

INSTITUT FÜR INFORMATIK

SyncCharts in C

Reinhard von Hanxleden

Bericht Nr. 0910

May 2009, Revised^a September 2009 (Draft)

Last compiled: 2009-10-07, 2:31 hrs (CET)

^aThe original, unrevised report is available at <http://rtsys.informatik.uni-kiel.de/~biblio/downloads/papers/report-0910-unrevised.pdf>



CHRISTIAN-ALBRECHTS-UNIVERSITÄT
ZU KIEL

Institut für Informatik der
Christian-Albrechts-Universität zu Kiel
Olshausenstr. 40
D – 24098 Kiel

SyncCharts in C

Reinhard von Hanxleden

Bericht Nr. 0910
May 2009, Revised^a September 2009 (**Draft**)
Last compiled: 2009-10-07, 2:31 hrs (CET)

e-mail:
rvh@informatik.uni-kiel.de

^aThe original, unrevised report is available at <http://rtsys.informatik.uni-kiel.de/~biblio/downloads/papers/report-0910-unrevised.pdf>

Technical Report

Abstract

Statecharts are a well-established visual formalism for the description of reactive real-time systems. The SyncCharts dialect of Statecharts, which builds on the synchrony hypothesis, has a sound formal basis and ensures deterministic behavior. This report presents SyncCharts in C (SC), an approach on how to seamlessly and efficiently embed SyncCharts constructs into a conventional imperative programming language. SC offers deterministic concurrency and preemption via a simulation of multi-threading, inspired by reactive processing.

SC can be used as a regular programming language, requiring just a C compiler; no special tools or hardware are needed. However SC's conciseness, completeness and semantic closeness to SyncCharts make it an attractive candidate in a number of other scenarios: 1) as an intermediate target language for synthesizing graphical SyncChart models into executable code, in a more traceable manner than the traditional path through Esterel; 2) as instruction set architecture for programming precision timed (PRET) or reactive architectures; or 3) as a virtual machine instruction set. A reference implementation of SC, based on light-weight C macros, is available as open source code.

Key words: SyncCharts, Statecharts, Esterel, synchronous programming, code synthesis, model-based design

Note: An abridged version of the original report appeared elsewhere [32].

Contents

1	Introduction	1
2	Introductory Examples	3
2.1	Reactive Control in SC—The PCO Example	3
2.2	Signals in SC—The grcbal3 Example	7
3	A Tour of SC	12
3.1	The SC Programming Model	12
3.1.1	Synchronous threading	12
3.1.2	Signals	14
3.2	Multithreading Simulation	15
3.2.1	Coarse program counters	15
3.2.2	The dispatcher	15
3.2.3	Thread and label structuring	16
3.2.4	Thread scheduling	20
3.3	SC Operators	24
3.3.1	SC thread handling operators	24
3.3.2	SC signal operators	26
3.3.3	SC sequential control operators	28
3.3.4	An example of expanded macros—ABRO	30
3.4	SC Structure	30
3.4.1	Program files	30
3.4.2	Functions	33
3.4.3	Types	34
3.4.4	Variables	35
4	Further Examples	36
4.1	Count2Suspend	36
4.2	Exits	36
4.3	FilteredSR	40
4.4	Shifter3	40
4.5	PreAndSuspend	42
4.6	Reincarnation	43
4.7	PrimeFactor	44
5	Related Work	46

6	Experimental Results	50
6.1	Conciseness of SC, Code Size	50
6.2	SC Performance	50
7	Conclusions and Outlook	53
A	The SC files	58

List of Figures

2.1	The PCO (Producer-Consumer-Observer) example.	4
2.2	The <code>grcbal3</code> example.	8
3.1	The status of the whole program	13
3.2	Execution status of a single thread	13
3.3	The ABRO example.	17
3.4	The <code>SurfDepth</code> example.	19
3.5	ABRO tick function after macro expansion	30
3.6	ABRO tick function after macro expansion, annotated and reformatted.	31
3.7	ABRO tick function after macro expansion, annotated and stripped of empty statements.	32
4.1	The <code>Count2Suspend</code> example.	37
4.2	The <code>Exits</code> example. See Fig. 4.3 for alternative tick functions.	38
4.3	Alternative variants for the SC tick function of the <code>Exits</code> example	39
4.4	The <code>FilteredSR</code> example.	40
4.5	The <code>Shifter3</code> example.	41
4.6	The <code>PreAndSuspend</code> example.	42
4.7	The <code>Reincarnation</code> example.	43
4.8	The <code>PrimeFactor</code> example.	44
6.1	Comparison of SC with two code synthesis variants of Esterel Studio.	51

List of Tables

2.1	SC thread operators	6
2.2	SC signal operators and sequential control operators	9

Listings, Outside Figures

3.1	<code>selectCid()</code> : Computation of id of thread to be dispatched (from <code>sc.c</code>)	15
3.2	<code>dispatch()</code> : Variable definitions for the dispatcher (from <code>sc.h</code>)	16
A.1	The header file <code>sc.h</code>	58
A.2	The main program file <code>sc.c</code>	64
A.3	The Makefile	64

Chapter 1

Introduction

The control flow of reactive systems typically entails not just the sequential control flow found in traditional programming languages, such as conditionals and loops, but also exhibits concurrency and preemption. This reactive control flow is naturally expressed by the Statechart formalism introduced by David Harel [13], which extends classical finite state machines by concurrency and hierarchy/preemption. These extensions allow to keep descriptions compact and avoid the classical state explosion problem.

The graphical Statechart formalism has been originally developed to let application experts precisely describe the behavior desired for an application. Its visual nature makes this formalism accessible to non-computer scientists, without the need to be versed in a traditional programming language. However, beyond this visual *syntax*, Statecharts offer important *concepts* that can be expressed in non-visual languages as well, such as the concepts of state-based control flow, hierarchy, concurrency, and its model of time. This model of computation (MoC) of Statecharts offers a powerful abstraction mechanism compared to classical programming models. For this reason, Statechart models are typically viewed as more abstract than, say, a program written in C.

A typical design flow may start with a graphical modeling tool, which synthesizes a Statechart model into a C program, which is further compiled into some executable. However, it is also quite common to bypass the visual modeling step. Just as the code generator of a modeling tool is able to express the Statechart MoC in a C program, so it is possible for a human programmer to express Statechart behavior as a C program [29, 34]. This does not offer the visual appeal of graphical Statecharts, but has other advantages:

- no need for a modeling tool,
- high portability, and
- seamless integration with a fully featured, widely used programming language, including the type system, expression handling, control flow, access to low-level I/O, pre-processors, etc.

Even if one assumes a design flow that starts at a graphical modeling tool that supports Statecharts, it is of interest how Statechart behavior can be expressed concisely in a traditional programming language. For a number of reasons, we would like to be able to generate code that preserves the structure of the graphical model:

- it simplifies the development of the code synthesizer of the modeling tool;

- it facilitates back-annotations from the executable code into the graphical model, which allows visual animations of the running code and allows to set break points in the model; and
- it simplifies code certification for safety-critical embedded systems.

This report describes *SyncCharts in C* (SC), which is a light-weight approach to express SyncCharts [3] in C programs. SC combines the formal soundness of SyncCharts, including deterministic concurrency and preemption, with the efficiency and wide support for the C language. The main idea of SC is to emulate multi-threading, and is inspired by reactive processing [33]. As we do not have direct access to the program counter at the C language level, we keep track of individual threads via state labels, implemented as usual C program labels. These labels can also be viewed as continuations [5], or coroutine [9, 15] re-entry points. Precedence among transitions, respecting strong/weak abortions and hierarchy, and the adherence to signal dependencies are achieved by checking transition triggers in the proper order as well as assigning appropriate thread ids and priorities.

To write and execute an SC application requires neither specific tools nor special execution platforms, although both may support this concept further. All that is needed to get started is an understanding of SyncCharts (see *e. g.* the tutorial provided by André [3]), a C compiler, and the SC files. The SC files consist of one header file (`sc.h`), to be included by the application code, and one C-file (`sc.c`), to be linked in by the application. They are open source and available for free download¹.

As the name suggests, SC has been developed with the SyncCharts execution model in mind. However, SC can also be viewed as a generic approach for programming light-weight, deterministic concurrent programs in C, without using SyncChart-specifics such as (valued) signals or (weak) abortions. For example, SC appears to be a suitable candidate for writing concurrent C programs that have predictable functionality and timing on PRET-like architectures [19], without having to resort to low-level synchronization mechanisms based on physical timing characteristics.

In this report, we will first work through two examples that give an overview how SC programmers can implement reactive control and signal-based communication, followed by a full tour of SC in Chapter 3 and further examples in Chapter 4. Chapter 5 discusses related work, experimental results are presented in Chapter 6. The report concludes in Chapter 7. Appendix A lists the SC files.

¹<http://www.informatik.uni-kiel.de/rtsys/sc/>

Chapter 2

Introductory Examples

This chapter gives a first practical introduction to SC by working through two examples. The first presents the fundamental reactive control flow mechanism supported by SC, namely concurrency and preemption. The second example makes use of signal handling and illustrates how SC supports intricate thread inter-dependencies.

2.1 Reactive Control in SC—The PCO Example

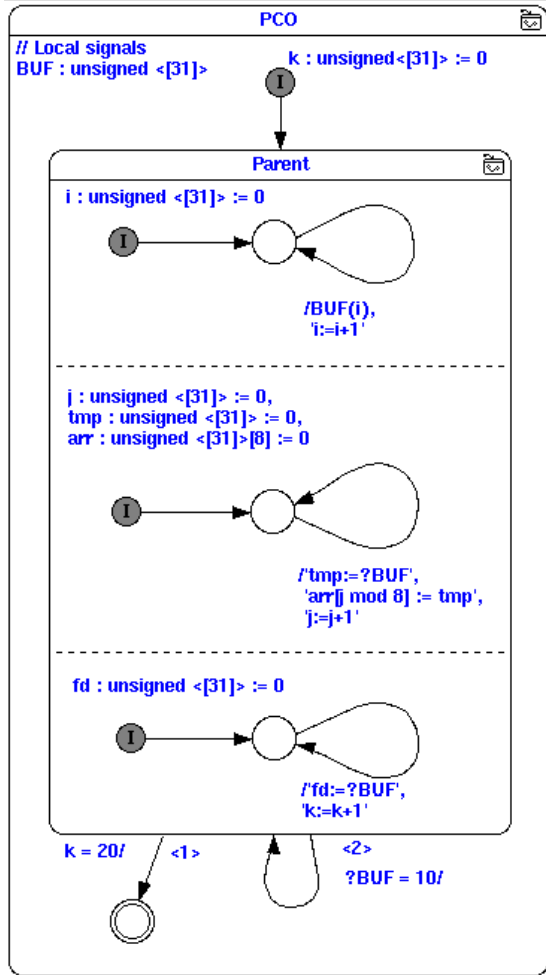
This section covers

- the general structure of SC programs,
- how SC macros are embedded in regular C code,
- the concept of deterministic, label-based simulated multi-threading, and
- deterministic preemptions.

We will illustrate these points with PCO, shown in Fig. 2.1, a simple producer-consumer example with an observer, inspired by Lickly *et al.* [19]. In addition to the original example, PCO also has a parent thread that restarts production/consumption once the buffer has the value 10, and which terminates after 20 iterations.

The SyncCharts version (Fig. 2.1a) shows a **Parent** macrostate, which is an AND (parallel) state that consists of three substates, corresponding to the producer, consumer and observer. Each substate consists of a state with a self-transition, which is triggered unconditionally and performs some action. For example, the producer state writes the current value of *i* into a buffer **BUF**, a *valued signal* in SyncCharts parlance. The consumer state reads the value of **BUF** into some variable **tmp** and then writes **tmp** into an array **arr**. The observer also reads from **BUF**. The **Parent** state re-enters itself when **BUF** has the value 10, and transitions to some final state when **k**, incremented by the observer, has reached the value 20.

Compared to an implementation that would try to achieve the same behavior with, say, Java threads, the interesting aspect of the SyncChart implementation is that the concurrency is deterministic. The three substates of **Parent** execute in lock step, and the SyncCharts semantics requires that in each execution, **BUF** must be written before it is read. Hence, the code generator of EsterelStudio, which generates C code from this (via Esterel), must schedule



(a) SyncChart, produced with EsterelStudio

Producer	Consumer	Observer
<pre>int main() { DEAD (28); volatile unsigned int * buf = (unsigned int*) (0x3F800200); unsigned int i = 0; for (i = 0; ; i++) { DEAD (26); *buf = i; } return 0; }</pre>	<pre>int main() { DEAD (41); volatile unsigned int * buf = (unsigned int*) (0x3F800200); unsigned int i = 0; int arr[8]; for (i=0; i<8; i++) arr[i] = 0; for (i = 0; ; i++) { DEAD (26); register int tmp = *buf; arr[i%8] = tmp; } return 0; }</pre>	<pre>int main() { DEAD (41); volatile unsigned int * buf = (unsigned int*) (0x3F800200); volatile unsigned int * fd = (unsigned int*) (0x8000600); unsigned int i = 0; for (i = 0; ; i++) { DEAD (26); *fd = *buf; } return 0; }</pre>

(b) PRET version, without preemption (from [19])

```

1  #include "sc.h"
2
3  int BUF, fd, i, j, k = 0, tmp, arr [8];
4
5  // ===== MAIN FUNCTION =====
6  int main()
7  {
8      int notDone, init = 1;
9
10     do {
11         notDone = tick(init); // Call tick function
12         //sleep(1);           // Slow down by 1 sec
13         init = 0;
14     } while (notDone);
15     return 0;
16 }
17
18 // ===== TICK FUNCTION =====
19 int tick(int isInit)
20 {
21     TICKSTART(isInit, 1);
22
23     PCO:
24     FORK(Producer, 4);
25     FORK(Consumer, 2);
26     FORK(Observer, 3);
27     FORKE(Parent);
28
29     Producer:
30     for (i = 0; ; i++) {
31         PAUSE;
32         BUF = i;
33     }
34
35     Consumer:
36     for (j = 0; j < 8; j++)
37         arr[j] = 0;
38     for (j = 0; ; j++) {
39         PAUSE;
40         tmp = BUF;
41         arr[j % 8] = tmp;
42     }
43
44     Observer:
45     for ( ; ; ) {
46         PAUSE;
47         fd = BUF;
48         k++;
49     }
50
51     Parent:
52     while (1) {
53         PAUSE;
54         if (k == 20) TRANS(Done);
55         if (BUF == 10) TRANS(PCO);
56     }
57
58     Done:
59     TERM;
60     TICKEND;
61 }

```

(c) Complete SC program

Figure 2.1: The PCO (Producer-Consumer-Observer) example.

the producer before the consumer and the observer. Similarly, the transitions leaving `Parent` have deterministic behavior; in this example, they are so-called *weak abortions*, meaning that the body of the parent gets to finish its current execution before a transition is taken. An implementation with classical Java threads offers none of these assurances. To achieve the same effect would require explicit barrier synchronization. Note also that for example using Java’s `synchronized` to protect access to the shared buffer does not help, as this would only guarantee exclusive access, but no ordering.

One approach suggested recently to enforce this synchronization is to use explicit low-level time-triggered scheduling. The PRET architecture [19] offers a `DEAD` instruction which guarantees a (minimal) delay before a thread proceeds. Fig. 2.1b shows the PRET version of a reduced variant of PCO that does not have preemptions. In this PRET version, the buffer access is coordinated by giving the producer a head start before the consumer and observers (`DEAD 28` vs. `DEAD 41`), and then keeping all three running at the same rate (`DEAD 26`). To guarantee proper synchronization this way requires a timing analysis of the code and the underlying architecture, and the resulting program is fairly non-portable.

The SC version of PCO is shown in Fig. 2.1c. The `main` function contains a `while` loop that calls a `tick` function. This function computes one reaction by simulating all *enabled* threads for one tick. The return value of `tick` indicates whether the program has terminated, *i. e.*, whether all threads have become *disabled*. The `while` loop of `main` continues as long as any thread is still enabled. In this example, a call to `sleep(1)` results in a reaction rate of—approximately—once per second.

The `tick` function consists of regular C code and some macros. These *SC macros* are declared in `sc.h`, included in line 1. An overview of the SC Thread Handling Operators, which perform the multi-threading simulation and form the core of SC, is given in Table 2.1. The remaining SC operators are introduced in Sec. 2.2, Table 2.2. A full discussion of all SC is presented in Sec. 3.3.

The first SC macro used in PCO, `TICKSTART`, performs some book keeping, depending on whether this is the initial tick or not. This is followed by a sequence of `FORK`/`FORKE` macros, which fork off the children of the current thread. The current thread, started when entering `tick`, is the `Main` thread. The forked threads are `Prod`, `Cons`, and `Obs`.

As the forked threads are associated with the `Parent` state of the SyncChart, we will also refer to these as `Parent`’s children; however, the thread that is forking them is the `Main` thread. In this example, the `Main` thread only forks these children, as the `Parent` macrostate is the only macrostate ever entered by `Main`.

Each `FORK` gives a thread its initial priority (here all 0), a starting label, and an id. `FORKE` specifies a priority for the current thread (again 0), a continuation label (`ParentMain`), and the set of children that were just forked. Sets of threads are encoded as a bit vector, `id2b` maps a thread into this vector. This set is needed to properly abort `Main`’s children when `TRANS` is called, see below.

Threads are declared with the `idtype` enumeration type (line 4).

The starting point of each thread is declared with an ordinary C label, named after the thread. This is just a convention; from a C perspective, these labels and the thread names have different name spaces and are different objects: one is a memory address, the other is an enumeration type index.

The code for each thread is regular C code, except that each thread contains a `PAUSE` macro. `PAUSE` indicates that a thread becomes inactive and is ready to relinquish control to

Mnemonic, Operands	Notes
TICKSTART*(<i>isInitial</i> , <i>p</i>)	Start (initial) tick, assign main thread priority <i>p</i> .
TICKEND	Finalize tick, return 1 iff there is still an enabled thread.
PAUSE* ⁺	Deactivate current thread for this tick.
PAUSEG*(<i>l</i>)	Shorthand for PAUSE; GOTO(<i>l</i>).
HALT* ⁺	Shorthand for <i>l</i> : PAUSE; GOTO(<i>l</i>).
TERM*	Terminate current thread.
ABORT(<i>l</i>)	Abort descendant threads.
TRANS(<i>l</i>)	Shorthand for ABORT; GOTO(<i>l</i>).
SUSPEND*(<i>cond</i>)	Suspend (pause) thread and its descendants if <i>cond</i> holds.
SUSPENDGOTO*(<i>l</i>)	Suspend (pause) thread and its descendants, continue at <i>l</i> holds.
FORK(<i>l</i> , <i>p</i>)	FORK creates a thread with start address <i>l</i> and priority <i>p</i> .
FORKE*(<i>l</i>)	FORKE specifies the continuation address <i>l</i> for the spawning thread.
JOINE* ⁺ (<i>l_{else}</i>)	If descendant threads have terminated normally, proceed; else pause, jump to <i>l_{else}</i> .
JOIN* ⁺	Shorthand for <i>l_{else}</i> : JOINE(<i>l_{else}</i>). Waits for descendant threads to terminated normally.
PRIO* ⁺ (<i>p</i>)	Set current thread priority to <i>p</i> .
PPAUSE* ⁺ (<i>p</i>)	Shorthand for PRIO(<i>p</i>); PAUSE (saves one call to dispatcher).
JPPAUSE* ⁺ (<i>p</i> , <i>l_{then}</i> , <i>l_{else}</i>)	Shorthand for JOINE(<i>l</i>); GOTO(<i>l_{then}</i>); <i>l</i> : PRIO(<i>p</i>); PAUSE; GOTO(<i>l_{else}</i>); <i>l_{then}</i> : (saves another call to dispatcher).

Table 2.1: SC thread operators—tick delimiters, fork/join, priority handling, and abortion and suspension. Operators marked with an asterisk* may call the thread dispatcher, *i. e.*, can result in a thread context switch. Operators marked with a plus⁺ automatically generate continuation labels (visible in the program after macro expansion and in execution traces).

the *dispatcher*.

The dispatcher, called by PAUSE, selects a thread for resumption. In PCO the dispatcher selects from the *active* threads, which still have work to do in the current thread, the one with the highest *thread id*. The dispatcher may also consider dynamic priorities, see Sec. 2.2, but in PCO these are all 0. Threads are mapped to their ids with the *ids* array (line 5). The TickEnd thread, which must be present in any SC program and must have the lowest id (0), returns from tick if none of the other threads are active anymore.

Taking a look at the Main thread continuation at the Parent label (line 47), we note that the transitions triggered by inspecting first *k* and then BUF are implemented with a TRANS macro (lines 49 and 51). This macro transfers control to the argument label, and also aborts Parent’s child threads. Finally, TERM terminates the current thread (Main), and TICKEND does last book keeping before leaving tick again.

To summarize, we simulate multi-threading by keeping track of continuation points and calling a dispatcher whenever a context switch might occur. In the example, the dispatcher is called by PAUSE (thread becomes inactive for the current tick), FORKE (children have been created, current thread may have changed priority), and TERM (thread has terminated). The context of a thread is very light-weight: it consists of its id (static), its continuation label (dynamic), and a priority (dynamic). Everything else is shared. The thread id encodes the order in which threads are dispatched. In PCO, the producer has to run before the consumer

and the observer, hence `Prod` gets the highest id, which is 4. For a full discussion of PCO’s precedence constraints, see Sec. 3.2.4, p. 23.

All threads are included in one C tick function, just as for example a SyncChart or Esterel program is usually synthesized into a single reaction function. This makes data sharing and communication trivial (compare for example with the PRET communication in Fig. 2.1b), but limits modularization. This is a consequence of the label-based continuation encoding, since in C, we cannot transfer control to a label across function calls. Alternatives, such as encodings based on `setjmp/longjmp`, would provide more flexibility, but would also incur higher overhead. Note, however, that modularization is still possible insofar as “instantaneous” functionality, without any SC operator that calls the dispatcher, can still be compartmentalized into function calls. This suggests a programming model where the thread structure and their scheduling logic is summarized in a top-level tick function, and thread-local activities and data-intensive computations are modularized as function calls.

