

Verification and Optimization of a PLC Control Schedule

Ed Brinksma¹, Angelika Mader¹, Ansgar Fehnker² *

¹ Faculty of Computer Science, University of Twente, Netherlands, e-mail: {brinksma,mader}@cs.utwente.nl

² Electrical & Computer Engineering, Carnegie Mellon University, USA, e-mail: ansgar@ece.cmu.edu

Abstract. We report on the use of model checking techniques for both the verification of a process control program and the derivation of optimal control schedules. Most of this work has been carried out as part of a case study for the EU VHS project (Verification of Hybrid Systems), in which the program for a Programmable Logic Controller (PLC) of an experimental chemical plant had to be designed and verified. The original intention of our approach was to see how much could be achieved here using the standard model checking environment of SPIN/Promela. As the symbolic calculations of real-time model checkers can be quite expensive it is interesting to try and exploit the efficiency of established non-real-time model checkers like SPIN in those cases where promising work-arounds seem to exist. In our case we handled the relevant real-time properties of the PLC controller using a time-abstraction technique; for the scheduling we implemented in Promela a so-called *variable time advance procedure*. To compare and interpret the results we carried out the same case study with the aid of the real-time model checker UPPAAL, enhanced with facilities for cost-guided state space exploration. Both approaches proved sufficiently powerful to verify the design of the controller and/or derive (time-)optimal schedules within reasonable time and space requirements.

1 Introduction

Nowadays, the verification of hybrid systems is a popular topic in the formal methods community. The pres-

ence of both discrete and continuous phenomena in such systems poses an inspiring challenge for our specification and modelling techniques, as well as for our analytic capacities. This has led to the development of new, expressive models, such as timed and hybrid automata [3, 16], and new verification methods, most notably model checking techniques involving a symbolic treatment of real-time (and hybrid) aspects [11, 17, 7].

An important example of hybrid (embedded) systems are process control programs, which involve the digital control of processing plants, e.g. chemical plants. A class of process controllers that are of considerable practical importance are those that are implemented using Programmable Logic Controllers or PLCs. Unfortunately, both PLCs and their associated programming languages have no well-defined formal models or semantics, which complicates the design of reliable controllers and their analysis.

To assess the capacity of state-of-the-art formal methods and tools for the analysis of hybrid systems, the EU research project VHS (Verification of Hybrid Systems) has defined a number of case studies. One of these studies concerns the design and verification of a PLC program for an experimental chemical plant.

In this article we report on the use of two model checkers for the verification of a process control program for the given plant and the derivation of optimal control schedules. It is a companion paper to [13], which concentrates on the correct design of the process controller. The original intention of our approach was to see how much could be achieved using the standard model checking environment of SPIN/Promela [8]. As the symbolic calculations of real-time model checkers can be quite expensive it is interesting to try and exploit the efficiency of established non-real-time model checkers like SPIN in those cases where promising work-arounds seem to exist. In our case we handled the relevant real-time properties of the PLC controller using a time-abstraction technique;

* The work reported here was carried out while the second and third authors were employed by the Computer Science Department of the University of Nijmegen, Netherlands. The second author was supported by an NWO postdoc grant, the third author by an NWO PhD grant, and both were supported by the EU LTR project VHS (project nr. 26270).

for the scheduling we implemented in Promela a so-called *variable time advance procedure* [15]. For this case study these techniques proved sufficient to verify the design of the controller and derive (time-)optimal schedules with very reasonable time and space requirements. A first report on our experiences has been published as [5], whose findings are further extended and elaborated in this article.

One of the conclusions of our initial experiments as reported in the initial publication [5] was that “... it would be useful to be able to influence the search strategy of the model checker more directly and guide the search first into those parts ... where counterexamples are likely to be found.” Since this publication a new version of the real-time model checking tool UPPAAL has become available that employs a cost-guided evaluation strategy for state-space exploration, viz *cost-optimal UPPAAL* [4]. This tool is a natural candidate to support the derivation of optimal control schedules in a real-time environment. This motivated us to carry out the optimisation part of the case study again with cost-optimal UPPAAL, both as an interesting exercise in its own right, and to collect data to interpret and compare with the results obtained with SPIN.

