.7 Global value with linear type

Global values with linear type can only be used in the global scope. They cannot be used inside a function scope. For example, the following code cannot pass the type checking.

```
staload
UNSAFE = "prelude/SATS/unsafe.sats"

absvtype VT

extern fun create (): VT
extern fun use (v: !VT): void

val v = create ()

fun temp (): void = let
  val () = use (v)
in
end

implement main0 () = ()
```

**ATS compiler complains that** regexp\_main.dats: 177(line=15, offs=17) – 178(line=15, offs=18): error(3): the linear dynamic variable [v\$64(-1)] is expected to be local but it is not. regexp\_main.dats: 177(line=15, offs=17) – 178(line=15, offs=18): error(3): a linear component of the following type is abandoned: [S2Ecst(VT)]. patsopt(TRANS3): there are [2] errors in total. exit(ATS): uncaught exception: \_2home\_2alex\_2programs\_2ats2\_github\_2ATS\_2dPostiats\_2src\_2pats\_error\_2esats\_\_FatalErrorExn(1025)

Usually we would cast the global value of linear type into non-linear type, and cast it back inside a function scope, which is shown below:

```
staload
UNSAFE = "prelude/SATS/unsafe.sats"

absvtype VT

extern fun create (): VT
extern fun use (v: !VT): void

val v = create ()

val ele = $UNSAFE.castvwtp0{ptr}(v)

fun temp (): void = let
  val v1 = $UNSAFE.castvwtp0{VT}(ele)
  val () = use (v1)
  prval ((*void*)) = $UNSAFE.cast2void (v1)
in
end

implement main0 () = ()
```

---

## ATS-Postiats' syntax

---

Contents:

### 4.1 Mapping from source code to data structure

#### 4.1.1 Function declaration and implementation

We use the following source code as an example.

```
abstype ty (int, int)
extern fun foo {x: int} {y: int} (x: ty (x, y), y: int y):
  [q: int] int q

implement foo {x}{y}(x, y) = 3
```

For the function declaration, we would have a *D2Cdcstdecs*, which contains a list of *d2cst*. Each *d2cst* has type information of the defined constant. In this example, the type of function *foo* is quite complicated, which involves universal type (*S2Euni*) as well as existential type (*S2Eexi*).

ss



---

## Model Checking ATS

---

In this project, we focus on integrating model checking techniques seamlessly into the development of ATS program and ultimately build a practical system for verifying concurrent ATS program.

### 5.1 Tutorial

Before going to the details, let's have a quick look of how the methodology looks like. The producer-consumer problem is a classic one in field of concurrent programming. A recommended implementation is described in the documentation of ATS<sup>1</sup>, which exploits the type system of ATS to better ensure the correctness of the program. Certain mistakes, as stated below, can be avoided, which is great. However, a well-typed implementation for producer-consumer problem in ATS may still cause deadlock. And it's very difficult to solely rely on type system to capture such errors. Therefore, we start seeking help from other techniques, among which model checking is our pick here. It can help detect bugs related to temporal properties in concurrent systems and provide corresponding counterexamples. To apply the model checking technique, we need to form up the precise semantics of ATS programs, which in turn requires the precise semantics of those concurrency primitives related to communication and synchronization. We form up a collection of such primitives, based on which programmers can build concurrent program in ATS with semantics meaningful to the model checking techniques we employ here. Such collection is given in the file `conats.sats`. It also contains some other primitives used for model checking, which we shall explain as we see more examples.

The complete implementation for producer-consumer problem can be found here `16_reader_writer.dats`. You can also read, modify, and verify the implementation via our website [Model Checking ATS](#). We illustrate some of the code snippets below.