2.2 Signals in SC—The `grcbal3` Example

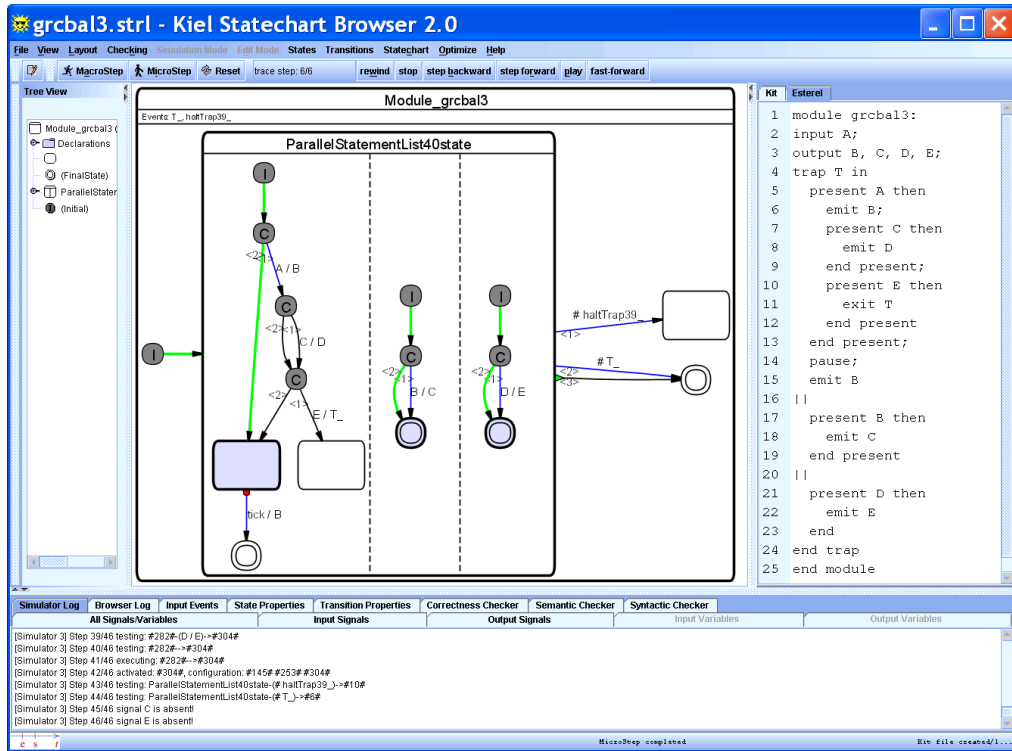
This section covers

- more elaborate thread scheduling via the use of dynamic thread priorities,
- signal handling,
- a synthesis path from Esterel to SC, and
- how SC macros alone suffice to write a tick function.

Again we use an example, `grcbal3`, to illustrate these issues. Originally, this example was programmed in Esterel, and has been presented by Edwards and Zeng in their description of the Columbia Esterel Compiler [10]. Hence the name of the benchmark: GRC is the Graph Code intermediate representation of the CEC, BAL is the Bytcode Assembly Language of a virtual machine (VM) targeted by the CEC. The `grcbal3` Esterel code has been transformed into a SyncChart using KIEL [24]. Fig. 2.2a shows the Esterel version, on the right, with the generated SyncChart, in the midst of an animated simulation—the initial tick has just been executed, with no inputs present.

The Esterel program illustrates the use of signals to synchronize threads. It has an input signal `A` and output signals `B...E`. There are three concurrent threads, which are enclosed in a trap triggered by `T`. Esterel’s trap construct provides exception handling; in the example, the `exit T` statement (line 11) throws the exception. The three threads communicate back and forth via signals; for example, if `A` is present, the first thread emits a `B`, which causes the second thread to emit `C`, which in turn causes the first thread to emit a `D`.

The SyncChart synthesized by KIEL is equivalent to the Esterel version. However, as SyncCharts do not provide traps, they have to be emulated with weak abortions. This translation is always possible, and in `grcbal3` this can be done in a straightforward fashion, via a weak abort triggered by a fresh signal `T_`. The transition that implements this is shown in the lower right of the SyncChart, which leads to a final state (double circle). The #-mark means that the transition is *immediate*, meaning it can be triggered from the initial instant on. Note that the synthesis process produces a superfluous state, reachable via `haltTrap39`—which



(a) Screen shot of KIEL [25], as it synthesizes a SyncChart from the original Esterel code [10]

```

1  TICKSTART(isInit, 1);
2  FORK(T1, 6);
3  FORK(T2, 5);
4  FORK(T3, 3);
5  FORKE(TMmain);
6
7  T1:
8  PRESENTElse(A, T1B);
9  EMIT(B);
10 PRIO(4);
11 PRESENTEmit(C, D);
12 PRIO(2);
13 PRESENTElse(E, T1B);
14 EMIT(T_);
15 TERM;
16 T1B: PAUSE;
17 EMIT(B);
18 TERM;
19
20 T2:
21 PRESENTEmit(B, C);
22 TERM;
23
24 T3:
25 PRESENTEmit(D, E);
26 TERM;
27
28 TMmain:
29 PRESENTElse(T_, TJoin);
30 ABORT;
31 TERM;
32 TJoin: JOINE(TMmain);
33 TICKEND;

```

```

1  TICKSTART(isInit, 1);
2  FORK(T1, 6);
3  FORK(T2, 5);
4  FORK(T3, 3);
5  FORKE(TMmain);
6
7  T1: if (PRESENT(A)) {
8    EMIT(B);
9    PRIO(4);
10   if (PRESENT(C))
11     EMIT(D);
12   PRIO(2);
13   if (PRESENT(E)) {
14     EMIT(T_);
15     TERM; }
16 }
17 PAUSE;
18 EMIT(B);
19 TERM;
20
21 T2: if (PRESENT(B))
22   EMIT(C);
23 TERM;
24
25 T3: if (PRESENT(D))
26   EMIT(E);
27 TERM;
28
29 TMmain: if (PRESENT(T_)) {
30   ABORT;
31   TERM; }
32 JOINE(TMmain);
33 TICKEND;

```

```

1  ===== TICK 0 STARTS, inputs = 01, enabled
2  = 00
3  ===== Inputs (id/name): 0/A
4  ===== Enabled (id/state): <init>
5  FORK: 1/<init> forks 6/T1, active = 0103
6  FORK: 1/<init> forks 5/T2, active = 0143
7  FORK: 1/<init> forks 3/T3, active = 0153
8  FORKE: 1/<init>
9  PRESENT: 6/T1 determines A/0 as present
10 EMIT: 6/T1 emits B/1
11 PRIO: 6/T1 set to priority 4
12 PRESENT: 5/T2 determines B/1 as present
13 EMIT: 5/T2 emits C/2
14 TERM: 5/T2 terminates, enabled = 073
15 PRESENT: 4/.L73 determines C/2 as present
16 EMIT: 4/.L73 emits D/3
17 PRIO: 4/.L73 set to priority 2
18 PRESENT: 3/T3 determines D/3 as present
19 EMIT: 3/T3 emits E/4
20 TERM: 3/T3 terminates, enabled = 017
21 PRESENT: 2/.L76 determines E/4 as present
22 EMIT: 2/.L76 emits T_-/5
23 TERM: 2/.L76 terminates, enabled = 07
24 PRESENT: 1/TMain determines T_-/5 as present
25 ABORT: 1/TMain disables 070, enabled = 03
26 TERM: 1/TMain terminates, enabled = 03
27 ===== TICK 0 terminates after 22 instructions.
28 ===== Enabled (id/state): 0/.L.TICKEND
===== Resulting signals (name/id): 0/A, 1/B,
2/C, 3/D, 4/E, 5/T_, Outputs OK.

```

(d) Example trace

(b) SC tick function (SC operators only)

(c) SC tick function (mixed with standard C)

Figure 2.2: The grcbal3 example.

Mnemonic, Operands	Notes
SIGNAL(S)	Initialize a local signal S .
EMIT(S)	Emit signal S .
SUSTAIN*+(S)	Sustains signal S . Shorthand for l : EMIT(S); PAUSE; GOTO(l).
PRESENT(S)	Evaluates to 1 if S is present, else evaluates to 0.
PRESENTELSE(S, l_{else})	If S is present, proceed normally; else, jump to l_{else} . If PRESENT is available: shorthand for if (!PRESENT(s)) GOTO(l_{else}).
PRESENTEMIT(S, T)	If S is present, emit T . Shorthand for if (PRESENT(S)) EMIT(T).
AWAIT*+(S)	Wait (non-immediately) for signal S . Shorthand for l_{else} : PAUSE; PRESENT(s, l_{else}).
AWAITI*+(S)	Wait (immediately) for signal S . Shorthand for GOTO(l); l_{else} : PAUSE; l : PRESENT(s, l_{else}).
EMITINT(S, val)	Emit valued signal S , of type integer, with value val .
EMITINTMUL(S, val)	Emit valued signal S , of type integer, combined with multiplication, with value val .
VAL(S)	Retrieve value of signal S .
VALREG(S, reg)	Retrieve value of signal S , into register/variable reg .
PRESENTPRE(S, l_{else})	Evaluates to 1 if S was present in previous tick, else evaluates to 0. If S is a signal local to thread t , consider last preceding tick in which t was active, <i>i. e.</i> , not suspended.
PRESENTPREELSE(S, l_{else})	If S was present in previous tick, proceed normally; else, jump to l_{else} . If PRESENTPRE is available: shorthand for if (!PRESENTPRE(s)) GOTO(l_{else}).
VALPRE(S)	Retrieve value of signal S at previous tick.
VALPREREG(S, reg)	Retrieve value of signal S at previous tick, into register/variable reg .
GOTO(l)	Jump to label l .
CALL(l)	Call function l (eg, an on exit function).
RET	Return from function call.
ISAT(id, l_{state}, l)	If thread id is at state l_{state} , then proceed to next instruction (<i>e. g.</i> , an on exit function of associated with id at state l_{state}). Else, jump to label l .
ISATCALL($id, l_{state}, l_{action}$)	Shorthand for ISAT(id, l_{state}, l); CALL(l_{action}); l ;

Table 2.2: SC signal operators (pure signals, valued signals, and accesses to the previous tick) and SC sequential control operators (jumps and exit actions). See Table 2.1 on the asterisk* and plus+ annotations.

is nowhere emitted, hence it can be safely eliminated. This is a result of the general rule for transforming traps, which has to handle nested traps and trap actions [24], and a lack of a subsequent optimization in KIEL that would remove such clearly unreachable states.

Fig. 2.2c shows the tick function of the SC version of `grcbal3`. In addition to the SC concurrency operators already introduced in Sec. 2.1 and Table 2.1, `grcbal3` makes use of SC *signal operators*. An overview of these and some other, sequential control operators is given in Table 2.2.

To better understand this example’s operation, consider also the execution trace shown in Fig. 2.2d. All SC macros (apart from `TICKSTART` and `TICKEND`) log their operation to `stdout` if instructed to do so via a preprocessor directive. The trace illustrates the operation of `grcbal3` in case input signal `A` is present. The first line shows the input signals (`A`) and the enabled threads (initially none) as bit vector, in octal notation with leading 0. `TICKSTART`, `FORK`, and `FORKE` are as explained for the `PCO` example (Sec. 2.1). One difference, however, is that threads `A1`, `A2` and `A3`, which correspond to the three concurrent substates embedded in the macrostate in the SyncChart version, are started with priorities 3, 2, and 1, respectively. This priority is used by the dispatcher, which always resumes the active thread with the highest priority; if there are multiple such threads with the same, highest priority, then the highest thread id decides. In `PCO`, all threads had priority 0, hence there only the thread id matters to the dispatcher.

After `Main` has forked its children, `FORKE` calls the dispatcher, see line 5 in the program, line 7 in the trace. This starts `A1` (thread id 2), as it has the highest priority. `A1` determines `A` as present and emits signal `B`. The `PRIO` directive lowers `A1`’s priority to 2, and calls the dispatcher. Now `A2` (id 3) is started, as it has the same priority as `A1`, but a higher thread id. `A2` determines `B` as present and hence emits `C`. Then the `TERM` operator *terminates* `C`, meaning that it is deactivated (does not resume in the current tick) and disabled (will not be resumed in the next tick). Therefore `TERM` calls the dispatcher, without specifying a continuation label. The set of remaining enabled threads is encoded in a bit vector, see line 13 of the trace. The vector octal 027, binary 10111, has bits 0 (rightmost bit, indicating thread `TickEnd`), 1 (`Main`), 2 (`A1`) and 4 (`A3`) set.

In this fashion, control is passed back and forth between `Parent`’s children until they have all have completed their tick, and the `Main` thread, running at priority 0, resumes; see line 23 of the trace. It determines that `T_` is present, which corresponds in the original Esterel program to a thrown exception (`exit T`), hence the program has to terminate. This is done by first aborting `Parent`’s children with `TRANS` (in this case unnecessary, as they have all terminated already), transferring control to label `B`, and then terminating `Main`.

As the trace indicates (line 26), a total of 23 SC instructions have been executed, and solely the always-enabled `TickEnd` thread is still enabled. The trace also shows the signals emitted by the reaction. In this example, the `main` function calling the tick function not only sets the inputs (currently read in from an array), but also compares the generated output to a reference output (“`Outputs OK`”). See the original code (`grcbal3.c`, in the SC distribution) for how this is done.

One last operator in `grcbal3` not explained yet is the `JOIN` in line 29. Here the `Main` thread checks whether all of its children have terminated. If so, then `Main` also terminates, according to the semantics of SyncChart macrostates, and similarly Esterel’s concurrency operator `||`.

To summarize, `grcbal3` illustrates how thread ids and priorities can be used to schedule threads in an arbitrary fashion. In this case, we have used this to schedule threads such that signal dependencies, imposed by the Esterel/SyncCharts semantics, are adhered to. This semantics requires that within a tick all potential signal emitters run before a signal is tested. This is similar to the situation in the producer-consumer example, just that in `grcbal3` there is not just one buffer to synchronize on, but four output signals.

This example is, admittedly, fairly intricate, as it has also been designed to illustrate the scheduling challenges that Esterel poses to a compiler. For an inexperienced SC programmer it may therefore be non-obvious how to assign priorities and thread ids properly such that

signal dependency rules are adhered to; see Sec. 3.2.4 for a full discussion. There are several possible alternatives to unchecked manual priority/id assignment:

- one might relegate thread id and priority assignment to a separate analysis pass, similar to an Esterel compiler (feasible, but it would require a separate tool);
- one might use SyncCharts—or Esterel—as entry language for SC, and do the signal dependence analysis there (also possible, but this would lose the direct embedding in C); or,
- one might add run-time checks to the SC operators that ensure that no signals that have been tested already in a tick are emitted in a tick (a reasonable consistency check, easy to implement—but it does not offer a guarantee as a static analysis would do).

However, one should also note that such intricate dependencies appear to be rather rare. We can distinguish three types of programs:

Dynamically scheduled programs that require dynamic scheduling of threads, which entails run-time alterations of thread priorities (via `PRIORITY`);

Statically scheduled programs that require just static scheduling, which can be handled with thread id assignment; and

Unscheduled programs that do not impose scheduling constraints at all.

From the 10 benchmarks currently included in the SC distribution, most provided by André [3], only `grcbal3` and `Exits` belong to the first category.

Chapter 3

A Tour of SC

SC consists of a programming model, which is implemented with simulated multi-threading, a set of SC operators, and a convention on how to structure an SC program. These concepts are explained in the next sections.

3.1 The SC Programming Model

SC programs follow the *synchronous programming model* established in SyncCharts and other synchronous programming languages. This programming model is characterized by two main concepts explained in the following, the synchronous threading model and signals. The first concept is essential to SC; the second one is also provided by SC, but can be regarded as optional for the SC programmer.

3.1.1 Synchronous threading

The main concept that must be understood to program in SC is that of a *logical tick*, or *logical instant*. An SC program conceptually consists of a number of concurrent threads, whose concurrent execution is grouped (synchronized) by a progression of logical ticks. A tick boundary of a thread is usually denoted by the PAUSE operator, which thus denotes implicit synchronization points among thread; no thread can start the next tick, before all concurrent threads have completed the current tick.

In Statecharts parlance, the tick boundaries (PAUSE operators) collectively reached by the currently active threads form the *configuration* of a system. The system progresses from one stable configuration to the next. The *synchrony hypothesis* states that the computation of a reaction does not consume any time. In other words, the progression from one configuration to the next configuration is considered to not consume physical time; physical time only advances while the system rests in a stable configuration. This synchrony hypothesis is of course an abstraction from reality, where computations of course do consume time, but this abstraction allows a compositional, formally grounded semantics.

Program execution states

Figure 3.1 illustrates the execution states of the whole program, using the SyncChart formalism. Upon program start, the main thread is enabled (forked), and the program is considered

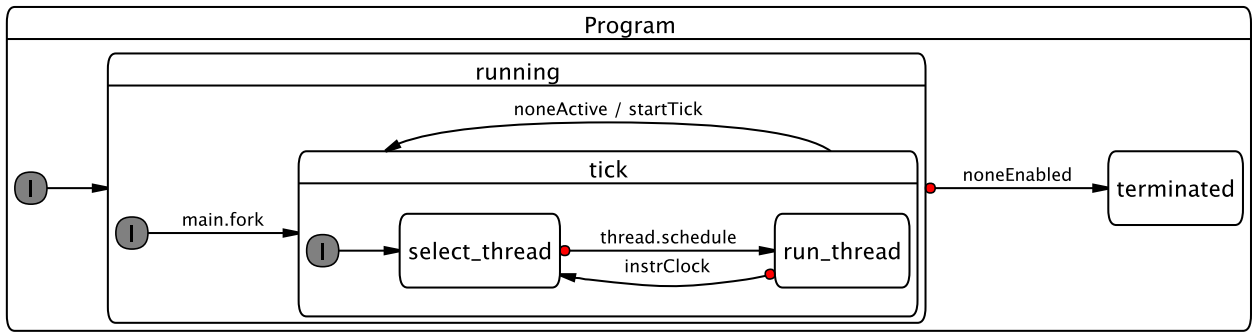


Figure 3.1: The status of the whole program (from [18]).

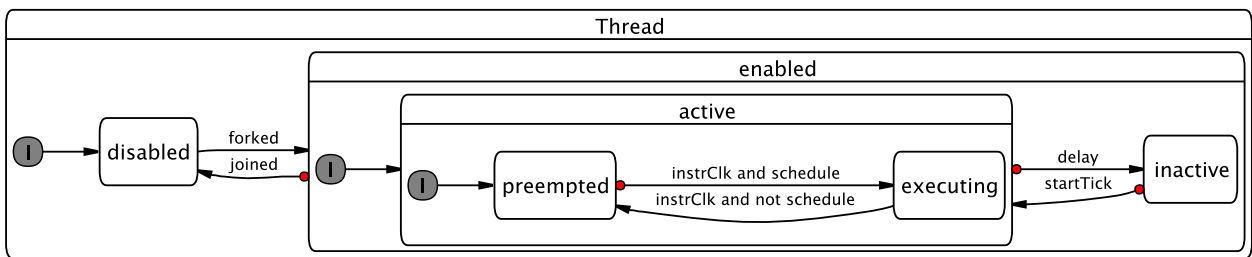


Figure 3.2: Execution status of a single thread (from [18]).

running. Entering the tick state corresponds to a call to the tick function, described in Section 3.4. The state is left (the tick function terminates) when no threads are active anymore. If all threads have become disabled (all have terminated), the whole program becomes terminated.

Note that the thread logic and the overall program logic of SC are very close to the Kiel Esterel Processor [18]. A slight difference arises within the tick state: in the KEP, each instruction cycle is started by a thread selection step, whereas in SC, threads run freely until they encounter an SC operator that implies a call to the dispatcher and a possible thread context switch. The operators in question that can thus possibly cause a thread context switch are marked in Table 2.1.

Thread execution states

The execution status of a thread is illustrated in Figure 3.2. Two flags are needed to describe the status of a thread. One flag indicates whether the thread is *disabled* or *enabled*. Initially, only the main thread is enabled. Other threads become enabled whenever they are forked (with `FORK`, see Section 3.3.1), and become disabled again when they terminate themselves (`TERM`) or get aborted by a transition that leaves the parent state (using `TRANS`). The other flag indicates whether the thread should still be scheduled within the current logical tick (the thread is *active*) or not (*inactive*).

A thread is *active* if it still has work to do in the current tick, otherwise it is *inactive*. The order of execution among active threads is statically determined by a *thread id* and a *thread*

priority. The *thread dispatcher* starts/resumes the active thread with the highest priority, this thread then becomes *executing*. Threads that are active but not executing are considered *preempted*.

Priorities can be shared among threads; if there are multiple threads with the same, highest priority, the thread with the highest id wins. To ensure that there is a unique thread to be chosen, the thread ids must not be shared among threads that can be active concurrently. SC requires each thread to be given a thread id and initial priority upon its creation (with `FORK`). SC also provides an operator to change the priority of a running thread (`PRIO`). With these mechanisms, the programmer can enforce arbitrary, deterministic thread schedules. For simplicity, these schedules are usually determined statically; however, in terms of expressiveness of the SC operators, it would also be possible to create dynamic schedules.

3.1.2 Signals

Another concept that is characteristic of synchronous languages is that of *signals*, which can be used for broadcast communication among threads. SC programmers do not have to use signals, they might achieve the same effect with the appropriate use of standard C variables. However, the explicit use of signals for thread control, for example to trigger preemptions, might help to clarify the interaction and synchronization patterns across thread boundaries.

Signals are *absent* per default. They become *present* for the current tick if a thread *emits* the signal in this tick. Any thread (including the emitting thread) can test for the presence of a signal and can change control flow accordingly, including not only conditional branches but also various forms of preemption.

SC provides a full range of signal handling operators, including local, valued, and combined signals, and tests for signal presence across tick boundaries (the `PRE` operator). SC assumes that signals become visible within the tick they are emitted, and also, unlike some other approaches [8], allows to test for signal absence in the current tick, not just in the next tick.

A word of caution: it is a common assumption of (strictly) synchronous programs that signals have a unique, well-defined presence/absence status for the duration of a tick. This effectively means that we must not test for the presence of a signal if it may still be emitted within that tick; see also Sec. 3.2.4. In other words, all writes must be performed before any reads are done. Compilers dedicated to synchronous languages perform a static *signal dependency analysis* of the program and try to compute a—usually static—schedule that orders threads (or thread segments) accordingly. A compiler rejects the program if it cannot find such a schedule. As mentioned in Sec. 2.2, one could envisage an analysis tool that performs a signal dependency analysis on an SC program and checks that the encoded schedule respects all signal dependencies. In the presence of arbitrary C control flow, this analysis would have to be conservative. If we were to use SC as intermediate language for synthesizing code from a visual SyncCharts model, it would be the responsibility of the code synthesis tool to perform a dependency analysis on the model and to schedule threads accordingly. Using plain SC, as presented here, we do not perform such a compilation or analysis; we just use a regular C compiler and C does not have a concept of signal dependencies and consistency. It is thus the responsibility of the programmer to schedule threads accordingly, using thread ids and priorities. On the other hand, this also provide the options to weaken the requirement of strict synchrony, and to program with a signal model that corresponds *e. g.* to the original Statecharts semantics. Note also that the program will in any case be deterministic, there are no race conditions that can produce different outputs for the same inputs.

3.2 Multithreading Simulation

To simulate multi-threading, we must be able to keep track of the locus of control of each thread, and we need a dispatcher that performs the context switches.

3.2.1 Coarse program counters

In a VM or hardware implementation of the SC operators, one could have direct access to a program counter that denotes the locus of control. As we are working here at the C level, we do not have that option. Instead, we annotate the C program with regular C labels, at all possible thread continuation points. In a way, these denote thread-level “basic blocks,” but unlike traditional basic blocks, they do not denote sequences of straight line code, but instead they delineate sequences of code in which no thread context switch can happen.

Using gcc’s computed goto extension, we can store these program labels in an ordinary C array. In SC, this array of *coarse program counters* is `pc[idMax]`, declared in `sc.h` (see Section 3.4.1).

Whenever a thread calls an SC operator that might result in a context switch to another thread, we must save a *continuation point* for the thread in its program counter. In our implementation, this operation is folded into the SC operators, so that it suffices to pass the continuation point of the thread along as argument to the SC operator, which then performs the book keeping.