The rest of this paper is organised as follows: section 2 gives a description of the batch plant, the nature of PLCs and a description of the control program that was systematically designed in previous work [13]. Section 3 describes the Promela models for the plant and the control process, and their use for its formal verification and optimization. Section 4 then introduces a cost-optimal UPPAAL model of the same processes, and presents the optimization results that were obtained using it. Section 5, finally, evaluates the work and presents our conclusions.

2 Description of the system

The system of the case study is basically an embedded system, consisting of a batch plant and a Programmable Logic Controller (PLC), both of which are described in more detail below. The original goal of the case study was to write a control program such that the batch plant and the PLC with its control program together behave as intended. The intended behaviour is in the first place that new batches can always be produced, and, in the second place, optimality of the control schedule.

2.1 Description of the batch plant

The batch plant (see Figure 1) of the case study is an experimental chemical process plant, originally designed for student exercises. We describe its main features below; a more detailed account can be found in Kowalewski’s description of the plant [10].

It “produces” batches of diluted salt solution from concentrated salt solution (in container B1) and water (in container B2). These ingredients are mixed in container B3 to obtain the diluted solution, which is subsequently transported to container B4 and then further on to B5. In container B5 an evaporation process is started. The evaporated water goes via a condenser to container B6, where it is cooled and pumped back to B2. The remaining hot, concentrated salt solution in B5 is transported to B7, cooled down and then pumped back to B1.

The controlled batch plant is clearly a hybrid system. The discrete element is provided by the control program and the (abstract) states of the valves, mixer, heater and coolers (open/closed, on/off). Continuous aspects are tank filling levels, temperatures, and time. The latter can be dissected into real-time phenomena of the plant on the one hand, such as tank filling, evaporation, mixing, heating and cooling times, and the program execution and reaction times (PLC scan cycle time), on the other. The controller of the batch plant is a nice example of an *embedded system*: the controlling, digital device is part of a larger physical system with a particular functionality.

For the case study we decided to fix the size of a batch: it consists of either a volume of 4.2 liter salt solution with a concentration of 5 gram/liter and a volume of 2.8 liter pure water distributed over two containers, or, when mixed, a volume of 7 liter salt solution of 3 gram/liter concentration in a single container. With these batch sizes containers B1, B2, B4, B6 and B7 have a capacity of two “units” of the relevant volumes, and B3 and B5 of only one such “unit”. The plant description [10] gives durations for the transport steps from one tank to another. In our (timed) plant model we used these durations as our basis, although the actual durations might possibly be different.¹

2.2 Programmable Logic Controllers

PLCs are special purpose computers designed for control tasks. Their area of application is enormous. Here, we briefly emphasize the main characteristics of PLCs in comparison to “usual” computers.

The most significant difference is that a program on a PLC runs in a permanent loop, the so called *scan cycle* (see Figure 2). In each scan cycle the program in the PLC is executed once, where the program execution may depend on variable values stored in the memory. The length of a scan cycle is in the range of milliseconds, depending on the length of the program. Furthermore, a part of each scan cycle is dedicated to data exchange with the environment: a PLC has *input points* connected

¹ Note that these durations, given in Table 1, differ from those used in the initial report [5], which were chosen to be more comparable to the figures used in similar work by Niebert and Yovine [14].