As indicated in<sup>1</sup>, a shared object contains a linear object, which in this example is a linear buffer. The primitives provided in `conats.sats` do not support such type. Therefore we define a linear type `lin_buffer` as well as corresponding functions for manipulating objects of such type, which is given below.

```

1 // Define linear buffer to prevent resource leak.
2 absviewtype lin_buffer (a:t@type)
3
4 local
5   assume lin_buffer (a) = atomref (a)
6 in
7   fun lin_buffer_create {a:t@type} (
8     data: a): lin_buffer a = let
9     val ref = conats_atomref_create (data)
10    in
11    ref
12  end

```

<sup>1</sup> <http://ats-lang.sourceforge.net/EXAMPLE/EFFECTIVATS/Producer-Consumer/main.html>

```

13
14 fun lin_buffer_update {a:t@ype} (
15   lref: lin_buffer a, data: a): lin_buffer a = let
16   val () = conats_atomref_update (lref, data)
17   in
18   lref
19   end
20
21 fun lin_buffer_get {a:t@ype} (
22   lref: lin_buffer a): (lin_buffer a, a) = let
23   val v = conats_atomref_get lref
24   in
25   (lref, v)
26   end
27 end

```

Three functions *conats\_atomref\_create*, *conats\_atomref\_update*, and *conats\_atomref\_get* are declared in `conats.sats`. Intuitively, they are used for creating a mutable object whose content can be accessed in an atomic manner.

In our example, we only need a linear buffer whose content is an integer. The following code defines the type *demo\_buffer* for such linear buffer and some auxiliary functions for accessing it.

```

1 // Define linear integer buffer for demonstration.
2 viewtypedef demo_buffer = lin_buffer int
3
4 fun demo_buffer_isful (buf: demo_buffer): (demo_buffer, bool) = let
5   val (buf, len) = lin_buffer_get (buf)
6   in
7   (buf, len > 0) // Assume the buffer can only hold 1 elements.
8   end
9
10 fun demo_buffer_isnil (buf: demo_buffer): (demo_buffer, bool) = let
11   val (buf, len) = lin_buffer_get (buf)
12   in
13   (buf, len <= 0)
14   end
15
16 fun demo_buffer_insert (buf: demo_buffer): demo_buffer = let
17   val (buf, len) = lin_buffer_get (buf)
18   val buf = lin_buffer_update (buf, len + 1)
19   in
20   buf
21   end
22
23 fun demo_buffer_takeout (buf: demo_buffer): demo_buffer = let
24   val (buf, len) = lin_buffer_get (buf)
25   val buf = lin_buffer_update (buf, len - 1)
26   in
27   buf
28   end

```

One thing worth mentioning is the number 1 we choose as the capacity of the virtual buffer shared by producer and consumer. In reality, a shared buffer may have a large capacity. But a big number may cause model checking not to be able to detect the potential bugs. Arguably, if our implementation is correct for a small capacity of shared buffer, it has better chances to be correct as well for large capacity.

Now we can create the linear buffer holding integer and then put it into a shared object which can be accessed by multiple threads. The corresponding code is shown below.

```

1 // Create a buffer for model construction.
2 val db: demo_buffer = lin_buffer_create (0)
3
4 // Turn a linear buffer into a shared buffer.
5 val s = conats_shared_create {demo_buffer} (db)

```

`conats_shared_create` is a function declared in `conats.sats`, whose semantics is about creating an shared object protecting its content via mutex and condition variables.

We now give out the code for producer and consumer. For the purpose of model checking, *producer* is actually a function which keeps increasing the counter inside the linear buffer whenever possible. If the capacity is reached, the producer would wait until the consumer takes out (by decreasing the counter) something out of the buffer. The same idea applies to the *consumer* functions. Notably, both *producer* and *consumer* would wake up the potentially waiting counterpart by sending a signal.

```

1 // Keep adding elements into buffer.
2 fun producer (x: int):<fun1> void = let
3   val db = conats_shared_acquire (s)
4
5   fun insert (db: demo_buffer):<cloref1> demo_buffer = let
6     val (db, isful) = demo_buffer_isful (db)
7   in
8     if isful then let
9       val db = conats_shared_condwait (s, db)
10    in
11      insert (db)
12    end else let
13      val (db, isnil) = demo_buffer_isnil (db)
14      val db = demo_buffer_insert (db)
15    in
16      if isnil then conats_shared_signal (s, db)
17      else db
18    end
19  end
20
21  val db = insert (db)
22  val () = conats_shared_release (s, db);
23 in
24  producer (x)
25 end
26
27 // Keep removing elements from buffer.
28 fun consumer (x: int):<fun1> void = let
29   val db = conats_shared_acquire (s)
30
31   fun takeout (db: demo_buffer):<cloref1> demo_buffer = let
32     val (db, isnil) = demo_buffer_isnil (db)
33   in
34     if isnil then let
35       val db = conats_shared_condwait (s, db)
36     in
37       takeout (db)
38     end else let
39       val (db, isful) = demo_buffer_isful (db)
40       val db = demo_buffer_takeout (db)
41     in
42       if isful then let
43         // Omitting the following would cause deadlock