3.2.2 The dispatcher

As explained in Section 3.1.1, the dispatcher starts/resumes the active thread with the highest priority; if there are multiple active threads with the same, highest priority, the thread with the highest id wins.

As the dispatcher may be called rather frequently—namely, whenever we perform an SC operator that can result in a context switch, see also Table 2.1—we should strive for an efficient implementation of the dispatcher. The dispatcher should be as general as necessary and as fast as possible. As the demands of SC programs on the dispatcher may vary—in particular, they may or may not use priorities—SC provides different dispatchers, which can be selected by the application, via a `#define USEPRIO C` preprocessor directive.

The dispatcher consists of two parts:

1. Computation of the current (to be dispatched) thread id, `cid`.
2. A jump to the corresponding program counter, stored in `pc[cid]`.

Listing 3.1: `selectCid()`: Computation of id of thread to be dispatched (from `sc.c`)

```
1 void selectCid () {
2   int act;
3
4   _cid = 0;
5   for (act = active; act > 1; act >>= 1)
6     _cid++;
7 }
```

The computation of `cid` is implemented in `selectCid`, see Listing 3.1. A loop iterates through thread ids, starting from the highest id (given by `idHi`, see Section 3.4.4) downwards to id

0. In the loop body, we check whether the currently examined thread id is active, and if so, whether it has a higher priority than the highest priority encountered so far. The run time of this implementation is linear in the number of thread ids in use.

This function only considers whether a thread is active or not. In the current SC implementation, this information is stored in a bit vector `active`. Hence it suffices to set `cid` to the position of the highest set bit in `active`. The implementation of `selectCidNoprio` uses the obvious algorithm, with a run time linear in the position of the highest bit. Note that there are also alternatives that run logarithmic to bit vector size¹. Which algorithm is actually faster depends on the application.

Listing 3.2: `dispatch()`: Variable definitions for the dispatcher (from `sc.h`)

```

1  #if ((defined __i386__ || defined __amd64__ || defined __x86_64__) && defined __XGNUC__)
2  // Version 1: x86 + gcc available .
3  // Use fast Bit Scan Reverse assembler instruction .
4  #define dispatch() \
5  __asm volatile (" bsr %1,%0\n" \
6  : "=r" (_cid) \
7  : "c" (active) \
8  ); \
9  goto *_pc[_cid]
10
11 #else
12 // Version 2: x86 + gcc not available .
13 // Call function, defined in sc.c.
14 #define dispatch() \
15 selectCid (); \
16 goto *_pc[_cid]
17 #endif

```

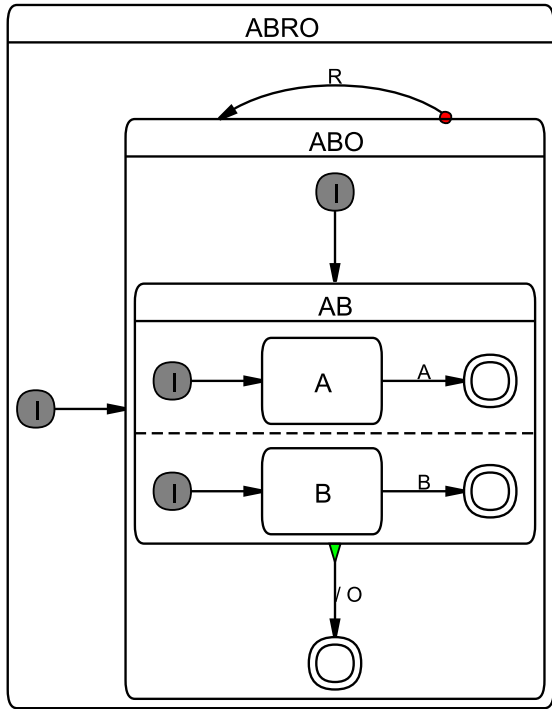
Fortunately, many processor instruction sets provide an assembler instruction that does exactly this, to detect the index of the highest set bit. The x86 does this with the Bit Scan Reverse (BSR) instruction. Using the `gcc` assembler escape, we can embed this instruction into C and thus obtain an even faster dispatcher. This consists of just a couple of instructions and has constant run time. This dispatcher variant is not implemented as a separate function, but instead as a macro, to be expanded/inlined by the preprocessor. This does not unduly increase the code size, and saves the function call overhead at run time. It also alleviates the need to link against `sc.c` that defines the alternative dispatcher functions (see Section 3.4.1). Listing 3.2 shows how the dispatcher is defined.

Note that the current implementation of SC, based on bit vectors implemented as simple integers, assumes that the number of concurrent threads does not exceed the word size. The same limitation applies to signals, whose presence/absence status is also implemented as integer-based bit vectors. Neither limitation has posed any problems in the applications considered so far. However, it should be rather straightforward to lift either limitation, and to use bit vectors of arbitrary size or some other unrestricted data structure.

3.2.3 Thread and label structuring

To illustrate how the thread structure is derived from a SyncChart, consider the ABRO example in Fig. 3.3 [3, Fig. 5-12]. ABRO is arguably the “hello-world” program of synchronous programming and has the following behavior: Two concurrent threads wait for signals A and

¹See eg <http://graphics.stanford.edu/~seander/bithacks.html#IntegerLog>



(a) SyncChart

<pre> 1 TICKSTART(islnit, 4); 2 3 do { 4 FORK(AB, 1); 5 FORKE(ABOMain); 6 7 AB: 8 FORK(WaitA, 2); 9 FORK(WaitB, 3); 10 FORKE(ABMain); 11 12 WaitA: 13 AWAIT(A); 14 TERM; 15 16 WaitB: 17 AWAIT(B); 18 TERM; 19 20 ABMain: 21 JOIN; 22 EMIT(O); 23 TERM; 24 25 ABOMain: 26 AWAIT(R); 27 ABORT; 28 } while (1); 29 30 TICKEND;</pre>	<pre> 1 TICKSTART(islnit, 4); 2 3 ABO: 4 FORK(AB, 1); 5 FORKE(ABOMain); 6 7 AB: 8 FORK(WaitA, 2); 9 FORK(WaitB, 3); 10 FORKE(ABMain); 11 12 WaitA: 13 AWAIT(A); 14 TERM; 15 16 WaitB: 17 AWAIT(B); 18 TERM; 19 20 ABMain: 21 JOIN; 22 EMIT(O); 23 TERM; 24 25 ABOMain: 26 AWAIT(R); 27 ABORT; 28 TRANS(ABO); 29 30 TICKEND;</pre>
--	---

(b) SC tick function

(c) SC (only) tick function

Figure 3.3: The ABRO example.

B; once these have occurred, in any order, output O is emitted. If R is present, the behavior is reset. As the transition triggered by R is a *strong abort*, the transition takes priority over the internal behavior of ABO : if R is present in a tick, ABO does not get to execute in that tick.

A note on the SC implementation: as there is normal termination leaving ABO , there is no need for a `JOIN` on thread AB .

Thread structure

Each macrostate of degree of concurrency n has n *embedded* threads. In $ABRO$:

- ABO has one embedded thread (AB), and
- AB has two embedded threads ($WaitA$ and $WaitB$).

Naming threads and their initial label

- `TickEnd` is the thread returning from the tick function, at macro `TICKEND`. It must have id 0, as it has to run after all other threads.
- `Main` is the main thread, activated in the initial tick upon entering the tick function.

- Other threads are named after their first state.
- The initial label of a thread is named after the thread.

Note that the last two rules are in most cases redundant. However, it can be the case that a thread does not commence directly at its first state, in particular if the initial transition has to perform some action; see for example the initial transition of thread **S1** in **PrimeFactor** (Fig. 4.8). In such cases, the recommended convention is:

- The thread and its entry point should still be named after its first state S , according to the SyncChart diagram.
- However, the entry point of the state should be renamed to S_{surf} .

The latter part is derived from the surface/depth distinction, elaborated on in the following.

Surface vs. depth

In SyncCharts, as well as in Esterel, one distinguishes

Immediate transitions which can potentially be taken in the same tick as their source state is entered, and

Delayed transitions which will only become enabled from the next tick onwards.

Transitions are by default delayed; immediate transition triggers are indicated by a $\#$ -mark (see also Fig. 2.2a).

One also distinguishes the

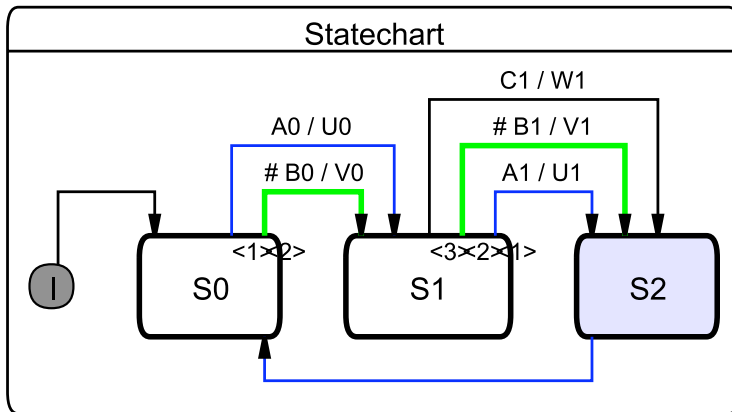
Surface of a statement, which is what is executed in the initial tick and which includes only the immediate transitions, and the

Depth of a statement, which is where execution commences in subsequent ticks and which includes immediate as well as delayed transitions.

Consider for example in **PrimeFactor** (Fig. 4.8b) state **S1**, which has an immediate transition triggered by **B**, and a delayed transition triggered by **A**. The former is tested in the surface of **S1**, at label **S1surf**, as well as later at the depth (which is commenced at label **S1depth**), whereas the latter transition is only tested at the depth.

It appears that in most cases code can be structured such that there is no need for code duplication between surface and depth. However, this cannot always be avoided. Consider the **SurfDepth** example in Fig. 3.4. The transitions from **S0** to **S1** can be ordered such that the transition-number priority can be honored (test **A0** before **B0**) as well as the immediate (**B0**)/non-immediate (**A0**) distinction. However, this is not possible with the transitions from **S1** to **S2**. The only immediate transition, triggered by **B1**, has a transition number between the two other, delayed transitions. Hence we must duplicate the test for **B1**, once at the surface label **S1surf**, once in the depth code starting at **S1depth**.

Note that this duplication concerns not only the test of the transition trigger, but also possible transition actions. In **SurfDepth** this applies to the emission of **V1**. We here chose



(a) SyncChart, during simulation—also illustrating KIEL’s horizontal/vertical layout.

```

1  TICKSTART(islnit, 1);
2
3  while (1) {
4
5  // State S0 – surface
6  while (1) {
7  if (PRESENT(B0)) {
8  EMIT(V0);
9  break; }
10 PAUSE;
11
12 // State S0 – depth
13 if (PRESENT(A0)) {
14 EMIT(U0);
15 break; }
16 }
17
18 // State S1 – surface
19 if (PRESENT(B1)) {
20 EMIT(V1);
21 } else {
22 while (1) {
23 PAUSE;
24
25 // State S1 – depth
26 if (PRESENT(A1)) {
27 EMIT(U1);
28 break; }
29 if (PRESENT(B1)) {
30 EMIT(V1);
31 break; }
32 if (PRESENT(C1)) {
33 EMIT(W1);
34 break; }
35 }
36 }
37
38 // State S2
39 PAUSE;
40 }
41
42 TICKEND;
  
```

(b) SC tick function

Figure 3.4: The SurfDepth example.

to minimize code size and to follow the write-things-once principle as much as possible, by sharing one `EMIT(V1)`₂₀ statement between the transition tests at lines 14 and 19, with the `GOTO(L2)`₁₅ statement. An alternative would be to have another `EMIT(V1)` directly after the test at line 14, followed by a `GOTO(S2)`. This increases the program size by one statement (the `EMIT(V1)`), but saves one statement at run time (the `GOTO(S2)`).

Label naming

- The (surface) entry point of a macrostate S gets label S , or exceptionally (see above) label S surf.

Example: label `ABO`.

Example: in `PrimeFactor` (Fig. 4.8), label `S1surf`.

- The depth entry point of a state S gets label S depth.

Example: in `PrimeFactor` (Fig. 4.8), label `S1depth`.

- For macrostate S , the entry point of the code that checks transitions attached to S gets label S main.

Examples: `ABmain`, `ABOmain`.

3.2.4 Thread scheduling

As SC is embedded into plain C, SC programs are deterministic, and there is no need for classical synchronization among threads using semaphores or similar concepts. However, as mentioned in Sec. 3.1.2, certain *scheduling rules* must be followed when encoding a specific SyncChart in SC, to adhere to the original, synchronous semantics.

This is related to the scheduling problem that a compiler for synchronous languages faces, and one might use similar concepts to address this. For example, one might transform a SyncChart into something like a CKAG (Concurrent KEP Assembler Graph [18]) that expresses scheduling constraints, and then transcribe this into an SC program. As we here—so far—assume a human programmer that writes an SC program, we do not describe the scheduling task in algorithmic terms, but instead give precedence constraints that must be fulfilled. As stated in Sec. 2.2, this is a non-trivial problem in the general case; however, it appears that most SC programs exhibit relatively few scheduling constraints.

Precedence of operations

Let Op_1 and Op_2 be two operations that must be performed in an SC program. For example, in `ABRO`, let Op_1 be the test for the presence of signal `R` of the `Main` thread in line 24, abbreviated as $Op_1 = \text{Main}_{24, \text{PRESENT}(\text{R})}$, and let $Op_2 = \text{WaitA}_{12, \text{PRESENT}(\text{A})}$. In this case, Op_1 , which corresponds to the strong abort transition on `ABO`, must be executed before Op_2 , which is a transition nested within `ABO`. We say that Op_1 *has precedence* over Op_2 , and write this as $Op_1 \succ Op_2$, in this example $\text{Main}_{24, \text{PRESENT}(\text{R})} \succ \text{WaitA}_{12, \text{PRESENT}(\text{A})}$.

In the following, we will use the terms statements and operations interchangeably. A note on notation: we may use the line-number-in-subscript notation throughout the report to refer to specific statements in a program. We may also abbreviate statements, for example by omitting label arguments.

We are now ready to define the precedence constraints imposed by a SyncChart. It is $Op_1 \succ Op_2$ if

1. Op_1 and Op_2 can be executed in the same tick, and
2. one of the following conditions holds:

Outer-inner precedence Op_1 tests the trigger of a strong abort or suspension associated with a state S , and Op_2 belongs to a descendant (inner state) of S .

Example: in ABRO, $\text{Main}_{24,\text{PRESENT}(\text{R})} \succ \text{WaitA}_{12,\text{PRESENT}(\text{A})}$.

Inner-outer precedence Op_1 belongs to a descendant (inner state) of S , and Op_2 tests for normal termination or tests the trigger of a weak abort associated with a state S

Example: in ABRO, $\text{WaitA}_{12,\text{PRESENT}(\text{A})} \succ \text{AB}_{19,\text{JOIN}}$.

Transition-number precedence Op_1 and Op_2 are associated with transitions that are associated with the same state, and the transition Op_1 is associated with has a higher priority than Op_2 . Here, we refer to transition priorities indicated in the SyncChart with numbers (*increasing* number for *decreasing* priority).

Note that SyncCharts already impose some ordering on the transition numbers within the transitions associated with the same state. Highest priority (lowest number) have strong aborts, followed by suspension, followed by weak aborts, followed by normal termination.

Example: in Exits (Fig. 4.2), considering the transitions associated with state M10, the strong abort has precedence over normal termination, *i. e.*, $\text{M10}_{27,\text{PRESENT}(\text{A})} \succ \text{M10}_{20,\text{JPPAUSE}}$; note that the JOIN of the normal termination is folded in with PRIO and PAUSE.

Write-read precedence Op_1 writes (emits) a signal, which is read (tested for presence) by Op_2 .

Here, with “signal” we refer to general shared variables for which writer-reader precedence should be respected within a tick. This applies to signals in the sense of SyncCharts or Esterel, operated on via the SC signal operators (Table 2.2), but also to shared C variables, such as the buffer BUF in the PCO example (Fig. 2.1).

Example: in grcbal3 (Fig. 2.2c), $\text{A1}_{8,\text{EMIT}(\text{B})} \succ \text{A2}_{19,\text{PRESENT}(\text{B})}$.

Example: in PCO (Fig. 2.1c), $\text{Prod}_{33,\text{BUF}=1} \succ \text{Cons}_{39,\text{tmp}=\text{BUF}}$.

We summarily refer to the first three types of precedence constraints as *structural constraints*, whereas the last one is a *signal constraint*.

Fulfillment of precedence constraints

We also classify a precedence constraint $Op_1 \succ Op_2$ as follows:

Intra-thread precedence Op_1 and Op_2 belong to the same thread.

In this case, the constraint must be fulfilled via the sequential ordering of the operations within a thread.

Inter-thread precedence Op_1 and Op_2 belong to concurrent threads.

In this case, the constraint must be fulfilled via an appropriate assignment of static thread ids and, if necessary, dynamic priorities.

Thread precedence

We can lift the notion of precedence from individual operations to the threads that they belong to. For an operation Op , let $thrd(Op)$ be the thread associated with Op . For example, in ABRO, it is $thrd(\text{WaitA}_{12, \text{PRESENT}(A)}) = \text{WaitA}$. Then, for Op_1, Op_2 with $Op_1 \succ Op_2$ and $t_1 = thrd(Op_1), t_2 = thrd(Op_2)$, this implies $t_1 \succ t_2$. In ABRO, $\text{Main}_{24, \text{PRESENT}(R)} \succ \text{WaitA}_{12, \text{PRESENT}(A)}$ implies $\text{Main} \succ \text{WaitA}$. In other words, **Main** should be scheduled before **WaitA**.

In some cases it is convenient to use a mixed notation that orders an individual operation with another thread. For example, $\text{Main}_{24, \text{PRESENT}(R)} \succ \text{WaitA}$ expresses that the presence test on R must run before thread **WaitA**.

Static vs. dynamic scheduling

For an SC program P derived from a SyncChart and a pair of operations in P , the SyncChart either specifies a fixed order in which the operations must be performed, or it does not specify an order at all. All precedence constraints on individual operations are static. In other words, \succ is a partial order with respect to individual operations in P .

However, at the thread level, it may be the case that for a pair of threads T_1 and T_2 in P there exist operations in T_1 and T_2 that induce $T_1 \succ T_2$, and simultaneously other operations that induce $T_2 \succ T_1$. In other words, \succ is not necessarily a partial order with respect to threads in P .

If \succ is a partial order in P at the thread level, then it is possible to schedule all threads statically by just assigning them thread ids that respect \succ . There is no need for dynamic priorities, all thread priorities can remain at 0. As noted in Sec. 2.2, it appears that most programs belong to this category of statically schedulable programs.

If \succ is not a partial order in P at the thread level, then one should still assign thread ids in a way that static precedences between threads are met; however, one must use positive thread priorities as well to resolve the remaining dynamic precedences.

Furthermore, immediate transitions must be properly distinguished from delayed transitions—see also Sec. 3.2.3.

In the following, we will illustrate how precedence constraints are met in SC programs with the examples introduced so far, ABRO, PCO, and grcbal3. Chapter 4 provides further examples.

Precedence constraints in ABRO

In ABRO, there are the following structural inter-thread constraints at the thread level.

1. Strong abortion on ABO (*outer-inner*):

$\text{Main} \succ \text{AB}, \text{Main} \succ \text{WaitA}, \text{Main} \succ \text{WaitB}$

2. Normal termination on **AB** (*inner-outer*):

$\text{WaitA} \succ \text{AB}, \text{WaitB} \succ \text{AB}$

This thread precedence relation induces a partial order, which can be fulfilled with the following thread id assignment: $\text{AB} = 1, \text{WaitB} = 2, \text{WaitA} = 3, \text{Main} = 4$. We generally omit stating explicitly the id of **TickEnd**, since it always must be 0; however, the program must still declare the **TickEnd** thread and assign id 0 to it.

Note that the order between **WaitB** and **WaitA** could as well have been reversed. Here, **WaitA** has been given the higher priority such that the order of execution between **WaitA** and **WaitB** is consistent with the order in which they appear in the program, as an aid in helping to understand the execution trace. Apart from this small consideration for the human observer, it makes no difference in which order threads are executed that do not have a precedence constraint between them. The order in which the threads appear in the program has no semantic relevance.

Precedence constraints in PCO

In PCO (Fig. 2.1c, p. 4), there are the following inter-thread constraints at the thread level.

1. Weak abortions on **Parent** (*inner-outer*):

$\text{Prod} \succ \text{Main}, \text{Cons} \succ \text{Main}, \text{Obs} \succ \text{Main}$

2. Writer on **BUF** before reader on **BUF** (*write-read*):

$\text{Prod} \succ \text{Cons}, \text{Prod} \succ \text{Obs}$

Note that the first constraint is again a structural constraint, but the second is a signal constraint. Again, this precedence relation induces a partial order at the thread level, which is observed by the following thread id assignment: $\text{Main} = 1, \text{Cons} = 2, \text{Obs} = 3, \text{Prod} = 4$.

Precedence constraints in grcbal3

As pointed out in Sec. 2.2, **grcbal3** (Fig. 2.2) is a relatively complex example that requires dynamic scheduling. In other words, \succ is not a partial order at the thread level. We will therefore use the mixed operation/thread notation to capture the constraints as concisely as possible while still permitting an ordering, without contradictions.

1. Weak abortion and normal termination on macrostate (*inner-outer*):

$\text{A1} \succ \text{Main}, \text{A2} \succ \text{Main}, \text{A3} \succ \text{Main}$

This inter-thread structural constraint is met by assigning **Main** the thread id 1 (the lowest possible, apart from the **TickEnd** thread) and priority 0.

2. Precedence of weak abortion over normal termination (*transition-number*):

$\text{Main}_{27, \text{PRESENT}(\top)} \succ \text{Main}_{29, \text{JOIN}}$

This intra-thread structural constraint is met by ordering the operations in the program accordingly.

3. Communication via signal **B** (*write-read*):

$$A1_{8,EMIT(B)} \succ A2_{19,PRESENT(B)}$$

This constraint is met by executing the first operation at priority 3, as induced by the $Main_{2,FORK(3,A1,ids[A1])}$ statement, and the second at priority 2.

4. Communication via signal **C** (*write-read*):

$$A2_{20,EMIT(C)} \succ A1_{10,PRESENT(C)}$$

This constraint is met by executing both operations at priority 2, as induced by $A1_{9,PRIO(2)}$ and $Main_{3,FORK(2,A2,ids[A2])}$, and assigning **A2** a higher thread id than **A1**.

5. Communication via signal **D** (*write-read*):

$$A1_{11,EMIT(D)} \succ A3_{23,PRESENT(D)}$$

This constraint is met by executing the first operation at priority 2 and the second at priority 1.

6. Communication via signal **E** (*write-read*):

$$A3_{24,EMIT(E)} \succ A1_{13,PRESENT(E)}$$

This constraint is met by executing both operations at priority 1 and assigning **A3** a higher thread id than **A1**.

3.3 SC Operators

There are three classes of SC operators: *SC Thread Handling Operators*, *SC Signal Operators*, and *SC Sequential Control Operators*.