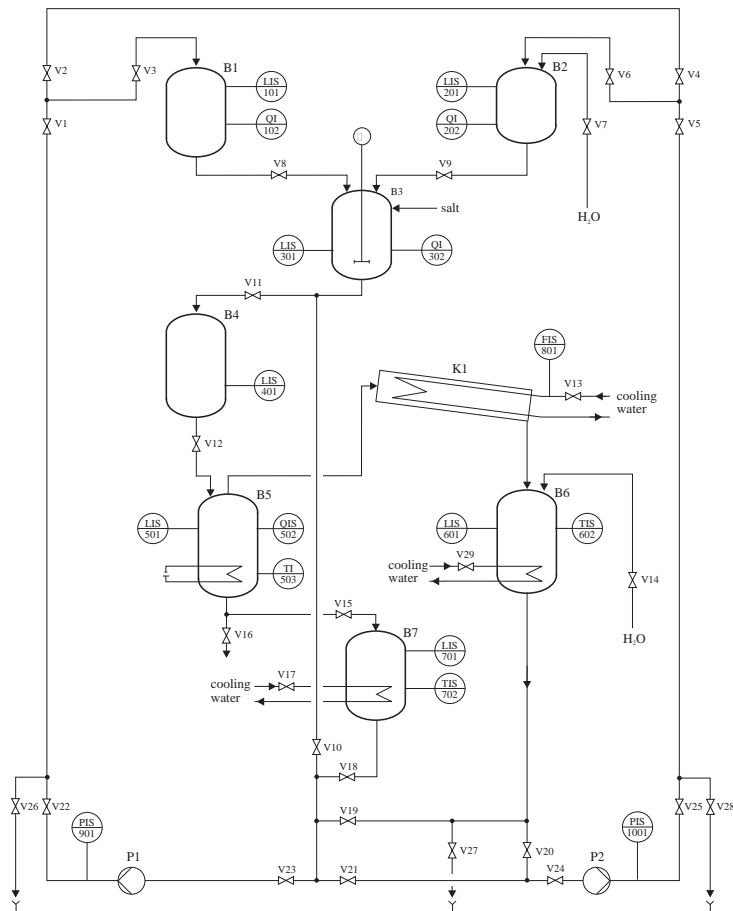


Fig. 1. The P/I-Diagram of the Batch Plant

B1–B3	B2–B3	B3–B4	B4–B5	heat B5	B5–B7	cool B6	cool B7	B6–B2	B7–B1
320	240	600	330	1470	260	300	600	240	220

Table 1. Duration of plant processes in seconds




Fig. 2. A PLC scan cycle

via an interface with a dedicated *input area* of its memory, and the *output area* of the memory is connected via an interface with the *output points* of the PLC. On the input points the PLC receives data from sensors, on the output points the PLC sends data to actuators. Finally, there are some activities of the operating system (self checks, watchdogs etc.) that take place in a scan cycle. The operation system itself is small and stable, which is

prerequisite for reliable real-time control. PLC programs are developed and compiled on PCs in special programming environments and can be loaded to the PLC.

There are different domain specific programming languages collected in an IEC standard for PLCs [9]. In our application we used Sequential Function Charts (SFC), a graphical language that is related to Petri nets, and the program executed in each scan cycle depends on the places that are active at that moment. SFC thus provides for the higher-level structure and the actual instructions of our application are written in Instruction List (IL), an assembly-like language. The PLC languages offer *timers* as basic constructs, which also differs from most of the usual programming languages.

The scan cycle mechanism makes PLCs suitable for control of continuous processes (tight loop control). However, it has to be guaranteed that the scan cycle length is always below the minimal reaction time that is required by the plant to control the entire system. In this case study the scan cycle time is a few orders of mag-

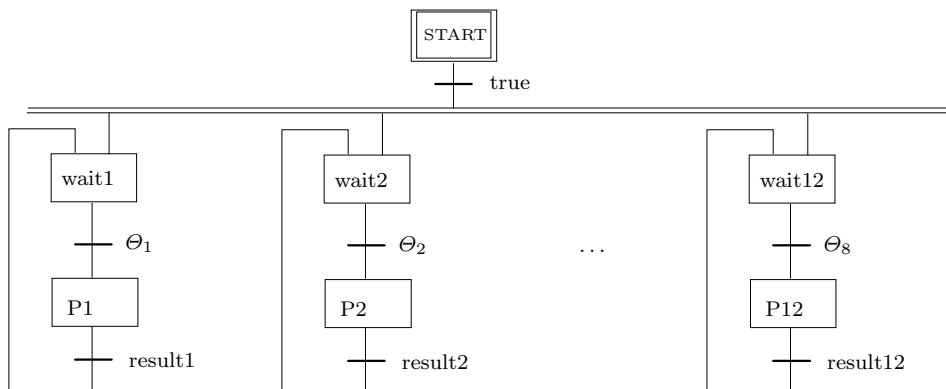


Fig. 3. The control program in Sequential Function Chart

nitude smaller than what the reaction time has to be. The execution time of a scan cycle is in the range of a few milliseconds. For some applications the timing behaviour in this range is relevant, e.g. for machine control. For our chemical plant it is not relevant: it does not really matter whether a valve closes 10 ms earlier or later. This property is relevant when modelling the whole system. Here, we can model the PLC as if executing “time-continuously”, i.e. a scan cycle takes place in zero time. In comparison to the PLC the plant is so “slow” that it cannot distinguish a real PLC with scan cycles from an ideal time-continuous control. For a more detailed discussion of modelling PLCs see [12].