```

```

44     // val db = conats_shared_signal (s, db)
45     in db end
46     else db
47     end
48 end
49
50 val db = takeout (db)
51 val () = conats_shared_release (s, db);
52 in
53     consumer (x)
54 end

```

Due to the usage of linear type of ATS, ATS compiler would complain if a programmer forgets to call *conats\_shared\_acquire* to acquire the mutex (which is inside the shared object) before updating the counter, or *conats\_shared\_release* to release the mutex. However, type checking won't be able to detect the potential deadlock if the producer or consumer doesn't call the *conats\_shared\_signal* function.

Model checking can help detect the aforementioned bug. However, unlike type checking, model checking can only be applied to a runnable program instead of a collection of functions. Therefore we set up the environment as follows so that we have a complete model. The model consists of two threads, one for producer and one for consumer. The *conats\_tid\_allocate* and *conats\_thread\_create* functions are provided by `conats.sats`. Intuitively, they are used for allocating thread id and creating new thread with a given function.

```

1  val tid1 = conats_tid_allocate ()
2  val tid2 = conats_tid_allocate ()
3
4  val () = conats_thread_create(producer, 0, tid1)
5  val () = conats_thread_create(consumer, 0, tid2)

```

Since model checking allows us to verify various properties of a program, we specify as follows that we want to verify that our program does not have deadlock.

```

1  %{$
2  #assert main deadlockfree;
3  %}

```

So far we have implemented the producer-consumer problem. With the appropriate implementations of functions from `conats.sats`, we can compile and run the ATS program. Due to the nondeterminism caused by concurrency, the potential deadlock may not happen during several runnings. But with model checking, we are guaranteed that there is no deadlock if our implementation can pass the model checking.

The model checking process goes as follows. We build a tool, which is able to extract a model from the ATS program given above. Currently, the extracted model is in the modeling language CSP#. We then use the state-of-art model checker [PAT](#) to check the generated model. To ease the whole process, we set up a website for readers to try this methodology on-line: [Model Checking ATS](#). The aforementioned example can be found under the name "16\_reader\_writer.dats" in the dropdown list "Select ATS Example". We are working on building tools to better relate the model checking result (counterexample) to the original ATS program. However, it's still quite informative just by inspecting the current result of the model checker since the extracted model in CSP# is quite readable. As for the example, if we omit *conats\_shared\_signal* in *consumer*, model checking would give out the following result including the trace leading to the deadlock. (We omit the detail of the trace here for clarity purpose.)

```

=====
Assertion: main() deadlockfree
*****Verification Result*****
The Assertion (main() deadlockfree) is NOT valid.
The following trace leads to a deadlock situation.
<init -> main_init -> main61_id_s1.0 -> lin_buffer_create_63_s1.0 -> main61_id_s2.0 -> .....

```

```

*****Verification Setting*****
Admissible Behavior: All
Search Engine: Shortest Witness Trace using Breadth First Search
System Abstraction: False

*****Verification Statistics*****
Visited States:2392
Total Transitions:4588
Time Used:0.3925891s
Estimated Memory Used:24059.904KB

```

Next we will illustrate more features of this methodology of combining type checking of ATS programming language with model checking technique to verify properties of concurrent programs.