3.3.1 SC thread handling operators

An overview of the SC Thread Handling Operators, which perform the multi-threading simulation and form the core of SC, is given in Table 2.1, p. 6.

Tick start and end

TICKSTART and TICKEND do some book keeping. For example, in the initial tick, TICKSTART initializes the TickEnd thread (see Section 3.4.3) and activates the **Main** thread; in subsequent ticks, TICKSTART activates the enabled threads.

TICKEND determines whether there are still any enabled threads, apart from the never disabled TickEnd thread.

Pausing, suspending, aborting and terminating a thread

PAUSE pauses the currently active thread. This entails setting the program counter of the current thread to the label provided as argument, to deactivate the current thread, and to call the dispatcher.

SUSPEND suspends (“freezes”, “steals the clock from”) the current thread and its descendants for the current tick and calls the dispatcher. For an example, see `Count2Suspend`, Fig. 4.1, p. 37.

Differences between PAUSE and SUSPEND:

- PAUSE deactivates just the current thread, whereas SUSPEND also deactivates its descendants. The latter exploits that the PCs of the descendants must reside at tick boundaries, *i. e.*, there is nothing more to do for the descendants in the current tick.
- Unlike PAUSE, SUSPEND must do some signal handling in case local signals and `pre` are used, as explained in Section 3.3.2.

TERM terminates the current thread by disabling it.

TRANS performs an abortion of the current thread and its descendants, by simply disabling them, and transfers control to the specified label l . In SyncCharts, this corresponds to a (weak or strong abort) transition from the current state to some other state.

Whether TRANS corresponds to a weak or strong abort is merely a question of whether TRANS is executed *before* the descendant threads have computed the tick (*strong abort*) or *after* the descendants have run (*weak abort*). See also the outer-inner vs. inner-outer precedences discussed in Sec. 3.2.4. Again, this is an implication of the SyncChart semantics and can be viewed as a (reasonable) convention. Nothing would prevent a programmer to break with this convention and schedule a TRANS arbitrarily, in the middle of the execution of the descendant. This would not break determinism (we still have a sequential C program), but it would probably make the flow of the program more difficult to comprehend.

To summarize, a thread voluntarily relinquishes control for the remainder of the tick via PAUSE, SUSPEND, or TERM. A thread may also be aborted when a (transitive) parent performs a TRANS.

Fork and join

A sequence of FORK statements, followed by a FORKE statement, together form a *fork*. Each FORK creates a child thread by initializing its program counter and its priority and enabling it. FORKE then registers the descendant threads with the current (parent) thread and calls the dispatcher. The parent thread must know about its descendant threads to detect their termination, and also possibly to terminate them in case the parent is aborted. The set of descendants includes the newly created child threads, and, in case these will possibly fork threads as well, their descendants (transitively) as well. Note that the latter is necessary for abortions, but not for normal termination, as normal termination should respect the hierarchical ordering (grandchildren should terminate before children terminate normally).

JOIN performs the corresponding *join* operation, which checks whether all descendant threads have terminated normally. If they have terminated, control transfers immediately to the l_{then} label. This corresponds to a *normal termination transition* in SyncCharts. By (reasonable) convention, normal termination transitions have the lowest priority, and there can be only one such normal termination transition. See also the notes on transition-number precedence, Sec. 3.2.4. This means that after performing an unsuccessful JOIN, there is nothing else the current thread has to do for the current tick, and it pauses. We exploit this by folding the PAUSE into the else-branch of the JOIN. That is, if the descendant threads have not terminated, we execute a `PAUSE(l_{else})`.

Notes:

- One might also decide to break with the SyncChart convention of pausing after an unsuccessful join, and to supply an SC JOIN variant that does not automatically pause. This would be trivial to implement, but so far there has no need arisen to do so.
- A JOIN is only required if the corresponding SyncChart does have a normal termination transition. If the parent thread never terminates, or if it is only terminated through abortion (via TRANS), no JOIN is required.
- Normally, a fork spawns off *child threads* of the current thread, and the current thread keeps executing, at the label specified by FORKE. However, we may also construct a fork without FORKE, which just consists of a sequence of FORK statements that effectively create *sibling threads* of the current thread. Here the current thread simply keeps executing after the FORK statements. Since there is no FORKE, the dispatcher will not be called, so the current thread should be the one with the highest priority/thread id of the sibling threads. See Shifter3 (Fig. 4.5) for an example.

Thread priority handling

The initial priority of a thread is assigned upon creation of the thread, as argument to FORK. It may be necessary to change the priority of a thread later at run time. This is done with the PRIO operator. Note that within a tick, it is only meaningful to lower the priority of a thread, not to raise it, since if a thread is already executing, there is no effect when raising its priority [18]. We can use priority lowering to yield to other threads. However, as thread priorities are preserved across tick boundaries, we may want to raise a priority at the end of a tick, to start the next tick with a higher priority.

The PRIO operator entails a call to the dispatcher, as now another active thread might be the one with the highest priority, or at the same priority but with a higher thread id. However, there are common situations where the next operator to be executed by the thread that has just called PRIO is another operator that necessitates the dispatcher. For example, in the aforementioned scenario where we raise a priority to start the next tick with a higher priority, PRIO is followed immediately by PAUSE. In this case, the first call to the dispatcher is superfluous. Therefore, there are two *combined operators*, PPAUSE and JPPAUSE, that combine PRIO with other operators. These are not just syntactic sugar, but optimize performance, and code size. For example, JPPAUSE combines PRIO with a JOIN and a PAUSE, thus reducing three potential calls to the dispatcher to just one call.

3.3.2 SC signal operators

If an SC program wants to use signals (see Section 3.1.2), it can use the operators shown in Table 2.2. Signals must be declared in the `signaltype` (see Section 3.4.3).

Global vs. local signals, reincarnation

We can classify signals as follows:

Global signals Signals get initialized once at the beginning of each tick. This is the default in SC.

Local signals Signals can be declared for the scope of a SyncChart macrostate. This implies that signals are initialized whenever the macrostate is entered. This is achieved with the **SIGNAL** operator, see below.

An interesting aspect of local signals is the possibility of *reincarnation*, or *schizophrenia*. A loop around the macrostate declaring a local signal may provoke the simultaneous existence of two different “incarnations” of the local signal [3]. This is illustrated in the **Reincarnation** example (see Figure 4.7).

Pure signals

As explained in Section 3.1.2, signals can be *present* or *absent*. *Pure signals* just have this presence status, unlike valued signals, which also carry a value (see next section).

The **SIGNAL** operator initializes a local signal, as explained above, by setting its status to absent.

The **EMIT** operator sets a signal present.

The **PRESENT** operator checks for the presence of a signal. If it is, control proceeds normally to the next statement (*then branch*), otherwise it jumps to the specified label l_{else} (*else branch*).

Valued signals

Valued signals carry a value of a certain type. So far, SC implements just integer valued signals, extensions to other types (or a more generic typing mechanism) would be straightforward. The **EMITINT** operator emits a signal S (makes it present) and assigns it a value val .

If an application uses valued integer signals, the signal declaration in **signaltype** (see Section 3.4.3) has to order the valued signals before the pure signals. The number of valued signals, say n , must be declared with a “**#define valSigIntCnt n** ” directive.

The **VAL** operator retrieves the value of S and stores it in a register (an ordinary C variable). In SyncCharts/Esterel, this is done with the $?S$ notation. It would also have been straightforward to implement **VAL** as a function that returns the value directly, which might seem a bit more natural from the C perspective. However, to stay in the spirit of operators that could also be used for an ISA, **VAL** requests an explicit “destination register.”

Combine functions Adhering to the synchronous, deterministic SyncCharts semantics, signals have a unique presence/absence status throughout a tick. This is no problem for pure signals, in so far as the execution of multiple **EMIT** statements within one tick has no further effect, we just set an already present statement to present again. For valued signals, the situation is slightly more complicated, as valued signals are considered to carry a unique value throughout a tick as well. This at first sight conflicts with the possibility of executing multiple valued emissions within one tick, as these valued emissions might occur with different values. But SyncCharts (as Esterel) offers an elegant way out of this dilemma, by way of *combine functions*. These functions must be binary, commutative, associative functions that can be used to combine multiple values into one uniquely determined value. For example, we may use addition or *max* as combine functions. Subtraction would not be allowed, as it is not associative, and we cannot, in general, make any assumptions on the order in which values are supplied to the combine function.

So far, SC implements multiplication as combine function. `EMITINTMUL` emits an integer signal, combined with multiplication. Again, it would be straightforward to extend this to other combine functions, or to implement a generic mechanism.

As mentioned above in the context of signal reincarnation, it is possible that statements are be executed multiple times within a macro tick. This can lead to interesting—but still explainable and deterministic—behavior when using combined valued signals, as illustrated in the `PrimeFactor` example (see Figure 4.8).

Crossing tick boundaries (PRE)

In general, we are interested in the presence status (and perhaps value) of a signal for the current tick. However, to implement delays, or sometimes to break “dependency cycles,” we may want to access the status/value of a signal in the previous tick. This functionality is provided in `SyncCharts/Esterel` with the `pre` operator, and SC provides this functionality as well.

If this functionality is used, SC has to do some further book keeping, and this has to be indicated in the application with a `#define usePRE` directive.

`PRESENTPRE` is like `present`, but refers to the presence status of S not in the current tick, but in the previous tick.

`VALPRE` is like `VAL`, but again refers to the previous tick.

Pre, suspend, and local signals There is an interesting interaction between `pre`, suspension, and local signal declaration. Recall that suspension “steals the clock” from a thread (Section 3.3.1). If a thread has declared local signals and wants to access their status in the previous tick (via `PRESENTPRE` or `VALPRE`), “stealing the clock” from a the thread means that in the next tick when the thread is not suspended any more, the “previous tick” refers to the previous tick in which the thread was not suspended yet. See the `PreAndSuspend` example (Figure 4.6) for illustration.

To handle this case properly, the SC program has to do some bookkeeping. Specifically, it must keep track of local signals of states that might be suspended. To let SC do this, the application must provide a mapping from thread ids to lists of signals that are declared local to the thread, or its descendants. This mapping must be given by the `sigsDescs[]` array, see the original code (`PreAndSuspend.c`, in the SC distribution) for how this is done.

Whenever a thread i is suspended, the signals given in `sigsDescs[i]` are added to a list of signals (`sigsFreeze`) whose status is preserved into the next tick.

3.3.3 SC sequential control operators

The lower part of Table 2.2 lists further SC operators dedicated to sequential control. The `GOTO` is just what it says, implemented directly as a C `goto`. It is listed as an SC operator merely for completeness.

`SyncCharts` allow entry and exit actions to be associated with a state, these can also be used in SC, as explained in the following.

Entry actions

An *entry action* associated with a state S is performed whenever S is entered. This can be implemented in SC basically as a code sequence that immediately precedes the entry point of S , and redirecting transitions to S to the beginning of the entry action. Hence, no special SC operators are needed for entry actions.

Exit actions

An *exit action* is performed whenever the state is left. This also includes abortions, of the state itself or one of its (transitive) parents. This makes exit actions more powerful than entry actions, and their implementation does require specific SC operators.

An aborted thread does not regain control, so the aborting thread must ensure that any exit actions associated with an aborted thread are still performed. Again, there is a clear rule on what should happen when multiple exit actions might be performed: when macrostates with exit actions are nested, the exit actions are executed in the innermost to outermost order.

SC provides two operators for writing exit actions. $\text{CALL}(l, l_{ret})$ is an unconditional function call to label l . As we do not have direct access to a program counter, we must also explicitly specify the return address l_{ret} . CALL can be used whenever it is clear that the exit action must be called. For an exit action associated with state S , this could be for example at a regular exit point of S , or before S aborts and transfers to another state via TRANS . It could also be at an abortion of a parent state T , if S must be active whenever T is, *i. e.*, there are no sibling states of S .

The RET instruction returns from a function call, by transferring control to the l_{ret} label supplied to the last call instruction. Note that since exit actions are not nested, there is no need for a return address stack, it suffices to just remember one return address (implemented as global variable `returnAddress`). However, should one want to use the SC call mechanism also for nested calls, it would be straightforward to implement a stack instead of a simple return address variable..

The interesting case, as already mentioned, are abortions. Consider the situation where S has some sibling states, and the parent T gets aborted. When T gets aborted the exit action of S must be performed if S is active; otherwise, when a sibling of S is active, the exit action of S must not be performed. To implement this behavior, the $\text{ISAT}(id, l_{state}, l)$ operator can be used. It checks whether thread id is at state l_{state} ; if this is the case, control proceeds to the next instruction, which then commences the `on exit` function of associated with id at state l_{state} . Else, control proceeds to label l .

SC also provides $\text{ISATCALL}(id, l_{state}, l_{action}, l)$ as a shorthand for $\text{ISAT}(id, l_{state}, l)$; $\text{CALL}(l_{action}, l)$. For example, in the `Exits` code (Fig. 4.2b), the ISATCALL at label `L3`, which is reached upon normal termination of state `M10`, conditionally calls the exit action of `M2`. An equivalent SC program that does not make use of this shorthand is shown in Fig. 4.3a.

```

1 | WaitA:
2 |   do { goto _LL66; _L66: if (!(signals & (1 << A))) { _LL66: _pc[_cid] = &&_L66; goto _L.PAUSEG; } } while (0);
3 |   do { goto _L.TERM; } while (0);
4 |
5 | WaitB:
6 |   do { goto _LL70; _L70: if (!(signals & (1 << B))) { _LL70: _pc[_cid] = &&_L70; goto _L.PAUSEG; } } while (0);
7 |   do { goto _L.TERM; } while (0);
8 |
9 | ABMain:
10 |  do { _L74: if (((enabled & _descs[_cid]) == 0)) { goto _LL74; } _pc[_cid] = &&_L74; goto _L.PAUSEG; _LL74: (void)
11 |    0; } while (0);
12 |  do { signals |= (1 << O); } while (0);
13 |  do { goto _L.TERM; } while (0);
14 |
15 | ABOMain:
16 |  do { goto _LL79; _L79: if (!(signals & (1 << R))) { _LL79: _pc[_cid] = &&_L79; goto _L.PAUSEG; } } while (0);
17 |  do { enabled &= ~_descs[_cid]; active &= ~_descs[_cid]; } while (0);
18 | } while (1);
19 |
20 | _L.TICKEND: return (enabled != (1 << 0)); _L.TERM: enabled &= ~(1 << _cid); _L.PAUSEG: active &= ~(1 << _cid);
    _L.dispatch: selectCid(); goto *_pc[_cid];
    }

```

Figure 3.5: ABRO tick function after macro expansion (produced by gcc -E).

3.3.4 An example of expanded macros—ABRO

Fig. 3.3.4 shows the ABRO tick function from Fig. 3.3.4 after macro expansion. For better readability, Fig. 3.3.4 shows the same function with added comments and reformatting, and Fig. 3.3.4 shows a version with stripped empty statements (`do ... while (0)`, `(void) 0`). All versions produce identical code (gcc does not require an optimization level to strip these empty statements).

3.4 SC Structure

3.4.1 Program files

There are two variants possible, the *minimal variant* that does not link in `sc.c`, and the (*extended variant*) that does link it in.

The minimal files variant

An SC program consists of at least the following files:

sc.h A header file that defines a number of types, global variables and the SC macros.

APP.c A C file that defines an application *APP* (for example, `ABRO.c`). This must include `sc.h`.

The above is sufficient, if no separate, alternative dispatcher routine is required, which in turn requires that

1. the application does not depend on thread priorities, and

```

1  int tick(int islnit )
2  {
3      // Thread ids: AB=1, WaitB=2, WaitA=3, Main=4
4
5      // TICKSTART(islnit);
6      if ( islnit ) {
7          tickCnt = 0;
8          pc[TickEnd] = &&.L.TICKEND;
9          pr[TickEnd] = 0;
10         enabled = (1 << ids[TickEnd]);
11         active = enabled;
12         cid = ids[Main]; ;
13         enabled |= (1 << cid);
14         active |= (1 << cid);
15     } else {
16         active = enabled;
17         goto _L.dispatch ;
18     };
19
20     ABO: // PAR(0, AB, ids[AB]);
21     do {
22         pc[ids[AB]] = &&AB;
23         pr[ids[AB]] = 0;
24         enabled |= (1 << ids[AB]);
25         active |= (1 << ids[AB]);
26     } while (0);
27
28     // PARE(0, ABomain, id2b(AB) | id2b(WaitA) | id2b(WaitB));
29     do {
30         pc[cid] = &&ABomain;
31         pr[cid] = 0;
32         descsc[cid] = (1 << ids[AB]) | (1 << ids[WaitA]) | (1 << ids[
33             WaitB]);
34         goto _L.dispatch ;
35     } while (0);
36
37     AB: // PAR(0, WaitA, ids[WaitA]);
38     do {
39         pc[ids[WaitA]] = &&WaitA;
40         pr[ids[WaitA]] = 0;
41         enabled |= (1 << ids[WaitA]);
42         active |= (1 << ids[WaitA]);
43     } while (0);
44
45     // PAR(0, WaitB, ids[WaitB]);
46     do {
47         pc[ids[WaitB]] = &&WaitB;
48         pr[ids[WaitB]] = 0;
49         enabled |= (1 << ids[WaitB]);
50         active |= (1 << ids[WaitB]);
51     } while (0);
52
53     // PARE(0, ABmain, id2b(WaitA) | id2b(WaitB));
54     do {
55         pc[cid] = &&ABmain;
56         pr[cid] = 0;
57         descsc[cid] = (1 << ids[WaitA]) | (1 << ids[WaitB]);
58         goto _L.dispatch ;
59     } while (0);
60
61     WaitA: // AWAIT(A);
62     do {
63         goto _LL74;
64     _LL74: if (!(signals & (1 << A))) {
65         _LL74: pc[cid] = &&.L74;
66         goto _L.PAUSEG; }
67     } while (0);
68
69     // TERM;
70     do { goto _L.TERM; } while (0);
71
72     WaitB: // AWAIT(B);
73     do {
74         goto _LL77;
75     _LL77: if (!(signals & (1 << B))) {
76         _LL77: pc[cid] = &&.L77;
77         goto _L.PAUSEG; }
78     } while (0);
79
80     // TERM;
81     do { goto _L.TERM; } while (0);
82
83     ABmain: // JOIN;
84     do {
85         _L80: if (((enabled & descsc[cid]) == 0)) {
86             goto _LL80; }
87         pc[cid] = &&.L80;
88         goto _L.PAUSEG;
89     _LL80: (void) 0;
90     } while (0);
91
92     // EMIT(O);
93     do {
94         signals |= (1 << O);
95     } while (0);
96
97     // TERM;
98     do {
99         goto _L.TERM;
100    } while (0);
101
102    ABOmain: // AWAIT(R);
103    do {
104        goto _LL84;
105    _LL84: if (!(signals & (1 << R))) {
106        _LL84: pc[cid] = &&.L84;
107        goto _L.PAUSEG; }
108    } while (0);
109
110    // TRANS(ABO);
111    do {
112        enabled &= ~descsc[cid];
113        active &= ~descsc[cid];
114        goto ABO;
115    } while (0);
116
117    // TICKEND;
118    _L.TICKEND: return (enabled != (1 << ids[TickEnd]));
119
120    // Dispatcher, part of TICKEND macro
121    _L.TERM:   enabled &= ~(1 << cid);
122    _L.PAUSEG: active &= ~(1 << cid);
123    _L.dispatch : _asm volatile (" bsr _%1,%0\n" : "=r" (cid) : "
124        c" (active) );
125    goto *pc[cid];
126 }

```

Figure 3.6: ABRO tick function after magno expansion, annotated and reformatted.

```

1  int tick(int islnit)
2  {
3      // Thread ids: AB=1, WaitB=2, WaitA=3, Main=4
4
5      // TICKSTART(islnit);
6      if ( islnit ) {
7          tickCnt = 0;
8          pc[TickEnd] = &&_L_TICKEND;
9          pr[TickEnd] = 0;
10         enabled = (1 << ids[TickEnd]);
11         active = enabled;
12         cid = ids[Main];
13         enabled |= (1 << cid);
14         active |= (1 << cid);
15     } else {
16         active = enabled;
17         goto _L_dispatch;
18     };
19
20     ABO: // PAR(0, AB, ids[AB]);
21     pc[ids[AB]] = &&AB;
22     pr[ids[AB]] = 0;
23     enabled |= (1 << ids[AB]);
24     active |= (1 << ids[AB]);
25
26     // PARE(0, ABomain, id2b(AB) | id2b(WaitA) | id2b(WaitB));
27     pc[cid] = &&ABomain;
28     pr[cid] = 0;
29     descsc[cid] = (1 << ids[AB]) | (1 << ids[WaitA]) | (1 << ids[
        WaitB]);
30     goto _L_dispatch;
31
32     AB: // PAR(0, WaitA, ids[WaitA]);
33     pc[ids[WaitA]] = &&WaitA;
34     pr[ids[WaitA]] = 0;
35     enabled |= (1 << ids[WaitA]);
36     active |= (1 << ids[WaitA]);
37
38     // PAR(0, WaitB, ids[WaitB]);
39     pc[ids[WaitB]] = &&WaitB;
40     pr[ids[WaitB]] = 0;
41     enabled |= (1 << ids[WaitB]);
42     active |= (1 << ids[WaitB]);
43
44     // PARE(0, ABmain, id2b(WaitA) | id2b(WaitB));
45     pc[cid] = &&ABmain;
46     pr[cid] = 0;
47     descsc[cid] = (1 << ids[WaitA]) | (1 << ids[WaitB]);
48     goto _L_dispatch;
50     WaitA: // AWAIT(A);
51     goto _LL74;
52     _L74: if (!(signals & (1 << A))) {
53         _LL74: pc[cid] = &&_L74;
54         goto _L_PAUSEG;
55     }
56
57     // TERM;
58     goto _L_TERM;
59
60     WaitB: // AWAIT(B);
61     goto _LL77;
62     _L77: if (!(signals & (1 << B))) {
63         _LL77: pc[cid] = &&_L77;
64         goto _L_PAUSEG;
65     }
66
67     // TERM;
68     goto _L_TERM;
69     ABmain: // JOIN;
70     _L80: if (((enabled & descsc[cid]) == 0)) {
71         goto _LL80;
72     }
73     pc[cid] = &&_L80;
74     goto _L_PAUSEG;
75     _LL80:
76
77     // EMIT(O);
78     signals |= (1 << O);
79
80     // TERM;
81     goto _L_TERM;
82
83     ABomain: // AWAIT(R);
84     goto _LL84;
85     _L84: if (!(signals & (1 << R))) {
86         _LL84: pc[cid] = &&_L84;
87         goto _L_PAUSEG;
88     }
89
90     // TRANS(ABO);
91     enabled &= ~descsc[cid];
92     active &= ~descsc[cid];
93     goto ABO;
94
95     // TICKEND;
96     _L_TICKEND: return (enabled != (1 << ids[TickEnd]));
97
98     // Dispatcher, part of TICKEND macro
99     _L_TERM:   enabled &= ~(1 << cid);
100    _L_PAUSEG: active &= ~(1 << cid);
101    _L_dispatch : __asm volatile (" bsr _%1,%0\n" : "=r" (cid) : "
        c" (active) );
102    goto *pc[cid];
103 }

```

Figure 3.7: ABRO tick function after macro expansion, annotated and stripped of empty statements.