2.3 The control program

This section gives an informal description of the control program as we used it in our verification and optimisation exercises. Its formal derivation and a description of our related other verification activities can be found in a previous publication [13].

In the plant we can identify a number of basic plant processes, such as e.g. the transport of 4.2 l salt solution from container B1 to B3. All possible transport processes, the evaporation process and two cooling processes lead to 12 basic processes, which run in parallel. The activities of each process are simply to open some valves, switch on a mixer, pump or heater, and when the process is finished, close and switch off everything again. Each process starts its activities if its *activation conditions* are fulfilled, and otherwise is in a wait state. An active process remains active until its postconditions are fulfilled. Then it goes back into its waiting state. This means that we have a so-called *closed loop control*: the criterion to change state is not that sufficient time has elapsed, but that a particular event occurs.

It is not difficult to recognise this structure in the SFC representation of the program given in Figure 3. Control starts in the state “START” and (because the transition condition is “true”) immediately distributes to the 12 parallel wait states, waiting for the corresponding

activation conditions θ_i to become true. In a wait state a process does nothing, i.e. it executes an empty program. If the activation condition becomes true, control reaches control state P_i . The program associated with P_i is executed in every scan cycle as long as control remains in P_i . The IL programs P_1, \dots, P_{12} are given in Figure 4. The instructions of IL are assembler-like. Here, we mainly load the constants true or false into the accumulator and write the accumulator value to one of the variables, e.g. V_i , representing valve number i . The *action qualifiers* P_1 and P_0 (to the left of each program) indicate when the corresponding program blocks are executed. P_1 -labelled code is only executed in the first scan cycle after the control location has been reached; P_0 indicates that the instructions are only executed in the last scan cycle that the control is at this location, i.e. when the postcondition evaluates to true – one last scan cycle is executed *after* the postcondition has become true.²

The main complexity of the program is hidden in the design of the activation conditions θ_i . We can assume to have predicates $P_i.X$ for each step P_i indicating whether control is at the corresponding step or not (these variables are available in PLC programs). The conditions to start a process (i.e. step) can now be stated informally as follows:

1. The filling levels of the tanks must allow, e.g., for a transport step: the upper tank must contain enough material, the lower tank must contain enough space, etc. These conditions are encoded in the predicates Φ_i of Figure 6.
2. We do not want a tank to be involved in two (or more) processes at a time. E.g., when transferring solution from B4 to B5 there should not be a concurrent transfer from B3 to B4. This requirement can be formulated by conditions on valves: when solution is transferred from B4 to B5 valve V11 must be closed for the duration of the transfer (invariant).

² There exist more action qualifiers, e.g. to express that code must be executed at every scan cycle; they are not needed for the control program code.