## 5.2 Table of Contents

### 5.2.1 Ghosts for Model Checking

ATS programming language allows programmers to write code constructing proofs along side with operational code which does the actual computation. The code for proof is seen only by the type checker of ATS and is erased completely before the compiler of ATS starts generating executable programs. We extend such concept by allowing programmers to write code related to model checking. Such code can be seen by the type checker of ATS, which facilitates the type checking process, and will be erased (just like proof code) before the compiler of ATS starts generating executable programs. We call such code ghost code. Ghost code comes in various forms: Ghost function, Ghost value, Ghost type, and etc. We will talk about them in details when meeting related examples.

Both operational code and ghost code have semantics to model checker we build. On the other hand, proof code is neglected by the model checker. The simplest but maybe the most stimulating example of ghost code is *mc\_assert* as shown below.

```

1 fun foo (): void = let
2   prval () = mc_assert (false)
3   in
4   end
5
6 val () = foo ()

```

The executable generated by ATS compiler from the code above simply does nothing since the only content of the function *foo* is a ghost function *mc\_assert*, which is only meaningful to our model checker. Intuitively, we can view *mc\_assert* as assertions normally used by programmers for runtime checking, but this time it's used at the stage of model checking.

Feel free to try the example using our [online model checker](#). To use the model checker, we need to append the following code to indicate that we want to verify that *mc\_assert* succeeds.

```

1 %{$
2 #assert main |= G sys_assertion;
3 %}

```

The complete code (with some routine code for header files) for the example can be found here [demo\\_01\\_mc\\_assert.dats](#)

With *mc\_assert*, we can use model checking as a test facility as shown in the following example.

```

1  typedef Int = [x:int] int x
2
3  fun sum {x:nat} (x: int x): Int = x * (x + 1) / 2
4
5  fun sum2 {x:nat}.<x>.(x: int x):<fun> Int =
6    if x > 0 then x + sum2 (x - 1)
7    else 0
8
9  val n = 3
10 val ret = sum (n)
11 prval mc_ret = sum2 (n)
12 prval () = mc_assert (ret = mc_ret)

```

In the example above, we have two versions of the summation function. We use model checking to verify that they generate the same output. *mc\_ret* is a ghost value (declared via the keyword *prval*), and accordingly the executable generated by ATS compiler doesn't contain the invocation of *sum2*. However *sum2* itself is not a ghost function, and it can be used at runtime if programmers choose to. (Just for technical detail, currently only pure functions in ATS can be used to calculate ghost values. That's why the type of *sum2* is a little bit verbose, since we prove that *sum2* always terminates.)

One thing worth mentioning is the type of *mc\_assert* as shown below:

```
prfun mc_assert {b: bool} (x: bool b):<fun> [b == true] void
```

The type of the argument of *mc\_assert* depends on a boolean value *b* in the statics of ATS. *mc\_assert* assures the type system that *b* is true and as such the type checking of the whole program may be established successfully. Also the validity of such assertion is checking during the process of model checking.

## 5.2.2 Towards Concurrent Program

The type system of ATS consists of both dependent types and linear types. Please refer to ATS' documentation for its application in constructing verifiably correct sequential program. Linear types are of help to certain extent in ensuring safety properties about mutual exclusion. However, in general, the type system of ATS has difficulty in specifying properties of concurrent programs, e.g. invariants of objects across threads, absence of deadlock, liveness properties, and etc. Such incapability triggers our research in incorporating model checking techniques into the verification process of ATS program.

### Primitives for Concurrent Programming

But first of all, programmers have to be able to write concurrent programs in ATS. Currently, we add a set of primitives into ATS programming language to support concurrent programming in a style similar to using *pthread*. In the near future, we will add more primitives (e.g. *channel*, *future*) supporting different styles of concurrent programming.