2. the dispatcher can be implemented with a Bit Scan Reverse (BSR) assembler instruction embedded in the code. This instruction is accessible on x86 architectures when using `gcc`.

In this minimal files version, the *APP.c* file must define a `main` function.

To produce an executable, it suffices to compile just *APP.c*. For example, “`gcc PCO.c -o PCO`” produces an executable `PCO`.

The extended files variant

This variant should be used if

- an alternative dispatcher is required, because
 - the application needs thread priorities, or
 - BSR is not available,
- or if one wants the convenience of using a pre-defined `main` function that for example compares the output of the tick function with a given sequence of reference outputs.

This extended files variant uses the following additional file:

sc.c A C file that contains the `main` function, a dispatcher function (`selectCid`), and an auxiliary function for tracing (`vec2names`) that converts a bit vector to a string of thread or signal names.

Note that the `main` function assumes that signals are used, and hence calls signal-related functions that must be provided by *APP.c* (see Section 3.4.2). This means that when the extended files variant is used, for example, because the application uses thread priorities and hence an alternative dispatcher function is needed, *APP.c* must define these signal-related functions (which can be empty). This is slightly awkward and could be avoided for example by spreading the functions in *sc.c* across several files. Another alternative would be to pre-define alternative `main` functions (or rather functions called by `main`, which in turn can be selected via a macro mechanism in *APP.c*, similar to the selection of the appropriate dispatcher). However, to keep things simple, the functions are at this point all in *sc.c*.

To produce an executable, *sc.c* and *APP.c* must be compiled and linked. For example, “`gcc ABRO.c sc.c -o ABRO`” produces an executable `ABRO`.

The files *sc.c* and *sc.h* are part of the SC software package, *APP.c* must be written by the SC programmer.

3.4.2 Functions

Minimal files variant

In the minimal files variant, *APP.c* must not provide any specific function—except, as usual in C, a `main` function. However, it is good practice to modularize the program by providing the following function:

tick() This function is the top-level function that describes the behavior of the application. One call to `tick` completes when all active threads have reached the end of a logical tick (indicated by the `PAUSE` operator) or have terminated (indicated by `TERM`). The `main` function, defined in `sc.c`, calls `tick` repeatedly, until all threads defined in `tick` have terminated.

Extended files variant

If the `main` function provided by `sc.c` is used, the `tick` function, described in Section 3.4.2, is not optional, but mandatory, as it is called by `main` defined in `sc.c`. In addition, `main` calls the following functions, which therefore must be provided in `APP.c`:

getInputs() This function is called before `tick` is called and defines input signals. If no signals are used, this function is empty.

checkOutputs() This function is called after the `tick` function and can be used to define reference outputs. These are then compared with the outputs actually computed. If no signals are used, this function is also empty.

printval(int id) A function to print valued signal, with index `id`. If no valued signals are used, this function is empty.

3.4.3 Types

Minimal files variant and no signal usage

An SC program has to define the following type:

idtype An enumeration type that declares the *thread names*. This must contain the name `TickEnd`.

`TickEnd` is a special, degenerated thread that does nothing but finish a tick. This is implemented by assigning its program counter the label defined by the `TICKEND` operator (see Section 3.3.1). The `TickEnd` thread should only execute when no other thread is active anymore. It therefore must be assigned the lowest thread id (statically, in `idtype`, see Section 3.4.3), and the lowest priority (at run time, by the `TICKSTART` operator, see Section 3.3.1).

We here exploit that C enumeration types correspond to a sequence of integers, starting at 0 and increasing by 1. Thus `idtype` serves as a mapping from *thread names*, used in the SC program, to *thread identifiers*. These identifiers in turn serve as indices to thread-related information, in particular their *thread id* (via the `ids` array, see below). Thread identifiers are unique to each thread occurring in the program, implicitly defined via the `idtype`. In contrast, thread ids may be shared between threads, as long as these threads cannot be concurrent. In other words, thread identifiers have to be unique at compile time (statically), whereas thread ids may be shared, but have to be unique at run time (dynamically).

In the examples used here, there is no sharing of thread ids, as all threads used may be concurrent. Furthermore, it is often (but not always) the case that the thread id is identical to the thread identifier.

Extended files variant, or usage of signals

In the extended files variant, or if signals are used, *APP.c* also must define the following type:

signaltype An enumeration type that declares the signal names.

3.4.4 Variables

Minimal files variant and no signal usage

An SC program has to define the following variables:

idHi Defines the highest thread id in use.

ids Integer array that maps thread indices to ids. This must map thread TickEnd (to be included in the **idtype**, see Section 3.4.3) to id 0.

id2threadname[] An array of strings that maps thread ids to thread names. This should correspond to the **idtype** enumeration type (see Section 3.4.3).

Extended files variant, or usage of signals

In the extended files variant, or if signals are used, *APP.c* also must define the following variable:

s2signame An array of strings that maps signal ids to signal names. This should correspond to the defined **signaltype** enumeration.

Furthermore, in the extended files variant, *APP.c* must define the following variables, which are used by **main**:

runMax The number of runs to be executed.

tickMax The maximal number of ticks to execute per run.

Chapter 4

Further Examples

This chapter contains a selection of further examples provided by André [3].

4.1 Count2Suspend

Count2Suspend [3, Fig. 8-5], shown in Fig. 4.1, illustrates the use of suspension. The 2-bit counter in macrostate `Cnt2` counts up whenever input signal `T` is present—except when `inhib` suspends (“freezes”) operation of `Cnt2`. The `Main87,SUSPEND` statement performs the according control of the execution of `Cnt2`, see also p. 24.

Precedence constraints

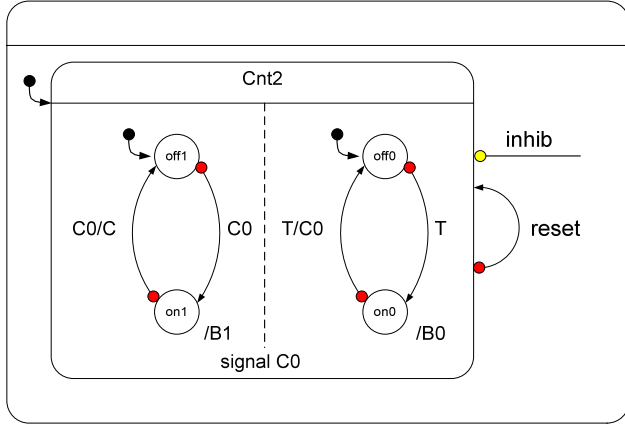
1. Strong abort and suspension on `Cnt2` (*outer-inner*):
`Main` \succ `Off0`, `Main` \succ `Off1`
2. Communication via a signal `C0` (*writer-reader*):
`Off0` \succ `Off1`

These constraints induce a partial order, met by the thread id assignment in the SC code.

4.2 Exits

Exits [3, Fig. 8-8], shown in Fig. 4.2, illustrates the handling of exit actions. These are implemented with the `CALL` operator, which calls exit actions unconditionally, and `ISATCALL`, which calls exit actions if the corresponding state is active (and now gets aborted). See also the descriptions of these operators on p. 29.

Alternative tick functions for `Exits` are shown in Fig. 4.3. The code shown in Fig. 4.3a differs from Fig. 4.2b only in that the shorthand `ISATCALL` is expanded into separate `ISAT` and `CALL` operations. The code in Fig. 4.3b inlines the exit actions. This violates the Write-Things-Once principle, but in this case makes the code shorter, as the exit actions consist of simple `EMIT` operations. However, rather surprisingly, this inlining actually degrades performance, by about 10%.



(a) SyncChart

```

1  TICKSTART(isInit, 3);
2
3  Cnt2:
4  FORK(Off1, 1);
5  FORK(Off0, 2);
6  FORKE(Cnt2Main);
7
8  Off1:
9  while (1) {
10   AWAIT(C0);
11   do {
12     EMIT(B1);
13     PAUSE;
14   } while (!PRESENT(C0));
15   EMIT(C);
16 }
17
18 Off0:
19 while (1) {
20   AWAIT(T);
21   do {
22     EMIT(B0);
23     PAUSE;
24   } while (!PRESENT(T));
25   EMIT(C0);
26 }
27
28 Cnt2Main:
29 while (1) {
30   PAUSE;
31   if (PRESENT(reset))
32     TRANS(Cnt2);
33   SUSPEND(PRESENT(inhib));
34 }
35
36 TICKEND;

```

(b) SC tick function

Figure 4.1: The Count2Suspend example.

Precedence constraints

1. Strong abort on M0 (*outer-inner*):

$\text{Main} \succ \text{M10}, \text{Main} \succ \text{M2}, \text{Main} \succ \text{M11}$

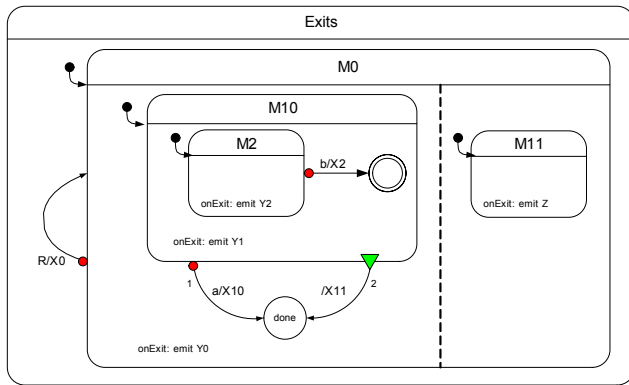
We meet this by assigning **Main** the highest thread id (4). Furthermore, as thread **M10**, a child of **Main**, will eventually raise its priority to 1, **Main** is also assigned this priority, in $\text{Main}_{6, \text{FORKE}}$, after forking **M10**.

2. Strong abort on M10 (*outer-inner*):

$\text{M10}_{27, \text{PRESENT}(A)} \succ \text{M2}$

This is met by assigning **M10** priority 1, with $\text{M10}_{20, \text{JPPAUSE}(1, \dots)}$, while **M2** has priority 0. Note that this strong abort is not immediate, but delayed. Hence it is part of the depth of **M10** (see Sec. 3.2.3), and it is sufficient if **M10** enters its depth with priority 1, but not its surface (label **M10main**).

3. Normal termination on M10 (*inner-outer*):



(a) SyncChart

```

1   TICKSTART(isInit, 4);
2
3   M0:
4     FORK(M10, 2);
5     FORK(M11, 1);
6     PRIO(6);
7     FORKE(M0main);
8
9   M10:
10    FORK(M2, 3);
11    FORKE(M10main);
12
13  M2:
14    PAUSE;
15  M2depth:
16    PRESENTEELSE(B, M2);
17    CALL(M2exit);
18    EMIT(X2);
19    TERM;
20  M2exit:
21    EMIT(Y2);
22    RET;
23
24  L2:
25    PRIO(2);
26  M10main:
27    JPPAUSE(5, M10depth);
28    ISATCALL(3, M2depth, M2exit);
29    CALL(M10exit);
30    EMIT(X11);
31    TRANS(Done);
32  M10depth:
33    PRESENTEELSE(A, L2);
34    ISATCALL(3, M2depth, M2exit);
35    CALL(M10exit);
36    EMIT(X10);
37    TRANS(Done);
38  Done:
39    HALT;
40  M10exit:
41    EMIT(Y1);
42    RET;
43
44  M11:
45    HALT;
46  M11exit:
47    EMIT(Z);
48    RET;
49
50  M0main:
51    AWAIT(R);
52    ISATCALL(3, M2depth, M2exit);
53    ISATCALL(5, M10depth, M10exit);
54    CALL(M11exit);
55    EMIT(Y0); // Only place to call exit action of M0
56    EMIT(X0);
57    TRANS(M0);
58
59  TICKEND;

```

(b) SC tick function

Figure 4.2: The Exits example. See Fig. 4.3 for alternative tick functions.

```

1 // Thread ids: M11=1, M10=2, M2=3, Main=4
2   TICKSTART(islnit);
3 M0:  PAR(0, M10, ids[M10]);
4     PAR(0, M11, ids[M11]);
5     PARE(1, M0main, id2b(M10) | id2b(M11) | id2b(
6       M2));
7 M10: PAR(0, M2, ids[M2]);
8     PARE(0, M10main, id2b(M2));
9
10 M2:  PAUSE;
11 M2depth:PRESENT(B, M2);
12     CALL(M2exit, L1);
13 M2exit: EMIT(Y2);
14     RET;
15 L1:   EMIT(X2);
16     TERM;
17
18 L2:   PRIO(0);
19 M10main:JPPAUSE(1, M10depth);
20     ISAT(ids[M2], M2depth, L4);
21     CALL(M2exit, L4);
22 L4:   CALL(M10exit, L5);
23 M10exit:EMIT(Y1);
24     RET;
25 L5:   EMIT(X11);
26     TRANS(Done);
27 M10depth:PRESENT(A, L2);
28     ISAT(ids[M2], M2depth, L7);
29     CALL(M2exit, L7);
30 L7:   CALL(M10exit, L8);
31 L8:   EMIT(X10);
32     TRANS(Done);
33 Done:  HALT;
34
35 M11:  HALT;
36 M11exit:EMIT(Z);
37     RET;
38
39 M0main: AWAIT(R);
40     ISAT(ids[M2], M2depth, L10);
41     CALL(M2exit, L10);
42 L10:  ISAT(ids[M10], M10depth, L11);
43     CALL(M10exit, L11);
44 L11:  CALL(M11exit, L12);
45 L12:  CALL(M11exit, L12);
46     EMIT(X0);
47     TRANS(M0);
48
49 TICKEND;

```

(a) Tick function without ISATCALL

```

1 // Thread ids: M11=1, M10=2, M2=3, Main=4
2   TICKSTART(islnit);
3 M0:  PAR(0, M10, ids[M10]);
4     PAR(0, M11, ids[M11]);
5     PARE(1, M0main, id2b(M10) | id2b(M11) | id2b(
6       M2));
7 M10: PAR(0, M2, ids[M2]);
8     PARE(0, M10main, id2b(M2));
9
10 M2:  PAUSE;
11 M2depth:PRESENT(B, M2);
12     EMIT(Y2);
13     EMIT(X2);
14     TERM;
15
16 L2:   PRIO(0);
17 M10main:JPPAUSE(1, M10depth);
18     ISAT(ids[M2], M2depth, L4);
19     EMIT(Y2);
20 L4:   EMIT(Y1);
21     EMIT(X11);
22     TRANS(Done);
23 M10depth:PRESENT(A, L2);
24     ISAT(ids[M2], M2depth, L7);
25     EMIT(Y2);
26 L7:   EMIT(Y1);
27     EMIT(X10);
28     TRANS(Done);
29 Done:  HALT;
30
31 M11:  HALT;
32
33 M0main: AWAIT(R);
34     ISAT(ids[M2], M2depth, L10);
35     EMIT(Y2);
36 L10:  ISAT(ids[M10], M10depth, L11);
37     EMIT(Y1);
38 L11:  EMIT(Z);
39     EMIT(Y0);
40     EMIT(X0);
41     TRANS(M0);
42
43 TICKEND;

```

(b) Tick function with inlined exit actions

Figure 4.3: Alternative variants for the SC tick function of the Exits example (Fig. 4.2).

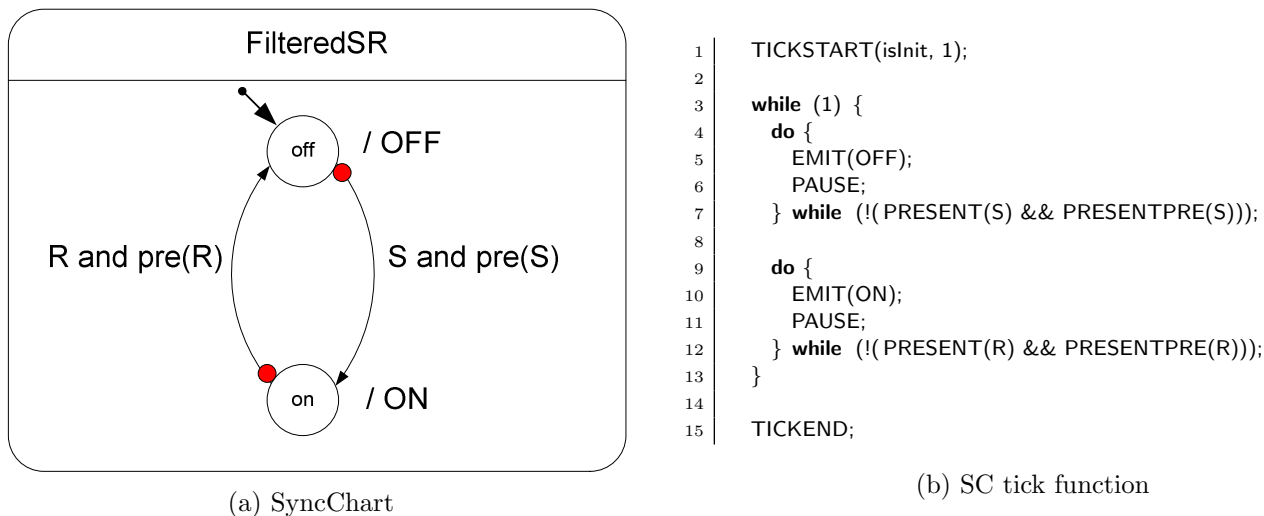


Figure 4.4: The FilteredSR example.

$M2 \succ M10_{20, JPPAUSE}$

This is met by assigning $M10$ a lower id than $M2$, and executing both $M2$ and the test for normal termination with priority 0. To ensure the latter, we set the priority of $M10$ at $M10_{9, FORKE}$, for the test in the tick when $M10$ is entered, and at $M10_{19, PRIO(0)}$, for the test in subsequent ticks.

4. Precedence of strong abortion over normal termination on $M10$ (*transition-number*):

$M10_{27, PRESENT(A)} \succ M10_{20, JPPAUSE}$

This intra-thread structural constraint is met by ordering the operations in the program accordingly.

4.3 FilteredSR

FilteredSR [3, Fig. 8-18], shown in Fig. 4.4, illustrates the use of PRE on pure signals.

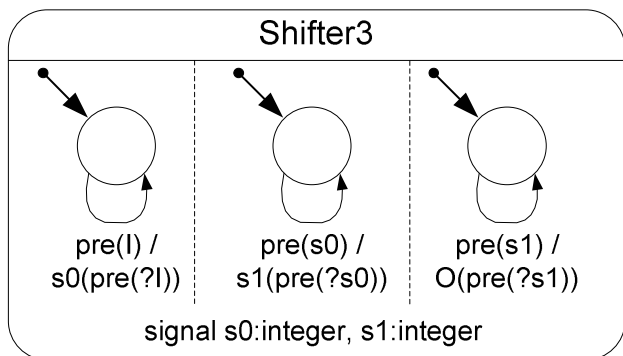
We here use *signal expressions*, in this case signal conjunction. In full regular C, such expressions can be built up the usual way with C's logical operators (`!`, `&&`, `||`). If we want to restrict ourselves to plain SC operators, we can encode these expressions with control flow, as is done here.

Precedence constraints As there is just the Main thread and each state has just one outgoing transition, there are no precedence constraints.

4.4 Shifter3

Shifter3 [3, Fig. 8-19], shown in Fig. 4.5, illustrates the use of PRE on valued signals.

There are three top-level concurrent threads. There are two alternatives to implement this:



(a) SyncChart

```

1 | TICKSTART(isInit, 1);
2 |
3 | FORK(Shift1, 2);
4 | FORK(ShiftO, 3);
5 |
6 | while (1) {
7 |   PAUSE;
8 |   if (PRESENTPRE(l))
9 |     EMITINT(S0, VALPRE(l));
10 | }
11 |
12 | Shift1 :
13 | while (1) {
14 |   PAUSE;
15 |   if (PRESENTPRE(S0))
16 |     EMITINT(S1, VALPRE(S0));
17 | }
18 |
19 | ShiftO :
20 | while (1) {
21 |   PAUSE;
22 |   if (PRESENTPRE(S1))
23 |     EMITINT(O, VALPRE(S1));
24 | }
25 |
26 | TICKEND;

```

(b) SC tick function

Figure 4.5: The Shifter3 example.


1. The Main thread is interpreted as the top-level macrostate **Shifter3**. It spawns off three children (**Shift0**, **Shift1**, **ShiftO**), and then terminates (with **TERM**), as it has nothing more to do.
2. The Main thread is interpreted as one of the concurrent subthreads of macrostate **Shifter3**, say **Shift0**. It spawns off two concurrent threads (**Shift1**, **ShiftO**); see also the last note in Sec. 3.3.1.

We here implement the second alternative, as it reduces the number of required threads by

1. Further notes on this methods of spawning concurrent threads:

- This is a fork without a **FORKE**, hence there is no call to the dispatcher before **Main** delves into the code region implementing the **Shift0** substate. In this case this is unproblematic, as there are no precedence constraints to be obeyed (see below). In general, if we use this technique of spawning concurrent threads (instead of child threads) and there is a particular thread that must be executed next, we must make sure that the current thread takes on the role this particular thread.
- If there are no precedence constraints, it is suggested to let the spawning thread continue with the code region that follows after the fork, in this case **Shift0**. If the spawning thread can start at the beginning of that code region, this saves a **GOTO**. In this example we do not have this saving, as we still have to jump to the **Shift0** label.

A further optimization implemented here: Starting the code fragment belonging to state **Shift0** with its depth (**Shift0depth**) allows to save a final **GOTO** by folding it into **PAUSE(Shift0depth)**. Similarly for the other concurrent states.



(a) SyncChart

Instant	1	2	3	4	5	6	7	8	9	10	11	12	13
T	-	+	-	+	-	+	-	-	+	-	+	+	-
B0	-	+	-	-	-	-	-	-	+	-	-	-	-
B1	-	-	-	+	-	-	-	-	-	-	+	-	-
C	-	-	+	+	-	+	+	-	+	+	-	+	+

(b) Sample execution trace [3]

```

1 | TICKSTART(isInit, 4);
2 |
3 | FORK(Cnt, 3);
4 | FORKE(Mod3CntMain);
5 |
6 | Cnt:
7 |   FORK(Off1, 1);
8 |   FORK(Off0, 2);
9 |   FORKE(CntMain);
10 |
11 | Off1:
12 |   AWAIT(C);
13 |   SUSTAIN(B1);
14 |
15 | Off0:
16 |   while (1) {
17 |     PAUSE;
18 |     EMIT(B0);
19 |     PAUSE;
20 |     EMIT(C);
21 |   }
22 |
23 | CntMain:
24 |   while (1) {
25 |     PAUSE;
26 |     if (PRESENTPRE(C))
27 |       TRANS(Cnt);
28 |   }
29 |
30 | Mod3CntMain:
31 |   while (1) {
32 |     PAUSE;
33 |     SUSPEND(!PRESENT(T));
34 |   }
35 |
36 | TICKEND;

```

(c) SC tick function

Figure 4.6: The PreAndSuspend example.

Precedence constraints There are no precedence constraints, even though there is signal-based communication via S_0 and S_1 ; all triggers are delayed (via PRE), hence any written signals will not be read before the next tick.