P1:	<table border="1"><tr><td>P1</td><td>LD true ST V8</td></tr><tr><td>P0</td><td>LD false ST V8</td></tr></table>	P1	LD true ST V8	P0	LD false ST V8	P2:	<table border="1"><tr><td>P1</td><td>LD true ST V9</td></tr><tr><td>P0</td><td>LD false ST V9</td></tr></table>	P1	LD true ST V9	P0	LD false ST V9
P1	LD true ST V8										
P0	LD false ST V8										
P1	LD true ST V9										
P0	LD false ST V9										
P3:	<table border="1"><tr><td>P1</td><td>LD true ST V8 ST Mixer</td></tr><tr><td>P0</td><td>LD false ST V8 ST Mixer</td></tr></table>	P1	LD true ST V8 ST Mixer	P0	LD false ST V8 ST Mixer	P4:	<table border="1"><tr><td>P1</td><td>LD true ST V9 ST Mixer</td></tr><tr><td>P0</td><td>LD false ST V9 ST Mixer</td></tr></table>	P1	LD true ST V9 ST Mixer	P0	LD false ST V9 ST Mixer
P1	LD true ST V8 ST Mixer										
P0	LD false ST V8 ST Mixer										
P1	LD true ST V9 ST Mixer										
P0	LD false ST V9 ST Mixer										
P5:	<table border="1"><tr><td>P1</td><td>LD true ST V11</td></tr><tr><td>P0</td><td>LD false ST V11</td></tr></table>	P1	LD true ST V11	P0	LD false ST V11	P6:	<table border="1"><tr><td>P1</td><td>LD true ST V12</td></tr><tr><td>P0</td><td>LD false ST V12</td></tr></table>	P1	LD true ST V12	P0	LD false ST V12
P1	LD true ST V11										
P0	LD false ST V11										
P1	LD true ST V12										
P0	LD false ST V12										
P7:	<table border="1"><tr><td>P1</td><td>LD true ST Heater</td></tr><tr><td>P0</td><td>LD false ST Heater</td></tr></table>	P1	LD true ST Heater	P0	LD false ST Heater	P8:	<table border="1"><tr><td>P1</td><td>LD true ST V15</td></tr><tr><td>P0</td><td>LD false ST V15</td></tr></table>	P1	LD true ST V15	P0	LD false ST V15
P1	LD true ST Heater										
P0	LD false ST Heater										
P1	LD true ST V15										
P0	LD false ST V15										
P9:	<table border="1"><tr><td>P1</td><td>LD true ST V17</td></tr><tr><td>P0</td><td>LD false ST V17</td></tr></table>	P1	LD true ST V17	P0	LD false ST V17	P10:	<table border="1"><tr><td>P1</td><td>LD true ST V29</td></tr><tr><td>P0</td><td>LD false ST V29</td></tr></table>	P1	LD true ST V29	P0	LD false ST V29
P1	LD true ST V17										
P0	LD false ST V17										
P1	LD true ST V29										
P0	LD false ST V29										
P11:	<table border="1"><tr><td>P1</td><td>LD true ST V18 ST V23 ST V22 ST V1 ST V3 ST Pump1</td></tr><tr><td>P0</td><td>LD false ST V18 ST V23 ST V22 ST V1 ST V3 ST Pump1</td></tr></table>	P1	LD true ST V18 ST V23 ST V22 ST V1 ST V3 ST Pump1	P0	LD false ST V18 ST V23 ST V22 ST V1 ST V3 ST Pump1	P12:	<table border="1"><tr><td>P1</td><td>LD true ST V20 ST V24 ST V25 ST V5 ST V6 ST Pump2</td></tr><tr><td>P0</td><td>LD false ST V20 ST V24 ST V25 ST V5 ST V6 ST Pump2</td></tr></table>	P1	LD true ST V20 ST V24 ST V25 ST V5 ST V6 ST Pump2	P0	LD false ST V20 ST V24 ST V25 ST V5 ST V6 ST Pump2
P1	LD true ST V18 ST V23 ST V22 ST V1 ST V3 ST Pump1										
P0	LD false ST V18 ST V23 ST V22 ST V1 ST V3 ST Pump1										
P1	LD true ST V20 ST V24 ST V25 ST V5 ST V6 ST Pump2										
P0	LD false ST V20 ST V24 ST V25 ST V5 ST V6 ST Pump2										

Fig. 4. Instruction List Programs for steps P1. . . . , P12

These requirements induce the conflict structure on the processes in Figure 5.

It is required that control is never at two conflicting processes at the same time. This condition is split into two parts: first, control cannot go to a process if a conflicting process is already active. These conditions are encoded in the predicates Ψ_i of Figure 7. Second, when conflicting processes could get control at the same moment only the one having *priority* gets it. These priorities are fixed, and their priority graph in figure 5 is cycle free. They induce the predicates Θ_i in figure 8.