The declarations of these primitives can be found in *conats.sats*. In the *Tutorial*, we already use *conats\_shared\_create*, *conats\_shared\_acquire*, *conats\_shared\_condwait*, *conats\_shared\_signal*, and *conats\_shared\_release*, whose types are declared as follows:

```

1  abstype shared_t (viewt@type, int)
2  typedef shared (a:viewt@type) = shared_t (a, 1)
3
4  fun conats_shared_create {a: viewt@type} (ele: a): shared (a)
5
6  fun conats_shared_acquire {a: viewt@type} {n:pos} (s: shared_t (a, n)): a
7  fun conats_shared_release {a: viewt@type} {n:pos} (s: shared_t (a, n), ele: a): void
8

```

```

9 fun conats_shared_signal {a: viewt@ype} (s: shared (a), ele: a): a
10 fun conats_shared_condwait {a: viewt@ype} (s: shared (a), ele: a): a

```

*conats\_shared\_create* creates a shared object (of type *shared a*) holding a linear buffer. The concept of shared object is similar to that of monitor<sup>2</sup>. The types of the aforementioned functions guarantee they are invoked appropriately in a non-object-oriented language like ATS. (E.g. we have to call *conats\_shared\_acquire* before invoking *conats\_shared\_condwait* and *conats\_shared\_signal*. From the perspective of pthread programming, a shared object consists of a mutex and a condition variable working together for synchronization purpose.

However, sometimes it's not enough for a shared object to contain just one condition variable. For example, in *Tutorial*, the shared object has only one condition variable, which is used to indicate two situations: buffer is full or empty. The example has no deadlock because it only has one producer and one consumer. Simply adding one more consumer to the example would lead to potential deadlock. The complete code can be downloaded here [16\\_1\\_producer\\_consumer\\_m\\_1.dats](#). You can also find this example at our website [Model Checking ATS](#). One example of the potential deadlock can be thought as follows: Initially, the shared buffer is empty. Two consumers come and wait on the shared object. The producer comes and puts one element into the buffer and then wake up one consumer. However, the newly active consumer doesn't execute instantly. Instead, the producer comes again, tries to put another element into the buffer, and has to wait on the shared object since the capacity of the buffer is 1. Then the newly active consumer gets one element out of the buffer and wakes up another consumer. At this moment, the buffer is empty, the producer is waiting, and two consumers won't signal the shared buffer any more. And this leads to a deadlock. The counterexample found by the model checker by Breadth First Search confirms with our speculation.

To solve this problem, we also provide another version of shared object which contains multiple condition variables. The types of related functions are shown below.

```

1  abstype shared_t (viewt@ype, int)
2
3  fun conats_sharedn_create {a: viewt@ype} {n:pos} (ele: a, n: int n): shared_t (a, n)
4
5  fun conats_sharedn_signal {a: viewt@ype} {i,n:nat | i < n} (s: shared_t (a, n), i: int i, ele: a):
6
7  fun conats_sharedn_condwait {a: viewt@ype} {i,n:nat | i < n} (s: shared_t (a, n), i: int i, ele: a)

```

With such shared object, we can now set up two condition variables handling both full and empty buffers separately. The complete code can be download here [16\\_2\\_producer\\_consumer\\_m\\_1\\_2cond.dats](#). A snippet of code for the producer is shown below.

```

1  // Keep adding elements into buffer.
2  fun producer (x: int):<fun1> void = let
3    val db = conats_shared_acquire (s)
4
5    fun insert (db: demo_buffer):<cloref1> demo_buffer = let
6      val (db, isful) = demo_buffer_isful (db)
7    in
8      if isful then let
9        val db = conats_sharedn_condwait (s, NOTEMP, db)
10     in
11       insert (db)
12     end else let
13       val (db, isnil) = demo_buffer_isnil (db)
14       val db = demo_buffer_insert (db)
15     in
16       if isnil then conats_sharedn_signal (s, NOTFUL, db)
17       else db
18     end
19   end

```

<sup>2</sup> [http://en.wikipedia.org/wiki/Monitor\\_%28synchronization%29](http://en.wikipedia.org/wiki/Monitor_%28synchronization%29)

```

20
21   val db = insert (db)
22   val () = conats_shared_release (s, db);
23   in
24   producer (x)
25   end

```

In this implementation, producer only signals the condition variable when the buffer is actually empty at that moment. This would lead to the missing of signal if we have multiple producers and consumers. (Please refer to section 8.2.2 of <sup>3</sup>.) In the example `16_3_producer_consumer_m_m_signal.dats`, there is two producers, each of which inserts only one element, and two consumers, each of which takes out one element. And the problem of miss signal would lead to deadlock. The model checker would demonstrate this by a counterexample. One remedy is to signal the condition variable every time. The other is to use `conats_sharedn_wait` instead of `conats_sharedn_signal`.