4.5 PreAndSuspend

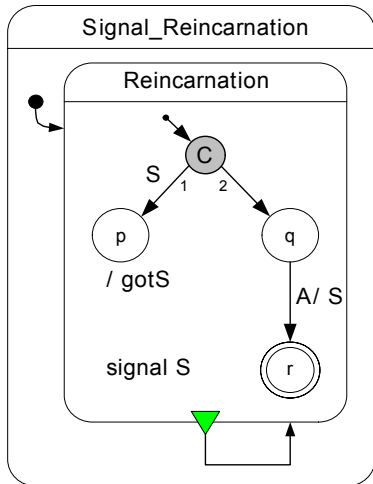
PreAndSuspend [3, Fig. 8-20], shown in Fig. 4.6, illustrates the proper handling of PRE in conjunction with suspension and local signals. See also the execution trace in Fig. 4.6b.

For the signal C , declared locally at state Mod3Cnt , $\text{PRE}(C)$ refers to the presence value of C in the previous tick in which Mod3Cnt was active (not suspended).

Precedence constraints

1. Suspension of Mod3Cnt (*outer-inner*):

Main \succ Cnt



(a) SyncChart

```

1 | TICKSTART(isInit, 1);
2 |
3 | while (1) {
4 |   SIGNAL(S);
5 |   if (PRESENT(S)) {
6 |     SUSTAIN(gotS);
7 |   }
8 |   AWAIT(A);
9 |   EMIT(S);
10 | }
11 |
12 | TICKEND;

```

(b) SC tick function

Figure 4.7: The Reincarnation example.

2. Strong abort on Cnt (*outer-inner*):
 $\text{Cnt} \succ \text{Off0}, \text{Cnt} \succ \text{Off1}$
3. Communication via C (*writer-reader*)
 $\text{Off0} \succ \text{Off1}$

These constraints induce a partial order, met by the thread id assignment in the SC code.

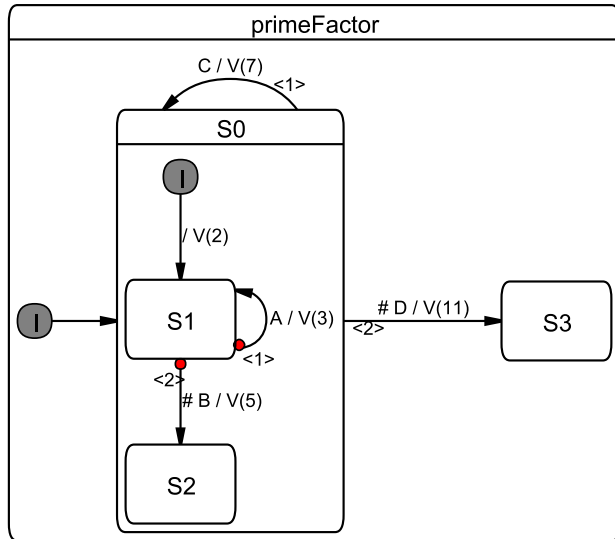
4.6 Reincarnation

Reincarnation [3, Fig. 8-22], shown in Fig. 4.7, illustrates the **SIGNAL** instruction to handle signal reincarnation.

The canonical encoding in SC would have the **Main** thread spawn off an inner thread that computes the behavior of the **Reincarnation** state. However, as the macrostate **Reincarnation** only has a normal termination transition attached to it, all the **Main** thread does is to reenter itself once **Reincarnation** has terminated. This can be streamlined by transferring control from all terminating states within **Reincarnation** to the entry of **Reincarnation**. In this case, there is just one such terminating state, namely **r**. Hence, there is no need anymore for a separate parent thread that checks for termination. In other words, the **Main** thread can directly run the **Reincarnation** state.

Precedence constraints

1. Normal termination of **Reincarnation** (*inner-outer*):
 This gets folded into the sequential code of **Main**.
2. There is also the conditional pseudo state with two outgoing transitions, which in principle constitutes a transition-number precedence. In this case, this corresponds to a simple if-then-else branch, encoded in $\text{PRESENT}(S, Q)_5$.



(a) SyncChart

```

1  TICKSTART(isInit, 1);
2
3  S0:
4  FORK(S1, 2);
5  FORKE(S0main);
6
7  S1:
8  EMITINTMUL(V, 2);
9  while (1) {
10   if (PRESENT(B)) {
11     EMITINTMUL(V, 5);
12     HALT;
13   }
14   PAUSE;
15   if (PRESENT(A))
16     EMITINTMUL(V, 3);
17 }
18
19 S0main:
20 while (1) {
21   if (PRESENT(D)) {
22     EMITINTMUL(V, 11);
23     TRANS(S3);
24   }
25   PAUSE;
26   if (PRESENT(C)) {
27     EMITINTMUL(V, 7);
28     TRANS(S0);
29   }
30 }
31
32 S3:
33 HALT;
34 TICKEND;

```

(b) SC tick function

Figure 4.8: The PrimeFactor example.

4.7 PrimeFactor

PrimeFactor [3, Fig. 8-25], shown in Fig. 4.8, illustrates the use of valued signals and the proper handling of reincarnation/schizophrenia.

- S0 needs no JOIN, as it never terminates normally.
- S2: PAUSE(S2) encodes a final, but non-terminating state; this corresponds to Esterel's halt.

Precedence constraints

1. Normal termination transitions of S0 (*inner -outer*):

S1 \succ Main

This is met by thread id assignment.

2. Ordering of transitions of S0 (*transition-number*):

$\text{Main}_{23,\text{PRESENT}(\text{C})} \succ \text{Main}_{19,\text{PRESENT}(\text{D})}$

This is met by statement ordering.

3. Ordering of transitions of **S1** (*transition-number*):

$\text{S1}_{13,\text{PRESENT}(\text{A})} \succ \text{S1}_{9,\text{PRESENT}(\text{B})}$

This is also met by statement ordering.

To summarize, all scheduling constraints are handled by proper ordering of the transition predicate tests, and by the fact that the id of the inner state (**S0**, id 1) is higher than the priority of the surrounding root thread.

Chapter 5

Related Work

Statechart variants Since the original Statecharts proposal [14], numerous dialects of Statecharts have been developed and Statecharts have also been incorporated into the Unified Modeling Language (UML). Statecharts are supported by a multitude of modeling tools; the first commercial tool was Statemate [14], other established tools today are SCADE/Esterel Studio (Esterel Technologies), Matlab/Simulink/Stateflow (The Mathworks), ASCET (ETAS), or Rational Rose (IBM). These tools all implement the fundamental Statechart concepts of concurrency, hierarchy, and signal broadcast. However, their underlying MoCs also have some subtle, but important differences, in particular regarding their handling of concurrency. In fact, while the visual *syntax* of Statecharts appears fairly simple and straightforward, it is not at all obvious what their *semantics* should be [6].

Most Statechart dialects in use today, including UML Statecharts, have the limitation that they do not offer deterministic concurrency. Concurrent states are often implemented as concurrent threads, thus inheriting the non-determinism associated with thread scheduling [17]. This can be alleviated by adopting a *strictly synchronous* semantics, which precisely states how computations should proceed [7]. The synchronous MoC implements the *synchrony hypothesis*, which abstracts from concrete run-time behavior by assuming that the computation of a reaction does not take any time. The *strict* interpretation of synchrony also adopts a fixed point semantics, which means that the status of events (sometimes also referred to as signals) must be consistent throughout a reaction. Strictly synchronous Statechart dialects are Argos [20] and SyncCharts [3], also known as Safe State Machines (SSMs). There is also the *loose* interpretation of synchrony, which does assume that physical time does not progress during a reaction, but does not require that the system progresses along fixed points. Instead, it allows the presence/absence status of events to change during a reaction. This is less restrictive than strict synchrony, but degrades compositionality and may lead to infinite computations. The original Statechart proposal implemented loose synchrony.

Expressing Statecharts in C/C++ As mentioned in the introduction, it is already common practice to express Statecharts in a classical programming language. Samek describes how to express UML Statecharts in C/C++ [29]. As in UML Statecharts, this approach does not provide deterministic concurrency. Wagner *et al.* describe how to implement FSMs in C [34], but these are flat automata without any concurrency.

Synchronous language extensions There have been several proposals to extend traditional programming languages by synchronous constructs. Reactive C [11] is an extension of C inspired by Esterel. It employs the concepts of computational instants (ticks) and preemptions, but does not provide true concurrency; Reactive C’s `merge` operator emulates concurrency by running threads sequentially, in their textual order.

FairThreads [8] extend this by true concurrency, implemented via native threads. They also offer macros to express automata. SC does not use native threads, but does its own, light-weight thread book keeping. Another difference is that the signal mechanism provided by FairThreads does not allow reaction to signal absence, whereas SC does allow this (see `grcbal3`).

The Esterel-C Language (ECL) [16] is another proposal to extend C by Esterel-like constructs. A C program is annotated with Esterel-like constructs for signal handling and reactive control flow, and from this program the ECL compiler derives an Esterel part and a purely sequential C part. SC is in the same spirit of annotating C with synchronous operators, but differs from ECL in that it does not resort to a separate language (Esterel).

Another recent proposal for a synchronous extension of C is Precision Timed C (PRET-C) [2, 26, 27]. PRET-C focuses on temporal predictability and assumes a target architecture with specific support for thread scheduling and abort handling. PRET-C provides a minimal set of C extensions, namely a concurrency operator, which runs threads with static priorities, a delayed abortion operator, and an EOT operator that delineates ticks. An associated compiler produces a corresponding intermediate format, the Timed Concurrent Control Flow Graph, where each thread at each EOT tests whether it is aborted or not with `Checkabort` nodes.

Lusteral, presented by Mendler and Pouzet [21], also tries to capture the essence of synchronous programming in a small number of operators. It combines elements of the synchronous languages Lustre, Esterel and Signal and embeds them in Haskell. As this is a functional language, it allows to express the semantics of the Lusteral operators nicely as higher-order functions.

Compiling synchronous programs As SC expresses synchronous, control-oriented concurrency by means of a—ultimately sequential—C program, executing an SC program raises similar issues as they arise when synthesizing a synchronous language into sequential code. There have been numerous proposals for this, in particular for the Esterel language [23, 10]. It is a common procedure to translate an Esterel program into a C program, but the resulting C program usually bears little resemblance to the original Esterel program. For example, the C code might be a flat automaton, or it might simulate a hardware circuit.

Probably the closest in spirit to SC is the BAL virtual machine [10], which proposes a high-level ISA that captures the Esterel semantics as closely as possible; see also the comparison done in Chapter 6.

Another interesting approach is the dynamic list code generation [10], which produces C code that executes concurrently running threads by dispatching small groups of instructions that can run without a context switch. These blocks are dispatched by a scheduler that uses linked lists of pointers to code blocks that will be executed in the current cycle. While the fundamentals of that code generation are very different from the SC approach, their use of pointers and `gcc`’s computed `gotos` has inspired the label-based “coarse grain program counter” approach presented here.

The PRET and SHIM programming models As discussed in Sec. 2.1, SC is also related to the programming model proposed for the Precision Timed Architecture (PRET) proposed by Edwards and Lee [19], but does not rely on low-level timing for synchronization.

Another related programming model is SHIM [31], proposed for software/hardware integration, which provides Kahn process networks with CSP-like rendezvous communication and exception handling. It uses a separate compiler to convert a SHIM program into sequential C code. SHIM, like SC, has been inspired by synchronous languages, but it does not use a synchronous programming model, instead relying on communication channels for synchronization.

Code generation from Statecharts/SyncCharts As SC can be used as a target format when synthesizing Statecharts into a sequential program, this work also relates to code generation from Statecharts. Three different methods of compiling Statecharts are common: compilation into an object oriented language using the state pattern [1], dynamic simulation [35], and flattening into finite state machines. Since flattening can suffer from state explosion, often a combination of flattening and dynamic simulation is used. All of these methods incur relatively high overhead and typically make use of a run time system to achieve concurrency, and usually the result is not deterministic.

For SyncCharts, it is also possible to translate the Statechart model into an equivalent textual Esterel program [12]. Such a translation was proposed by André [4] together with the initial definition of SyncCharts and their semantics. This transformation, with additional unpublished optimizations, is implemented in Esterel Studio. The resulting Esterel program can then be translated into software or hardware [23]. As discussed in Chapter 6, this path via Esterel to C is here used for experimental comparison. A drawback of this approach is that the original structure of SyncCharts cannot always be preserved in the Esterel code, as Esterel does not allow the arbitrary control flow that can be expressed by SyncChart transitions; this also can induce the need for additional signals, to encode the next active state. This structure is even less preserved in a C program compiled from the Esterel program.

Compilation for reactive processors One approach to synthesize SyncCharts into a textual program that does preserve the original structure is to generate code directly for a reactive processor [33], as done by the state machine to KEP compiler (**smakc!**) [30]. Unlike the instruction set architecture (ISA) of traditional processors, which provide only sequential control flow operators such as branches and jumps, the ISA of reactive processors directly expresses concurrency and preemption. The **smakc!** compiler targets the Kiel Esterel Processor [18], which implements synchronous concurrency via multi-threading. This multi-threading approach, which is also realized for example in the StarPro processor [36], has the advantage of allowing high degrees of concurrency without excessive resource requirements.

The SC operators have been inspired by the KEP ISA, and adopt the KEP's mechanism of priority-based multi-threading. However, the SC operators have been developed with SyncCharts in mind, rather than Esterel, and they make minimal assumptions on the execution platform. The main resulting differences between SC and the KEP ISA are:

- SC provides a **TRANS** operator that implements an arbitrary state transition;
- SC does not provide Esterel's exception handling via traps;
- SC does not rely on special watcher units to implement aborts.

A motivation for the KEP's watcher units was to avoid `Checkabort` instructions [28, 26], as these introduce an overhead—both in terms of code size as execution speed—at each tick, in all threads, proportional to the abort nesting depth. Interestingly, SC needs neither watchers nor `Checkaborts`, by giving parent threads the power to abort their descendants with the `TRANS` operator.

Chapter 6

Experimental Results

Note: the measurements have been performed with SC Release 1.2. In particular code sizes should be improved with the current release.

6.1 Conciseness of SC, Code Size

The main goal in developing SC was to develop a concise embedding of SyncChart behavior into C. It is difficult to measure “conciseness” precisely, as this compares a visual language against a textual one. A better point of reference might be Esterel code. For example, `grcbal3` in Esterel takes 25 lines (see Fig. 2.2a); in SC, it takes 28 lines (Fig. 2.2c). This indicates a comparable level of conciseness, which is remarkable in that the SC operators are embedded in the imperative, sequential programming model of C.


Another interesting point of comparison is the BAL VM instruction set, as it has been designed specifically to encode Esterel programs in as little memory as possible [10]. To encode `grcbal3`, BAL uses 74 instructions, of complexity comparable to the SC operators. The SC version makes do with 24 instructions, which—for this example—is more compact by a factor of three. Arguably, these instructions are also easier to relate to an Esterel program or a SyncChart than the BAL assembler. This makes SC an attractive alternative candidate for a VM instruction set.

Fig. 6.1a compares the size of the SC tick functions for a number of benchmarks, taken from André [3], with the size of the C code generated by EsterelStudio. Two synthesis variants are considered, one based on circuit simulation, the other based on GRC. As can be seen, SC is often less than half the size of the synthesized C code.


Fig. 6.1b compares sizes of object code for the tick function. Here, the SC code is larger on average, just in two cases it is slightly better than both E-Studio results. However, considering the sizes of the executable on an x86 architecture, shown in Fig. 6.1c, SC is ahead again. All results were obtained with `gcc -O3`.

6.2 SC Performance

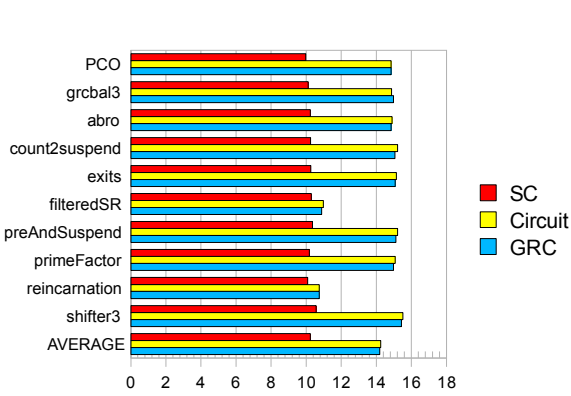
The development of SC has not been motivated primarily by performance concerns, but still it is interesting to see how it compares. On the negative side, SC basically just interprets a SyncChart, it cannot perform any global optimizations or partial evaluations at compile



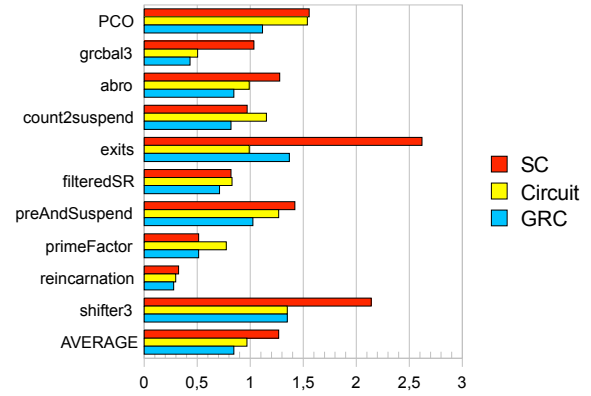
(a) Size of tick function in C source code, line count without empty lines and comments



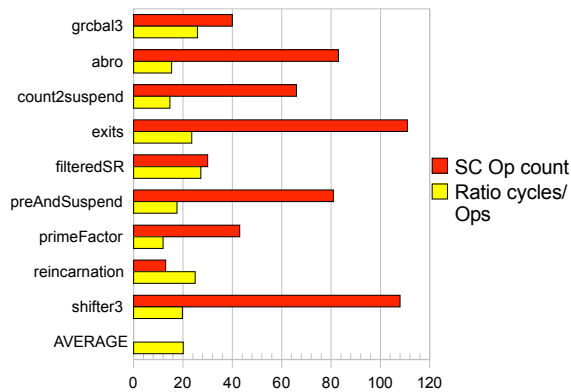
(b) Size of tick function object code, in Kbytes



(c) Size of executable, in Kbytes



(d) Accumulated run times of tick function, in thousands of clock cycles



(e) SC operations count, ratio to clock cycles

Figure 6.1: Comparison of SC with two code synthesis variants of Esterel Studio.

time, as do for example the EsterelStudio synthesis tools. On the positive side, SC code has no scalability problems, neither in terms of code size (like the flat automaton synthesis approach) nor in terms of run time. It only does work that needs to be done, in the sense that no unnecessary code regions are executed. This is different than for example the widely used circuit simulation approach, where always the whole circuit is simulated, irrespective of which regions are active. Furthermore, the SC context switches are very light weight, as 1) each thread requires very little information (see Sec. 2.1), and 2) the dispatcher is fast in typical scenarios (see Sec. 3.2.2). Therefore, SC certainly requires less overhead than a traditional thread-based implementation, where a context switch itself already takes thousands of instructions.

A more challenging point of reference are the monolithic C functions synthesized from SyncCharts. Figure 6.1d compares the run times of the tick functions, on an Intel Core 2 Duo architecture. For the measurements, a representative input trace was executed, outputs were compared against a reference trace, and the execution times of the individual calls to the tick functions were accumulated. Timings were done in numbers of processor cycles, using the x86 `rdtsc` (Read Time Stamp Counter) instruction. The machine runs at 2.4 GHz, so most of the runs took less than 1 μ s. As to be expected, SC does not beat any of the advanced synthesis techniques. In the `exits` example, which makes heavy use of exit actions, modularized into separate procedure calls rather than inlining (and possibly duplicating) them, the performance is even 2–2.5x worse. Overall, however, SC is roughly comparable, and in four of the ten benchmarks it is faster than the Circuit approach. As applications get larger, one should expect that SC stays comparable (at least), as again it does not have scalability problems. Also, one should expect that in practice, the run time of SC programs is dominated by regular C operations, not the SC operators.

A last statistic is shown in Figure 6.1e, which counts the SC operations executed, as listed in the output traces in the appendix. It also shows the ratio to the clock cycles from Fig. 6.1d. The ratio varies somewhat, but on average twenty clock cycles are needed to perform one SC operation.

Chapter 7

Conclusions and Outlook

SyncCharts in C are a light-weight approach to embed deterministic reactive control flow constructs into a widely used programming language. With a relatively small number of primitives it is possible to cover the complete SyncCharts language. The multi-threaded, priority-based approach has been inspired by synchronous reactive processing; hence, originally, this approach required a special compiler and a special architecture to implement. For example, the KEP has watchers that check for preemption in parallel to normal operation, a reactive processing unit that resolves control priorities on the fly, and a dispatcher that selects the next thread for execution at the beginning of each instruction cycle. Therefore, it was not obvious from the onset that it would be possible to achieve the same behavior by isolated SC operators, embedded in regular imperative sequential code, on a standard architecture, at a competitive performance. As it turns out, standard architectures already provide features that can be used to advantage, even if they are not directly available on the C level, such as the x86 `bsr` instruction that can be used for fast dispatching. A number of issues that pose challenges in implementing synchronous programs, such as schizophrenia or reaction to signal absence, are also unproblematic.

Considering the formal semantics of SC, as it is expressed in terms of C, one might take the stance that the semantics of the SC operators is expressed by the C statements they consist of, none of which touch on any of the many semantic uncertainties of C. In terms of mental complexity, this should not be as daunting as one might think; as of SC version 1.2, the file `sc.h` that defines all SC operators (except the general versions of the dispatcher, which are defined as functions in `sc.c`), is 567 lines long (see Listing A.1), of which 173 lines are comments, 49 lines are related to tracing, and 127 lines are empty. This leaves 218 lines of C code that explain what the 12 SC thread operators, 11 signal operators, and 5 sequential control operators do. Still, it should be worthwhile to formalize the semantics at a more abstract level, to allow formal reasoning about them.

SC is freely available, and can be used as is for writing reactive applications in C. However, there are a number of interesting further projects that should be pursued. As already mentioned, SC seems a viable candidate for synthesizing visual SyncCharts into code, especially if traceability is required, or as input language for PRET architectures. It would also be an interesting exercise to add something like a `DEAD` timing primitive [19] to SC. Unlike PRET architectures, traditional architectures probably cannot do this cycle-accurate; however, using something like `nanosleep` or the x86 `rtsc` instruction, it should be possible to get fairly close. One might use this to pad calls of the `tick` function to reduce the reaction jitter, replacing for example the crude call to `sleep` in PCO (Fig. 2.1c, line 15). A related issue is the WCRT

analysis for SC, which could build on earlier work [22,26]. Another question not addressed at all so far is how the SC approach could be used to extract true parallelism from a program, *e. g.* for programming multi-core processors. This should be feasible, *e. g.* by an alternative thread id/priority assignment scheme that expresses when things can be run in parallel; but it is an interesting question how to make this fast and how to minimize global synchronization overheads.

Acknowledgments

Numerous discussions have helped shape SC and this report. I would like to thank in particular Alain Girault, Michael Mendler, Partha Roop, Robert de Simone and Claus Traulsen.

Christian Schneider has conducted the experiments reported on in Chapter 6.