## Bibliography

### 5.2.3 Accessing Global Ghost Variables

#### Global ghost variables

To facilitate programmers to state the properties across the boundary of threads, we provide the concept of global ghost variables. (As being ghost, any flow of data from such variables to the operational state of the program is forbidden.) Programmers are also required to give identities to ghost variables, which provides a way to bridge model checking and type checking for program development. The following example demonstrates this concept.

```

1   stacst sid_init: sid
2   extern val mc_init: mc_gv_t sid_init
3
4   fun exec (x: int): void = let
5
6     fun foo {init: pos}(pf: int_value_of (sid_init, init) | x: int): int = x
7
8     prval (pf | init) = mc_get_int (mc_init)
9
10    // mc_assert cannot be omitted though it is ghost code.
11    prval () = mc_assert (init > 0)
12
13    val _ = foo (pf | x)
14  in
15  end
16
17  val tid1 = conats_tid_allocate ()
18
19  val () = conats_thread_create(exec, 0, tid1)
20
21  prval () = mc_set_int (mc_init, 1)
22
23
24  %{
25  // #assert main deadlockfree;
26
27  #assert main |= G sys_assertion;
28
29  %}

```

<sup>3</sup> The Art of Multiprocessor Programming

We explain this example in details as follows:

```
stacst sid_init: sid
extern val mc_init: mc_gv_t sid_init
```

*stacst* declares a constant *sid\_init* in the *statics* of ATS. This constant serves as the identifier of a ghost variable. *extern val* declares a value *mc\_init*, which is the counterpart of *sid\_init* in the dynamics of ATS. (In the future, only *stacst* is needed after we simplify the model generation process.)

```
fun foo {init: pos}(pf: int_value_of (sid_init, init) | x: int): int = x
```

The type of function *foo* states that a proof that ghost variable *sid\_init* used to be positive is needed to invoke the function. Therefore, if we erase the ghost code *mc\_assert* from the following code

```
fun exec (x: int): void = let

  fun foo {init: pos}(pf: int_value_of (sid_init, init) | x: int): int = x

  prval (pf | init) = mc_get_int (mc_init)

  // mc_assert cannot be omitted though it is ghost code.
  prval () = mc_assert (init > 0)

  val _ = foo (pf | x)
in
end
```

the type checker of ATS would complain that there exists unsolved constraint in the type checking process. From the example, we can see that a set of well designed interfaces for functions can force programmers to incorporate model checking method during the development process. The complete program can be downloaded here `24_global_ghost_variable.dats`. Because of the following code

```
val tid1 = conats_tid_allocate ()

val () = conats_thread_create(exec, 0, tid1)

prval () = mc_set_int (mc_init, 1)
```

The program stands a chance of making the *mc\_assert* (in function *exec*) to fail since the ghost variable is set to 1 after creating a new thread to execute function *exec*. The model checking process can help us detect such problem.

### Atomicity in ghost code

For specification purpose, sometimes it's necessary to group several ghost code into one atomic step. We provide two ghost primitives *mc\_atomic\_start* and *mc\_atomic\_end* to mark the scope for an atomic step, which we call an atomic scope, consisting of both ghost code and operational code. The following program shows the usage of the atomic step.

```
1  stacst mid: sid
2
3  extern val mc_m: mc_gv_t mid
4
5  fun foo1 (): void = let
6    prval () = mc_atomic_start()
7    prval () = mc_set_int (mc_m, 3)
8    prval () = mc_set_int (mc_m, 4)
9    prval () = mc_atomic_end()
10   in
11   end
```

```

12
13 fun foo2 (x: int): void = let
14     prval (pf | x) = mc_get_int (mc_m)
15     prval () = mc_assert (x <> 3)
16 in
17 end
18
19 val tid1 = conats_tid_allocate ()
20
21 val () = conats_thread_create(foo2, 0, tid1)
22
23 val () = fool ()

```

The *mc\_assert* in function *foo2* succeeds because the state in which the ghost variable *mc\_m* is set to 3 is not observable by other threads. The complete code can be downloaded here [18\\_atomic\\_opr.dats](#).

Currently, it's programmers' responsibility to make sure that only primitives for accessing ghost variables (*mc\_set\_int*, *mc\_get\_int*) and primitives for accessing global references (*conats\_atomref\_update*, *conats\_atomref\_get*, *conats\_atomarrayref\_update*, *conats\_atomarrayref\_get*) can appear in an atomic scope. Also they have to make sure that an atomic scope can contain at most one operational primitive for accessing global reference while there's no limit for the number of primitives for accessing ghost variables.

## 5.2.4 Virtual Lock

Improper handling of resources (e.g. memory) in a program may lead to various bugs (e.g. memroy leak) in sequential programs. The problem gets even worse when entering the concurrent domain, in which simultaneous access to shared resource by multiple threads is feasible. One example is that we may lose the integrity of data if two threads are using a shared memory to transfer data. Techniques for solving this problem generally rely on mutual exclusion principles to control access to shared resources. Mutual exclusion introduces a measure of synchronization, but with the cost of losing efficiency. With a deliberate design, sometimes we can remove the need for synchronization while maintaining the desired property of mutual exclusion. Simpson's four-slot fully asynchronous communication mechanism<sup>4</sup> demonstrates such idea. However, it's very difficult to verify that the deemed mutual exclusion property actually holds in the design. To tackle this problem, we provide two primitives supporting the concept of "virtual lock" to allow programmers to specify assumptions of mutual exclusion to various granularities according to their design. And such assumption can then be verified by our model checker.

Let's illustrate the usage of "virtual lock" using the following example of two-slot mechanism. Consider the scenario in which one writer and one reader try to communicate via a shared resource consisting of multiple memory regions (two in this example). Due to hardware constraint, access to each memory region cannot be done atomically. Therefore, reader may get inconsistent data if writer is writing the same region at the same time. The following code shows the proposed types for the shared resource (*dataslots\_t*) as well as the interfaces for accessing it.

```

1  abstype dataslots_t (t@type, int)
2
3  absviewtype own_slot_vt (int)
4
5  fun dataslots_create {a:t@type} {x:pos} (
6     x: int x, v: a): dataslots_t (a, x)
7
8  fun dataslots_update {a:t@type} {x,i:nat | i < x}
9     ( vpf: own_slot_vt (i)
10     | slots: dataslots_t (a, x), i: int i, v: a
11     ): (own_slot_vt i | void)
12
13 fun dataslots_get {a:t@type} {x,i:nat | i < x}

```

<sup>4</sup> H.R. Simpson, Four-slot fully asynchronous communication mechanism

```

14   ( vpf: own_slot_vt (i)
15     | slots: dataslots_t (a, x), i: int i
16   ): (own_slot_vt i | a)

```

The usage of linear type *own\_slot\_vt* states clearly that *dataslots\_update* and *dataslots\_get* require mutual exclusion on the memory region to be accessed. Normally, programmers ensure such property by the usage of synchronization primitives (e.g. mutex). However, in the following code, we try to gain mutual exclusion by the usage of a few global variables the access for which is atomic. The code is shown below.

```

1   typedef data_t = dataslots_t (int, 2)
2   val data: data_t = dataslots_create (2, 0)
3
4   typedef int2 = [i: int | i >= 0 && i <= 1] int i
5
6   // control variables
7   val latest = conats_atomref_create {int2} (0)
8
9   fun write (item: int): void = let
10    val index = 1 - conats_atomref_get (latest)
11
12    prval vpf = mc_acquire_ownership (index)
13    val (vpf | _) = dataslots_update (vpf | data, index, item)
14    prval () = mc_release_ownership (vpf)
15
16    val () = conats_atomref_update (latest, index)
17  in
18  end
19
20  fun read (): int = let
21    val index = conats_atomref_get (latest)
22
23    prval vpf = mc_acquire_ownership (index)
24    val (vpf | item) = dataslots_get (vpf | data, index)
25    prval () = mc_release_ownership (vpf)
26  in
27    item
28  end