Bibliography

- [1] J. Ali and J. Tanaka. Converting Statecharts into Java code. In *Proceedings of the Fourth World Conference on Integrated Design and Process Technology (IDPT '99)*, Dallas, Texas, June 2000. Society for Design and Process Science (SDPS).
- [2] S. Andalam, P. Roop, A. Girault, and C. Traulsen. PRET-C: A new language for programming precision timed architectures. Technical Report 6922, INRIA Grenoble Rhône-Alpes, 2009. <http://hal.inria.fr/docs/00/39/16/21/PDF/rr.pdf>.
- [3] C. André. Semantics of SyncCharts. Technical Report ISRN I3S/RR-2003-24-FR, I3S Laboratory, Sophia-Antipolis, France, April 2003.
- [4] C. André. Computing SyncCharts reactions. In *SLAP 2003: Synchronous Languages, Applications and Programming, A Satellite Workshop of ECRST 2003*, volume 88, pages 3 – 19, 2004.
- [5] A. W. Appel. *Compiling with Continuations*. Cambridge University Press, 2007.
- [6] M. v. d. Beeck. A comparison of Statecharts variants. In H. Langmaack, W. P. de Roever, and J. Vytupil, editors, *Formal Techniques in Real-Time and Fault-Tolerant Systems*, volume 863 of *Lecture Notes in Computer Science*, pages 128–148. Springer-Verlag, 1994.
- [7] A. Benveniste, P. Caspi, S. A. Edwards, N. Halbwachs, P. L. Guernic, and R. de Simone. The Synchronous Languages Twelve Years Later. In *Proceedings of the IEEE, Special Issue on Embedded Systems*, volume 91, pages 64–83, Jan. 2003.
- [8] F. Boussinot. Fairthreads: mixing cooperative and preemptive threads in C. *Concurrency and Computation: Practice and Experience*, 18(5):445–469, Apr. 2006.
- [9] M. E. Conway. Design of a separable transition-diagram compiler. *Communications of the ACM*, 6(7):396–408, 1963.
- [10] S. A. Edwards and J. Zeng. Code generation in the Columbia Esterel Compiler. *EURASIP Journal on Embedded Systems*, Article ID 52651, 31 pages, 2007.
- [11] Frederic Boussinot. Reactive C: An extension of C to program reactive systems. *Software Practice and Experience*, 21(4):401–428, 1991.
- [12] S. M. G. Berry and J.-P. Rigault. Esterel: Towards a synchronous and semantically sound high-level language for real-time applications. In *IEEE Real-Time Systems Symposium*, pages 30–40, 1983. IEEE Catalog 83CH1941-4.

- [13] D. Harel. Statecharts: A visual formalism for complex systems. *Science of Computer Programming*, 8(3):231–274, June 1987.
- [14] D. Harel, H. Lachover, A. Naamad, A. Pnueli, M. Politi, R. Sherman, A. Shtull-Trauring, and M. Trakhtenbrot. Statemate: A working environment for the development of complex reactive systems. *IEEE Transactions on Software Engineering*, 16(4):403–414, Apr. 1990.
- [15] G. Kahn and D. B. MacQueen. Coroutines and networks of parallel processes. In *IFIP Congress*, pages 993–998, 1977.
- [16] L. Lavagno and E. Sentovich. ECL: a specification environment for system-level design. In *DAC '99: Proceedings of the 36th ACM/IEEE conference on Design automation*, pages 511–516, New York, NY, USA, 1999. ACM Press.
- [17] E. A. Lee. The problem with threads. *IEEE Computer*, 39(5):33–42, 2006.
- [18] X. Li, M. Boldt, and R. von Hanxleden. Mapping Esterel onto a multi-threaded embedded processor. In *Proceedings of the 12th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS'06)*, San Jose, CA, October 21–25 2006.
- [19] B. Lickly, I. Liu, S. Kim, H. D. Patel, S. A. Edwards, and E. A. Lee. Predictable programming on a precision timed architecture. In *Proceedings of Compilers, Architectures, and Synthesis of Embedded Systems (CASES'08)*, Atlanta, USA, Oct. 2008.
- [20] F. Maraninchi and Y. Rémond. Argos: An automaton-based synchronous language. *Computer Languages*, 27(27):61–92, 2001.
- [21] M. Mendler and M. Pouzet. Uniform and modular composition of data-flow & control-flow in the lazy λ -calculus. Presentation at the International Open Workshop on Synchronous Programming (SYNCHRON'08), Aussois, France, Dec. 2008.
- [22] M. Mendler, R. von Hanxleden, and C. Traulsen. WCRT Algebra and Interfaces for Esterel-Style Synchronous Processing. In *Proceedings of the Design, Automation and Test in Europe (DATE'09)*, Nice, France, Apr. 2009.
- [23] D. Potop-Butucaru, S. A. Edwards, and G. Berry. *Compiling Esterel*. Springer, May 2007.
- [24] S. Prochnow, C. Traulsen, and R. von Hanxleden. Synthesizing Safe State Machines from Esterel. In *Proceedings of ACM SIGPLAN/SIGBED Conference on Languages, Compilers, and Tools for Embedded Systems (LCTES'06)*, Ottawa, Canada, June 2006.
- [25] S. Prochnow and R. von Hanxleden. Statechart development beyond WYSIWYG. In *Proceedings of the ACM/IEEE 10th International Conference on Model Driven Engineering Languages and Systems (MoDELS'07)*, Nashville, TN, USA, Oct. 2007.
- [26] P. S. Roop, S. Andalam, R. von Hanxleden, S. Yuan, and C. Traulsen. Tight WCRT analysis for synchronous C programs. In *Proceedings of the International Conference on Compilers, Architecture, and Synthesis for Embedded Systems (CASES'09)*, Grenoble, France, Oct. 2009.

- [27] P. S. Roop, S. Andalam, R. von Hanxleden, S. Yuan, and C. Traulsen. Tight WCRT analysis for synchronous C programs. Technical Report 0912, Christian-Albrechts-Universität Kiel, Department of Computer Science, Kiel, Germany, May 2009.
- [28] P. S. Roop, Z. Salcic, and M. W. S. Dayaratne. Towards Direct Execution of Esterel Programs on Reactive Processors. In *4th ACM International Conference on Embedded Software (EMSOFT'04)*, Pisa, Italy, Sept. 2004.
- [29] M. Samek. *Practical UML Statecharts in C/C++ Event-Driven Programming for Embedded Systems*. Newnes, 2008.
- [30] F. Starke, C. Traulsen, and R. von Hanxleden. Executing Safe State Machines on a reactive processor. Technical Report 0907, Christian-Albrechts-Universität Kiel, Department of Computer Science, Kiel, Germany, Mar. 2009.
- [31] O. Tardieu and S. A. Edwards. Scheduling-independent threads and exceptions in SHIM. In *Proceedings of the Proceedings of the International Conference on Embedded Software (EMSOFT'06)*, Seoul, Korea, Oct. 2006.
- [32] R. von Hanxleden. SyncCharts in C. In *Proceedings of the International Conference on Embedded Software (EMSOFT'09)*, Grenoble, France, Oct. 2009.
- [33] R. von Hanxleden, X. Li, P. Roop, Z. Salcic, and L. H. Yoong. Reactive processing for reactive systems. *ERCIM News*, 67:28–29, Oct. 2006.
- [34] F. Wagner, R. Schmuki, P. Wolstenholme, and T. W. Thomas. *Modeling Software with Finite State Machines: A Practical Approach*. Auerbach Publications, 2006.
- [35] A. Wasowski. On efficient program synthesis from Statecharts. In *Proceedings of the 2003 ACM SIGPLAN Conference on Language, Compilers, and Tools for Embedded Systems (LCTES'03)*, volume 38, issue 7, June 2003. ACM SIGPLAN Notices.
- [36] S. Yuan, S. Andalam, L. H. Yoong, P. S. Roop, and Z. Salcic. STARPro—a new multi-threaded direct execution platform for Esterel. In *Proceedings of Model Driven High-Level Programming of Embedded Systems (SLA++P'08)*, Budapest, Hungary, Apr. 2008.

Appendix A

The SC files

The complete SC package consists of two files:

sc.h The header file, to be included by each application `<application>.c`

sc.c The file including the main program, to be linked with `<application>.o`

This section also includes the Makefile and a run of calling make.

Listing A.1: The header file `sc.h`

```
1  /*! \ file sc.h
2  *
3  * Definition of SyncChart C macros.
4  *
5  * See README.txt for general information.
6  * See LICENSE.txt for licensing information.
7  * For further information, see
8  * http://www.informatik.uni-kiel.de/rtsys/sc/.
9  *
10 * @author Reinhard v. Hanxleden,
11 * rvh@informatik.uni-kiel.de
12 */
13 #include <stdio.h>
14
15 // =====
16 //! Instruction counting/Tracing
17
18 /*! Check whether externflags has been defined (eg from gcc command line).
19 * If so, suppress tracing and instruction counting.
20 * This then results in compact macro-expanded source code and executable.
21 */
22 #ifndef externflags
23 // Comment the following line out to suppress detailed tracing.
24 #define mytrace
25 #define instrCnt
26 #endif
27
28
29 //! Increment/decrement SC instruction counter.
30
31 /*! Decrement is needed in some places to avoid duplicate counting.
32 */
33 #ifndef instrCnt
34 #define instrCntIncr tickInstrCnt++;
35 #define instrCntIncr tickInstrCnt++;
36 #define instrCntDecr tickInstrCnt--;
37 #else
38 #define instrCntIncr
39 #define instrCntIncr
40 #define instrCntDecr
41 #endif
42
43
44 //! If tracing is turned on, print trace string.
45 #ifndef mytrace
46 #define trace0(f) printf(f);
47 #define trace2(f, a, b) printf(f, a, b);
48 #define trace3(f, a, b, c) printf(f, a, b, c);
49 #define trace4(f, a, b, c, d) printf(f, a, b, c, d);
50 #define trace5(f, a, b, c, d, e) printf(f, a, b, c, d, e);
51 #define trace6(f, a, b, c, d, e, g) printf(f, a, b, c, d, e, g);
52 #define trace7(f, a, b, c, d, e, g, h) printf(f, a, b, c, d, e, g, h);
53 #define trace8(f, a, b, c, d, e, g, h, i) printf(f, a, b, c, d, e, g, h,
54 i);
55 #define trace0c(f) printf(f),
56 #define trace1c(f, a) printf(f, a),
57 #define trace2c(f, a, b) printf(f, a, b),
58 #define trace3c(f, a, b, c) printf(f, a, b, c),
59 #define elsetrace else {
60 #define elsetraceend }
61 #endif
62 #define trace0(f)
63 #define trace1(f, a)
64 #define trace2(f, a, b)
65 #define trace3(f, a, b, c)
66 #define trace4(f, a, b, c, d)
67 #define trace5(f, a, b, c, d, e)
68 #define trace6(f, a, b, c, d, e, g)
69 #define trace7(f, a, b, c, d, e, g, h)
70 #define trace8(f, a, b, c, d, e, g, h, i)
71 #define trace0c(f)
72 #define trace1c(f, a)
73 #define trace2c(f, a, b)
74 #define trace3c(f, a, b, c)
75 #define elsetrace
76 #define elsetraceend
77 #endif
78
79
80 //! Count instruction (optionally), print trace string prefix (optionally).
81
82 /*! Trace string prefix takes a string s (typically denoting the
83 instruction)
84 * and identifies the executing thread, both by name and thread id.
85 */
86 #define traceThread(s) \
87 instrCntIncr \
88 trace3("%-9s_%d/%s_", s, _cid, state[_cid])
89
90 #define traceThreadc(s) \
91 instrCntIncr \
92 trace3c("%-9s_%d/%s_", s, _cid, state[_cid])
93
94 //! Print trace prefix + suffix
95
96 /*! s is string denoting instruction (eg, "PAUSE:")
97 * f is format string for trace suffix
98 * a, b, ... are arguments for format string
99 */
100 #define trace0t(s, f) traceThread(s) trace0(f)
101 #define trace1t(s, f, a) traceThread(s) trace1(f, a)
102 #define trace2t(s, f, a, b) traceThread(s) trace2(f, a, b)
103 #define trace3t(s, f, a, b, c) traceThread(s) trace3(f, a, b, c)
104 #define trace4t(s, f, a, b, c, d) traceThread(s) trace4(f, a, b, c, d)
105 #define trace5t(s, f, a, b, c, d, e) traceThread(s) trace5(f, a, b, c, d, e)
106 #define trace6t(s, f, a, b, c, d, e, f1) traceThread(s) trace6(f, a, b, c,
107 d, e, f1)
108 #define trace7t(s, f, a, b, c, d, e, f1, g) traceThread(s) trace7(f, a, b,
109 c, d, e, f1, g)
110 #define trace8t(s, f, a, b, c, d, e, f1, g, h) traceThread(s) trace7(f, a,
111 b, c, d, e, f1, g, h)
```

```

109 // Variants with trailing comma insted of semicolon:
110 // Variants with trailing comma insted of semicolon:
111 // Variants with trailing comma insted of semicolon:
112 #define trace0tc(s, f)          traceThreadc(s) trace0c(f)
113 #define trace1tc(s, f, a)       traceThreadc(s) trace1c(f, a)
114 #define trace2tc(s, f, a, b)    traceThreadc(s) trace2c(f, a, b)
115 #define trace3tc(s, f, a, b, c) traceThreadc(s) trace3c(f, a, b, c)
116
117 // =====
118 // Type definitions
119 // Type definitions
120 // Type definitions
121 typedef void *labeltype; //!< Computed goto — a la gcc
122 typedef unsigned int bitvector; //!< 32 bits on IA32
123 typedef bitvector signalvector; //!< 32 signals on IA32
124 typedef unsigned short threadtype; //!< Thread id/priority
125 typedef bitvector threadvector; //!< 32 threads on IA32
126
127 // =====
128 // Global variables
129 // Global variables
130 signalvector signals; //!< Bit mask for signals
131 threadvector enabled; //!< Bit mask for enabled threads
132 threadvector active; //!< Bit mask for active threads
133
134 #define _idMax 8*sizeof(threadvector) //!< Number of threads
135
136 int runCnt; //!< Counts program runs
137 int tickCnt; //!< Counts program ticks
138 int tickInstrCnt; //!< Instructions in one tick
139 threadtype _cid; //!< Id of current thread
140 labeltype _pc[_idMax]; //!< Pseudo program counters
141 threadvector _descs[_idMax]; //!< Descendants of thread
142 threadtype _parent[_idMax]; //!< Parent of thread
143 labeltype _returnAddress; //!< For function calls (eg Exit
    Actions)
144
145 #ifndef mytrace
146 char *statePrev[_idMax]; //!< State where thread resumed
    previous tick
147 char *state[_idMax]; //!< State where thread resumed
    current tick
148
149 #define clearPC(id) \
150     statePrev[id] = "<init>"; \
151     state[id] = "<init>"
152
153 #define initPC(p, label) \
154     _pc[p] = &&label; \
155     statePrev[p] = "<init>"; \
156     state[p] = #label
157
158 #define setPC(id, label) \
159     _pc[id] = &&label; \
160     statePrev[id] = state[id]; \
161     state[id] = #label
162
163 #else // No tracing
164 #define clearPC(id)
165 #define initPC(id, label) _pc[id] = &&label
166 #define setPC(id, label) _pc[id] = &&label
167 #endif // mytrace
168
169 // =====
170 // Declarations of constants and variables
171 // Declarations of constants and variables
172 // Declarations of constants and variables
173 // Constants defined in <application>.c
174 int runMax; //!< # of runs to execute
175 int tickMax; //!< # of ticks to execute
176
177 extern const char *s2signame[]; //!< Names of signals
178
179 // =====
180 // Declarations of functions defined in <application>.c:
181 // Declarations of functions defined in <application>.c:
182 // Declarations of functions defined in <application>.c:
183 //! Initialize signals to inputs for one tick.
184 void getInputs();
185
186 //! Set reference outputs and check valued signals, if there are any.
187 int checkOutputs(signalvector *tickOutputs);
188
189 //! Print value of a signal, if it has one.
190 void printVal(int id);
191
192 //! Compute one tick.
193 //! Returns 1 if some thread is still active in current tick.
194 */
195
196 int tick(int islnit);
197
198 //! Functions defined in sc.c
199 void selectCid();
200
201 // =====
202 //! Dispatcher
203 //! Dispatcher
204 //! Dispatcher
205 #if ((defined __i386__ || defined __amd64__ || defined __x86_64__) &&
    defined __XXGNUC__)
206 // Version 1: x86 + gcc available.
207 // Use fast Bit Scan Reverse assembler instruction.
208 #define dispatch() \
209     __asm volatile ("bsrl_%1,%0\n" \
210         : "=r" (_cid) \
211         : "c" (active) \
212         ); \
213     goto *_pc[_cid]
214 #else
215 // Version 2: x86 + gcc not available.
216 // Call function, defined in sc.c.
217 #define dispatch() \
218     selectCid(); \
219     goto *_pc[_cid]
220 #endif
221
222 // =====
223 // Low-level routines
224 // Low-level routines
225 // Low-level routines
226 //! Encoding of signal/thread 'u' (some non-negative int) in bitvector.
227 //! This implementation is fast and simple, BUT limits the max thread
228 //! ID and max signal ID to the word width of the machine (eg 32).
229 */
230 #define u2b(u) (1 << u)
231
232 // =====
233 // Keeping track of the thread status
234 // Keeping track of the thread status
235 // Keeping track of the thread status
236 //! Thread enabling/disabling
237 #define enable(id) \
238     enabled |= u2b(id); \
239     active |= u2b(id)
240
241 #define enableInit(id) \
242     enabled = u2b(id); \
243     active = enabled
244
245 #define disable(id) enabled &= ~u2b(id)
246 #define disableSet(idset) enabled &= ~idset
247 #define isEnabled(id) (enabled & u2b(id))
248 #define isEnabledNotOnly(id) (enabled != u2b(id))
249 #define isEnabledNoneOf(idset) ((enabled & idset) == 0)
250
251 //! Thread (de-)activation
252 #define activate(id) active |= u2b(id)
253 #define deactivate(id) active &= ~u2b(id)
254 #define deactivateSet(idset) active &= ~idset
255 #define isActive(id) (active & u2b(id))
256
257 // =====
258 // Tick start and end
259 // Tick start and end
260 // Tick start and end
261 #define _TickEnd 0 // Priority of TickEnd thread
262
263 //! Start a tick (an instant). 'p' denotes the main thread.
264 //! IF this is the initial tick ('islni' is set),
265 //! THEN initialize things and continue with following instruction,
266 //! ELSE call dispatcher to resume where we left off.
267 *
268 * This also initializes the _TickEnd thread. Note that
269 * _parent[_TickEnd] is undefined. Note also that _descs[_TickEnd]
270 * does not matter, as that thread should never perform an ABORT
271 * (TRANS) or JOIN.
272 */
273 #define TICKSTART(islni, p) \
274     static threadtype _pid, _ppid; \
275     static threadvector _fdescs = 0; \
276     static int _j; \
277     freezePreClear \
278     if ( islni ) { \
279         tickCnt = 0; \
280         initPC(_TickEnd, _L_TICKEND); \
281         enableInit(_TickEnd); \
282         _cid = p; \

```

```

283     _parent[_cid] = _TickEnd;           \
284     clearPC(_cid);                     \
285     enable(_cid);                       \
286     setPreInit                          \
287     setVallnit                          \
288     } else {                            \
289     active = enabled;                   \
290     dispatch_;                          \
291 }
292
293
294 //! Complete a tick.
295 //! Return 0 iff computation has terminated
296 */
297 #define TICKEND                          \
298     _L.TICKEND: setPre                   \
299     return isEnabledNotOnly(_TickEnd);   \
300     mergedDispatch
301
302
303 //! If inlineDispatch is defined, call dispatcher at each operator that
304 //! Otherwise, create shared code block for TERM/PAUSE/dispatch.
305 //! This can be included by TICKEND
306 */
307 #ifdef inlineDispatch
308 #define dispatch_ dispatch()
309 #define mergedDispatch
310 #else // Shared dispatch
311 #define dispatch_ goto _L.dispatch
312 #define mergedDispatch
313     _L.TERM: disable(_cid);
314     _L.PAUSEG: deactivate(_cid);
315     _L.dispatch: dispatch()
316
317 #endif
318
319
320
321 // =====
322 // Pausing, suspending, aborting and terminating a thread}
323
324 //! Construct a label, of the form "_L'line in source file ""
325
326 //! Originally contributed by Nicolas Berthier (nicolas.berthier@imag.fr)
327 *
328 * To avoid label clashes, one must
329 * - not generate multiple labels at the same line (this could be, eg,
330 * "PAUSE; PAUSE" in one line)
331 * - not include another files within a function (this appears
332 * unlikely anyway)
333 *
334 * Note that goto labels have function scope, hence it is ok to have
335 * identical labels in different files, as long as they belong to
336 * different functions.
337 *
338 * Note also that the ## preprocessing macro, which concatenates
339 * strings, prevents macro expansion of its arguments. Thus the
340 * construction using _CONCAT and _CONCAT_helper.
341 */
342 #define _CONCAT_helper(a, b) a ## b
343 #define _CONCAT(a, b) _CONCAT_helper(a, b)
344 #define _LABEL... _CONCAT(_L, _LINE...)
345 #define _LABELL... _CONCAT(_LL, _LINE...)
346
347
348
349 //! Pause a thread, resume at subsequent statement.
350 //! Semantically, this is the primitive operator.
351 * In terms of implementation, it is built with PAUSEG.
352 */
353 #define PAUSE
354 do {
355     trace1t("PAUSE:", "pauses,_active==_0%\n", active)
356     PAUSEG(_LABEL...); _LABEL...: (void) 0;
357 } while (0)
358
359 //! Shorthand for 'PAUSE; GOTO(label)'.
360 //! Pause a thread, resume at 'label'.
361 */
362 #define PAUSEG(label) do {
363     trace1t("PAUSEG:", "pauses,_active==_0%\n", active)
364     PAUSEG_(label);
365 } while (0)
366
367 //! Helper function (if/else-unsafe)
368 #ifdef inlineDispatch
369 #define PAUSEG_(label)
370     setPC(_cid, label);

```

```

454 //! Helper function (if/else-unsafe)
455 #ifndef inlineDispatch
456 #define TERM_ \
457     disable( _cid); \
458     deactivate( _cid); \
459     dispatch_
460 #else
461 #define TERM_ goto _L_TERM
462 #endif
463
464 // =====
465 // Handling concurrency
466
467 //! Spawn a thread at 'label', with priority 'p'.
468 //!
469 //!
470 #define FORK(label, p) do { \
471     FORK_(label, p); \
472     trace3t("FORK:", "forks_%d/%s_active==_0%\n", p, #label, active) \
473 } while (0)
474
475 //! Helper function (if/else-unsafe)
476 #define FORK_(label, p) \
477     initPC(p, label); \
478     _parent[p] = _cid; \
479     _fdescs |= u2b(p); \
480     enable(p)
481
482 //! Denote parent thread, starting at 'label'.
483 //! Must also calculate descendants.
484 //! Descendants are used
485 //! - to check for termination (with JOIN)
486 //! - to be disabled upon abortion (with TRANS)
487 //!
488 #define FORKE(label) do { \
489     trace0t("FORKE:", "\n") \
490     FORKE_(label); \
491 } while (0)
492
493 //! Helper function (if/else-unsafe)
494 #define FORKE_(label) \
495     setPC(_cid, label); \
496     _descs[_cid] = _fdescs; \
497     _fdescs = 0; \
498     _pid = _cid; \
499     while (( _ppid = _parent[_pid] ) != _TickEnd) { \
500         _descs[_ppid] |= _descs[_pid]; \
501         _pid = _ppid; \
502     } \
503     dispatch_
504
505 //! Join completed child threads.
506 //! IF all descendants have terminated,
507 //! THEN proceed after JOINE,
508 //! ELSE pause, resume at 'elselabel'.
509 //!
510 //! Semantically, this is the primitive Join-operator, from which the
511 //! others can be derived; hence JOINE appears first, and the other
512 //! operators are considered shorthands.
513 //!
514 //! In terms of implementation, the operators are built from JOINEG,
515 //! which semantically corresponds to JOINE + GOTO.
516 //!
517 #define JOINE(elselabel) do { \
518     JOINEG_( _LABEL_, elselabel); \
519     _LABEL_: (void) 0; \
520 } while (0)
521
522 //! Shorthand for 'JOINE(elselabel); GOTO(thenlabel)'.
523 //! IF all descendants have terminated,
524 //! THEN jump to 'thenlabel',
525 //! ELSE pause, resume at 'elselabel'.
526 //!
527 #define JOINEG(thenlabel, elselabel) do { \
528     JOINEG_(thenlabel, elselabel); \
529 } while (0)
530
531 //! Helper function (if/else-unsafe)
532 #define JOINEG_(thenlabel, elselabel) \
533     if (isEnabledNoneOf(_descs[_cid])) { \
534         trace1t("JOINEG:", "joins, transfers _to_%s\n", #thenlabel) \
535         goto thenlabel; \
536     }
537
538 trace1t("JOINEG:", "does_not_join,_pauses_at_%s\n", #elselabel) \
539 instrCntDecr \
540 PAUSEG_(elselabel)
541
542 //! Shorthand for 'elselabel: JOINE(elselabel)'. Join completed child
543 //! threads.
544 //! IF all descendants have terminated,
545 //! THEN proceed,
546 //! ELSE pause, resume at JOIN.
547 //!
548 #define JOIN do { \
549     _LABEL_: JOINEG_( _LABELLL_, _LABELL_); \
550     _LABELLL_: (void) 0; \
551 } while (0)
552
553 //! Set priority of a thread.
554 //! Semantically, this is the primitive operator.
555 //! In terms of implementation, PRIOG is the basic operator.
556 //!
557 #define PRIO(p) do { \
558     trace1t("PRIO:", "set_to_priority_%d\n", p) \
559     PRIOG_(p, _LABELL_); \
560     _LABELL_: (void) 0; \
561 } while (0)
562
563 //! Shorthand for 'PRIO(p); GOTO(label)'.
564 #define PRIOG(p, label) do { \
565     trace2t("PRIOG:", "set_to_priority_%d,_goto_%s\n", p, #label) \
566     PRIOG_(p, label); \
567 } while (0)
568
569 //! Helper function (if/else-unsafe)
570 #define PRIO_(p) \
571     deactivate( _cid); \
572     disable( _cid); \
573     _descs[p] = _descs[_cid]; \
574     _pid = _cid; \
575     while (( _ppid = _parent[_pid] ) != _TickEnd) { \
576         _descs[_ppid] &= "u2b(_cid); \
577         _descs[_ppid] |= u2b(p); \
578         _pid = _ppid; \
579     } \
580     _cid = p; \
581     enable( _cid);
582
583 //! Helper function (if/else-unsafe)
584 #define PRIOG_(p, label) \
585     PRIO_(p); \
586     setPC(_cid, label); \
587     dispatch_
588
589 // =====
590 // Efficient shorthands for thread handling
591
592 //! Efficient shorthand for 'PRIO(p); PAUSE; GOTO(label)'.
593 //! Set a priority, then pause (this sets "prionext"), resume at 'label'.
594 //!
595 //! This shorthand avoids the context switch immediately before the PAUSE.
596 //!
597 #define PPAUSEG(p, label) do { \
598     trace2t("PPAUSEG:", "sets_prio_to_%d,_pauses,_resumes_at_%s\n", p, # \
599     label) \
600     PPAUSEG_(p, label); \
601 } while (0)
602
603 //! Helper function (if/else-unsafe)
604 #define PPAUSEG_(p, label) \
605     PRIO_(p); \
606     instrCntDecr \
607     PAUSEG_(label)
608
609 //! Efficient shorthand for 'PRIO(p); PAUSE'.
610 //! Set a priority, then pause (this sets "prionext").
611 //!
612 //! This shorthand avoids the context switch immediately before the PAUSE.
613 //!
614 #define PPAUSE(p) do { \
615     trace1t("PPAUSE:", "sets_prio_to_%d,_pauses\n", p) \
616     PPAUSEG_( _LABELL_); \
617     _LABELL_: (void) 0; \
618 }

```

```

630     } while (0)
631
632
633 //! Efficient shorthand for 'JOIN(label); GOTO thenlabel; label: PRIO(p);
        PAUSE; GOTO(elselabel)'.
634 /*! IF all descendants have terminated,
635 * THEN jump to 'thenlabel',
636 * ELSE set priority, pause, and continue at 'elselabel'.
637 *
638 * This shorthand avoids the context switch immediately before the PAUSE.
639 */
640 #define JPPAUSEG(p, thenlabel, elselabel) do {
641     trace2t("JPPAUSEG:", "%s._prio_=%d\n",
642             isEnabledNoneOf(_descs[_cid]) ? "joins" : "does_not_join", p) \
643     JPPAUSEG_(p, thenlabel, elselabel);
644 } while (0)
645
646
647 //! Helper function (if/else-unsafe)
648 #define JPPAUSEG_(p, thenlabel, elselabel)
649     if (isEnabledNoneOf(_descs[_cid]))
650         goto thenlabel;
651     instrCntDecr
652     PPAUSEG_(p, elselabel)
653
654
655 //! Shorthand for 'JOIN(label); GOTO thenlabel; label: PRIO(p); PAUSE;
        GOTO(elselabel); thenlabel.'.
656 /*! IF all descendants have terminated,
657 * THEN proceed,
658 * ELSE set priority, pause, and continue at 'elselabel'.
659 *
660 * This shorthand avoids the context switch immediately before the PAUSE.
661 */
662 #define JPPAUSE(p, elselabel) do {
663     trace2t("JPPAUSE:", "%s._prio_=%d\n",
664             isEnabledNoneOf(_descs[_cid]) ? "joins" : "does_not_join", p) \
665     JPPAUSEG_(p, _LABEL_., elselabel);
666     _LABEL_.: (void) 0;
667 } while (0)
668
669
670 // =====
671 // Signal initialization, emission and testing
672
673 //! Initialize a local signal (handles reincarnation)
674 #define SIGNAL(s) do {
675     trace2t("SIGNAL:", "initializes _%s/%d\n", s2signame[s], s)
676     signals &= ~u2b(s);
677 } while (0)
678
679
680 //! Emission of a pure signal 's'
681 #define EMIT(s) do {
682     trace2t("EMIT:", "emits _%s/%d\n", s2signame[s], s)
683     signals |= u2b(s);
684 } while (0)
685
686
687 //! Sustain a pure signal 's'
688 #define SUSTAIN(s) do {
689     _LABEL_.: trace2t("SUSTAIN:", "emits _%s/%d\n", s2signame[s], s)
690     EMIT(s); PAUSEG(_LABEL_.);
691 } while (0)
692
693
694 //! Test for presence of signal 's'.
695 /*! IF 's' is present,
696 * THEN return 1,
697 * ELSE return 0.
698 */
699 #define PRESENT(s)
700     (trace3tc("PRESENT:", "determines _%s/%d._as_._%s\n",
701              s2signame[s], s, (signals & u2b(s)) ? "present" : "absent") \
702      (signals & u2b(s)))
703
704
705 //! Test for presence of signal 's'.
706 /*! IF 's' is present,
707 * THEN proceed to next instruction,
708 * ELSE jump to 'label'
709 */
710 #define PRESENTELSE(s, label) do {
711     if (!(signals & u2b(s))) {
712         trace3t("PRESENTELSE:", "determines _%s/%d._as_._absent_._transfers_to_..
713             %s\n", \
714                s2signame[s], s, #label)
715         goto label;
716     }
717 } while (0)
718
719
720     trace2t("PRESENTELSE:", "determines _%s/%d._as_._present\n", s2signame[
721     s], s) \
722     } while (0)
723
724 //! If signal 's' is present, emit 't'.
725 /*! Shorthand for 'PRESENTELSE(s, label); EMIT(t); label:'
726 */
727 #define PRESENTEMIT(s, t) do {
728     if (signals & u2b(s)) {
729         trace4t("PRESENTEMIT:", "determines _%s/%d._as_._present_._emits_._%s
730             /%d\n", \
731                s2signame[s], s, s2signame[t], t)
732         signals |= u2b(t);
733     }
734     elsetrace
735         trace2t("PRESENTEMIT:", "determines _%s/%d._as_._absent\n",
736                s2signame[s], s) \
737     elsetraceend
738     } while (0)
739
740 //! Await (immediately) signal 's'.
741 /*! IF 's' is present,
742 * THEN proceed to next instruction,
743 * ELSE pause.
744 */
745 #define AWAIT(s) do {
746     _LABEL_.: if (!(signals & u2b(s))) {
747         trace2t("AWAIT:", "determines _%s/%d._as_._absent_._waits\n",
748                s2signame[s], s)
749         PAUSEG(_LABEL_.);
750     }
751     trace2t("AWAIT:", "determines _%s/%d._as_._present_._proceeds\n",
752            s2signame[s], s) \
753     } while (0)
754
755 //! Await (non-immediately) signal 's'.
756 /*! Pause; Then, IF 's' is present,
757 * THEN proceed to next instruction,
758 * ELSE pause.
759 */
760 #define AWAIT_(s) do {
761     _LABEL_.: if (!(signals & u2b(s))) {
762         trace2t("AWAIT_:", "determines _%s/%d._as_._absent_._waits\n",
763                s2signame[s], s)
764         _LABELL_.: PAUSEG(_LABELL_.);
765     }
766     trace2t("AWAIT_:", "determines _%s/%d._as_._present_._proceeds\n",
767            s2signame[s], s) \
768     } while (0)
769
770 // =====
771 // Handling valued signals.
772 // The following is compiled conditionally depending on valSigIntCnt.
773 // <application>.c must define valSigIntCnt if valued signals are used.
774
775 #ifndef valSigIntCnt
776 //! At beginning of initial tick:
777 //! Initialize valued signals (-1 is for "undefined").
778 #define setValInit
779     for (i = 0; i < valSigIntCnt; i++)
780         valSigInt[i] = -1;
781
782 #else // #ifdef valSigIntCnt
783 #define setValInit
784 #endif
785
786 //! Emission of a valued signal 's', type integer.
787 #define EMITINT(s, val) do {
788     valSigInt[s] = val;
789     trace3t("EMITINT:", "emits _%s/%d._value_._%d\n",
790            s2signame[s], s, val)
791     signals |= u2b(s);
792 } while (0)
793
794 //! Emission of a valued signal 's', type integer, combined with *.
795 #define EMITINTMUL(s, val) do {

```

```

799     valSigIntMult[s] *= val; \
800     trace4t("EMITInt*:", "emits_%s/%d,_value_%d,_result_%d\n", \
801           s2signame[s], s, val, valSigIntMult[s]) \
802     signals |= u2b(s); \
803   } while (0)
804
805
806   Retrieve value of signal 's'.
807   #define VAL(s) ( \
808     trace3tc("VAL:", "determines_value_of_%s/%d_as_%d\n", \
809             s2signame[s], s, valSigInt[s]) \
810     valSigInt[s])
811
812   Retrieve value of signal 's' into 'reg'.
813   #define VALREG(s, reg) do { \
814     trace3t("VALREG:", "determines_value_of_%s/%d_as_%d\n", \
815            s2signame[s], s, valSigInt[s]) \
816     reg = valSigInt[s]; \
817   } while (0)
818
819
820 // =====
821 // Handling PRE
822 // The following is compiled conditionally depending on usePRE
823 // <application>.c must define usePRE if PRE is used
824
825
826 #ifdef usePRE
827 signalvector sigsPre; Signals from previous tick
828 signalvector sigsFreeze; Signals that are frozen, due to suspension
829
830 At beginning of initial tick:
831 Initialize previous signals.
832 #define setPreInit \
833     sigsPre = 0; \
834     setPreValInit;
835
836 At end of tick:
837 Copy current signals (unless frozen) to previous signals.
838 #define setPre \
839     sigsPre = (sigsPre & sigsFreeze) | (signals & ~sigsFreeze); \
840     setPreVal
841
842 When suspending current thread:
843 Add signals local to current thread or its descendants
844 to list of signals to freeze.
845 #define freezePre sigsFreeze |= sigsDescs[_cid];
846
847 At beginning of tick:
848 Clear list of signals to freeze.
849 #define freezePreClear sigsFreeze = 0;
850
851 #else ifndef usePRE
852 #define setPreInit
853 #define setPre
854 #define freezePre
855 #define freezePreClear
856 #endif ifndef usePRE
857
858
859 Test for presence of signal in previous tick.
860 IF 's' was present in previous tick,
861 THEN return 1,
862 ELSE return 0.
863 */
864 #define PRESENTPRE(s) \
865     (trace3tc("PRESENTPRE:", "determines_%s/%d_as_%s\n", \
866             s2signame[s], s, (sigsPre & u2b(s)) ? "present" : "absent") \
867     (sigsPre & u2b(s)))
868
869
870 Test for presence of signal in previous tick.
871 IF 's' was present in previous tick,
872 THEN proceed to next instruction,
873 ELSE jump to 'label'
874 */
875 #define PRESENTPREELSE(s, label) do { \
876     if (!(sigsPre & u2b(s))) { \
877         trace3t("PRESENTPRE:", "determines_previous_%s/%d_as_absent,_ \
878               transfers_to_%s\n", \
879               s2signame[s], s, #label) \
880         goto label; \
881     } \
882     trace2t("PRESENTPRE:", "determines_previous_%s/%d_as_present\n", \
883            s2signame[s], s) \
884   } while (0)
885
886 // =====

```

```

976 * Return to 'retlabel'
977 * Use this if an Exit Action _may_ have to be performed
978 * Shorthand for ISAT(id, statelabel, retlabel); CALL(label, retlabel);
979 */
980 #define ISATCALL(id, statelabel, label) do { \
981     if (isEnabled(id) && (_pc[id] == &&statelabel)) { \
982         trace1t("ISATCALL:", "calls_%s\n", #label) \
983         _returnAddress = &&_LABEL_; \
984         goto label; \
985     } \
986     trace1t("ISATCALL:", "does_not_call_%s\n", #label) \
987     _LABEL_: (void) 0; \
988 } while (0)

```

Listing A.2: The main program file sc.c

```

1  /*! \ file sc.c
2
3  * Main file for using SyncChart C macros
4
5  * See README.txt for general information.
6  * See LICENSE.txt for licensing information.
7  * For further information, see
8  * http://www.informatik.uni-kiel.de/rtsys/sc/.
9
10 * @author Reinhard v. Hanxleden,
11 * rvh@informatik.uni-kiel.de
12 */
13
14 #include "sc.h"
15
16 // =====
17 // Computing the id of next thread to be dispatched.
18 //! Uses obvious algorithm, run time linear in position of highest bit.
19 * Note that there are also alternatives that run logarithmic to bit vector
20 size.
21 * See eg http://graphics.stanford.edu/~seander/bithacks.html#IntegerLog.
22 * Which is actually faster depends on application.
23 */
24 void selectCid () {
25     int act;
26
27     _cid = 0;
28     for (act = active; act > 1; act >>= 1)
29         _cid++;
30 }
31
32 // =====
33 //! Tracing routines
34 void vec2names(char *prefix, char* suffix, bitvector ids, const char *
35 names[])
36 {
37 #ifdef mytrace
38     int id = 0;
39     int first = 1;
40
41     printf ("%s", prefix);
42     while (ids) {
43         if (ids & 1) {
44             if (first) {
45                 first = 0;
46             } else {
47                 printf (",");
48             }
49             printf ("_%d/", id);
50             if (names) {
51                 printf ("%s", names[id]);
52                 printfVal (id);
53             } else {
54                 printf ("%s", state[id]);
55             }
56             ids >>= 1;
57             id++;
58         }
59     }
60
61     if (first) {
62         printf ("<init>");
63     }
64     printf ("%s", suffix);
65 #endif
66 }
67
68 // =====

```

```

71 //! The main program
72 /*! Returns 0 iff outputs generated by program match reference trace
73 */
74 int main()
75 {
76     int runInstrCnt; // Instructions in one run
77     int runsInstrCnt = 0; // Instructions accumulated over all runs
78     int outputsOK = 1; // Outputs of simulation correct?
79     int notDone; // Current run not done yet?
80     int init; // Is initial tick?
81     signalvector tickInputs; // Input values for a tick
82     signalvector tickOutputs; // Reference output values for a tick
83     signalvector tickSignals; // Reference signal values for a tick
84
85 // Execute all runs
86 for (runCnt = 0; (runCnt < runMax) && outputsOK; runCnt++) {
87     printf ("####_RUN_%d_STARTS_#####\n",
88 runCnt);
89
90     runInstrCnt = 0;
91     tickCnt = 0;
92     init = 1;
93     enabled = 0;
94
95     do { // Execute all ticks of one run
96         tickInstrCnt = 0;
97         getInputs ();
98         tickInputs = signals;
99
100        trace3("====_TICK_%d_STARTS,_inputs_=%0o,_enabled_=%0o\n",
101            tickCnt, tickInputs, enabled);
102        vec2names("====_Inputs_(id/name):_", "\n", tickInputs, s2signame);
103        vec2names("====_Enabled_(id/state):_", "\n", enabled, 0);
104
105        notDone = tick(init); // Call automaton function
106        init = 0;
107
108        runInstrCnt += tickInstrCnt;
109        trace2("====_TICK_%d_terminates_after_%d_instructions.\n",
110            tickCnt, tickInstrCnt);
111        vec2names("====_Enabled_(id/state):_", "\n", enabled, 0);
112        vec2names("====_Resulting_signals_(name/id):_", "", signals,
113            s2signame);
114        outputsOK = checkOutputs(&tickOutputs);
115        if (outputsOK) {
116            tickSignals = tickInputs | tickOutputs;
117            if (signals == tickSignals) {
118                trace0(" ,_Outputs_OK.\n\n");
119            } else {
120                vec2names(" ,_Outputs_NOT_OK_-_expected_signals_", "!!\n\n",
121                    tickSignals, s2signame);
122                outputsOK = 0;
123            }
124        } else {
125            notDone = 0;
126        }
127
128        tickCnt++;
129        if (tickCnt >= tickMax) {
130            printf ("====_Executed_tickMax_=%d_ticks,_terminate.\n",
131                tickMax);
132            notDone = 0;
133        }
134    } while (notDone && outputsOK);
135
136    printf ("####_RUN_%d_terminates_after_%d_instructions\n",
137        runCnt, runInstrCnt);
138    runsInstrCnt += runInstrCnt;
139 };
140
141 printf ("####_All_runs_terminate,_after_%d_instructions\n",
142     runsInstrCnt);
143 return !outputsOK;
144 }

```

Listing A.3: The Makefile

```

1  sc-progs := ABRO Count2Suspend FilteredSR grcbal3 \
2  PreAndSuspend PrimeFactor Reincarnation Shifter3 SurfDepth
3
4  progs := $(sc-progs) $(sc-progs==sc) \
5  ABRO-expanded ABRO-expanded-annotated ABRO-expanded-
6  stripped \
7  Exits IfElseSafe PCO

```



```

8 | downloads := LICENSE.txt Makefile README.txt doxygen.conf sc.c sc.h
9 |   selectCid.c \
10 |     $(progs:=.c) $(progs:=.out)
11 | # CCFLAGS := -Wall -D inlineDispatch
12 | CCFLAGS := -O3
13 | # -Wall complains about unused label _L_TERM
14 | # CCFLAGS := -Wall -O3
15 |
16 | all : $(progs:=.out)
17 |
18 | make.trace:
19 |     (time make) >& $@
20 |
21 | allprogs : $(progs)
22 |
23 | PCO: PCO.c sc.h Makefile
24 |     gcc $(CCFLAGS) PCO.c selectCid.c -o PCO
25 |
26 | PCO2: PCO2.c sc.h Makefile
27 |     gcc $(CCFLAGS) PCO2.c selectCid.c -o PCO2
28 |
29 | # Want to compile ABRO-C without tracing, to not interfere with kbd input
30 | ABRO-C: ABRO-C.c sc.c sc.h Makefile
31 |     gcc $(CCFLAGS) -D externflags -D instrCnt ABRO-C.c sc.c -o $@
32 |
33 | %.out2: %.c sc.c sc.h Makefile
34 |     gcc $(CCFLAGS) $*.c sc.c -o $*
35 |     ./ $* > $@
36 |
37 | %.out: % %.c
38 |     ./ $* > $@
39 |
40 |
41 | # Preprocessing
42 | %-expanded.c: %.c sc.h Makefile
43 |     gcc $(CCFLAGS) -D externflags -E $*.c > $@
44 |
45 | %-expanded-flags.c: %.c sc.h Makefile
46 |     gcc $(CCFLAGS) -E $*.c > $@
47 |
48 |
49 | # Assembler
50 | %s: %.c sc.h Makefile
51 |     gcc $(CCFLAGS) -D externflags -S $*.c
52 |
53 | %-unopt.s: %.c sc.h Makefile
54 |     gcc $(CCFLAGS) -o $@ -D externflags -S $*.c
55 |
56 | %o: %.c sc.h Makefile
57 |     gcc $(CCFLAGS) -D externflags -c -o $*.o $*.c
58 |
59 | %.asm: %.o
60 | #     objdump -d $*.o > $@
61 |     otool -tv $*.o > $@
62 |
63 | %-linked.asm: %.c sc.c sc.h Makefile
64 |     gcc $(CCFLAGS) -D externflags $*.c sc.c -o $*
65 | #     objdump -d $*.exe > $@
66 |     otool -tv $* > $@
67 |
68 |
69 | # Default rule for executable
70 | %: %.c sc.c sc.h Makefile
71 |     gcc $(CCFLAGS) $*.c sc.c -o $*
72 |
73 | # Statistics
74 | all%.wc: $(progs:=%)
75 |     echo $@, CCFLAGS = $(CCFLAGS) > $@
76 |     date >> $@
77 |     wc $^ >> $@
78 |
79 | wc: all .c.wc all -expanded.c.wc all.o.wc
80 |     cat $^
81 |
82 |
83 | # Example to match either or expression :
84 | # grep -E 'trace|\' sc.h
85 | %stats:
86 |     echo "Line_count_of_$*:"
87 |     wc $*
88 |     echo "Comment_line_count_of_$*:"
89 |     grep -c "^_*/" $*
90 |     echo "Empty_line_count_of_$*:"
91 |     grep -c "^\$" $*
92 |     echo "Trace_related_line_count_of_$*,_discouting_multi_line_trace_
93 |         commands_(9),_counting_again_comments_(4):"
94 |     grep -c "trace" $*

```