```

In the example, the shared resource (*data\_t*) contains two regions (slots). *latest* is a global reference for an integer, which is created by the primitive *conats\_atomref\_create*. (Primitives *conats\_atomref\_create*, *conats\_atomref\_get*, and *conats\_atomref\_update* are provided as an extension to ATS to support concurrent programming.) To pass the type checking of ATS, we use two functions *mc\_acquire\_ownership* and *mc\_release\_ownership* to generate and destroy the linear ghost value (*vpf*), which serves as the warranty for mutual exclusion. *mc\_acquire\_ownership* and *mc\_release\_ownership* are not primitives. Instead, they are user-defined ghost functions. Their implementation is shown below.

```

1   prfun mc_acquire_ownership .<>. {i: nat}
2     (i: int i): own_slot_vt (i) = mc_vlock_get (i, 0, 1, 1)
3
4   prfun mc_release_ownership .<>. {i: nat}
5     (vpf: own_slot_vt (i)): void = mc_vlock_put (vpf)

```

The two ghost functions are built upon two primitives *mc\_vlock\_put* and *mc\_vlock\_get*. Intuitively, *mc\_vlock\_get* (*x*, *y*, *a*, *b*) indicates the acquisition of a *virtual lock* covering a rectangle with (*x*, *y*) as the upper left corner, *a* as the width (x-axis), and *b* as the height (y-axis), and *mc\_vlock\_put* indicates the release of the lock. And our model checker would check that under no circumstances would two threads try to acquire two virtual locks covering overlapping areas simultaneously. And this serves as the verification of mutual exclusion. To model checking the example, we would need to add the following code to implement those interfaces for accessing shared resource.

```
1 fun dataslots_create {a:t@type} {x:pos} (  
2   x: int x, v: a): dataslots_t (a, x) =  
3   conats_atomarrayref_create {a} (x, v)  
4  
5 fun dataslots_update {a:t@type} {x,i:nat | i < x}  
6   ( vpf: own_slot_vt (i)  
7   | slots: dataslots_t (a, x), i: int i, v: a  
8   ): (own_slot_vt i | void) = let  
9   val () = conats_atomarrayref_update (slots, i, v)  
10  in  
11   (vpf | ())  
12  end  
13  
14 fun dataslots_get {a:t@type} {x,i:nat | i < x}  
15   ( vpf: own_slot_vt (i)  
16   | slots: dataslots_t (a, x), i: int i  
17   ): (own_slot_vt i | a) = let  
18   val v = conats_atomarrayref_get (slots, i)  
19  in  
20   (vpf | v)  
21  end
```

In the aforementioned code, this implementation is actually based on the primitives for creating and accessing array. This is not necessary if our focus is to verify the validity of mutual exclusion.

The complete code can be downloaded here [20\\_1\\_two\\_slot\\_acm.dats](#). Without much thinking, we know that this implementation cannot pass model checking since writer just switches between two slots. Going further, the implementation ([20\\_2\\_three\\_slot\\_acm.dats](#)) using three slots doesn't work either. And based on the implementation using four slots ([20\\_3\\_four\\_slot\\_acm.dats](#)), we verify that Simpson's four-slot asynchronous mechanism possesses the acclaimed mutual exclusion property.

## Bibliography

### 5.3 Bibliography

---

## Indices and tables

---

- `genindex`
- `modindex`
- `search`



## V

Virtual Lock, 32