Fig. 5. Priority graph of the processes P1, . . . , P12; an arrow indicates that there is a conflict between processes, where arrows point to the processes with priority

$$\begin{aligned}
\Phi_1 &:= (B1 = \text{sol42C} \vee B1 = \text{sol82C}) \\
&\quad \wedge B3 = \text{empty} \\
\Phi_2 &:= (B2 = \text{water28C} \vee B2 = \text{water56C}) \\
&\quad \wedge B3 = \text{empty} \\
\Phi_3 &:= (B1 = \text{sol42C} \vee B1 = \text{sol82C}) \\
&\quad \wedge B3 = \text{water28C} \\
\Phi_4 &:= (B2 = \text{water28C} \vee B2 = \text{water56C}) \\
&\quad \wedge B3 = \text{sol42C} \\
\Phi_5 &:= B3 = \text{sol70C} \\
&\quad \wedge (B4 = \text{empty} \vee B4 = \text{sol70C}) \\
\Phi_6 &:= (B4 = \text{sol70C} \vee B4 = \text{sol140C}) \\
&\quad \wedge B5 = \text{empty} \\
\Phi_7 &:= B5 = \text{sol70C} \wedge (B6 = \text{empty} \vee \\
&\quad B6 = \text{water28C} \vee B6 = \text{water28H}) \\
\Phi_8 &:= B5 = \text{sol42H} \wedge (B7 = \text{empty} \vee \\
&\quad B7 = \text{sol42C} \vee B7 = \text{sol42H}) \\
\Phi_9 &:= B7 = \text{sol42H} \vee B7 = \text{sol84H} \\
\Phi_{10} &:= B6 = \text{water28H} \vee B6 = \text{water56H} \\
\Phi_{11} &:= (B7 = \text{sol42C} \vee B7 = \text{sol84C}) \\
&\quad \wedge (B1 = \text{empty} \vee B1 = \text{sol42C}) \\
\Phi_{12} &:= (B6 = \text{water28C} \vee B6 = \text{water56C}) \\
&\quad \wedge (B2 = \text{empty} \vee B2 = \text{water28C})
\end{aligned}$$

Fig. 6. The tank filling conditions

The execution mechanism of PLCs guarantees a synchronous execution of the parallel steps, in the sense that within each scan cycle each program that is attached to an active step is executed once. It is this mechanism that gives the conditions Θ_i their intended effect.

3 Verification and optimization with SPIN

To verify the correctness of the PLC program of Figure 3 we made a Promela model of the plant and the control program, and used SPIN to check that all execution sequences of their composition satisfy the property that “always eventually batches are produced”. This implies that under ideal circumstances in which no material is lost through leakage or evaporation, control is such

$$\begin{aligned}
\Psi_1 &:= \Phi_1 \wedge \neg P2.X \wedge \neg P4.X \wedge \neg P5.X \wedge \neg P11.X \\
\Psi_2 &:= \Phi_2 \wedge \neg P1.X \wedge \neg P3.X \wedge \neg P5.X \wedge \neg P12.X \\
\Psi_3 &:= \Phi_3 \wedge \neg P2.X \wedge \neg P4.X \wedge \neg P5.X \wedge \neg P11.X \\
\Psi_4 &:= \Phi_4 \wedge \neg P1.X \wedge \neg P3.X \wedge \neg P5.X \wedge \neg P12.X \\
\Psi_5 &:= \Phi_5 \wedge \neg P1.X \wedge \neg P2.X \wedge \neg P3.X \wedge \neg P4.X \\
&\quad \wedge \neg P6.X \\
\Psi_6 &:= \Phi_6 \wedge \neg P5.X \wedge \neg P7.X \wedge \neg P8.X \\
\Psi_7 &:= \Phi_7 \wedge \neg P6.X \wedge \neg P8.X \wedge \neg P10.X \wedge \neg P12.X \\
\Psi_8 &:= \Phi_8 \wedge \neg P6.X \wedge \neg P7.X \wedge \neg P9.X \wedge \neg P11.X \\
\Psi_9 &:= \Phi_9 \wedge \neg P8.X \wedge \neg P11.X \\
\Psi_{10} &:= \Phi_{10} \wedge \neg P7.X \wedge \neg P12.X \\
\Psi_{11} &:= \Phi_{11} \wedge \neg P1.X \wedge \neg P3.X \wedge \neg P8.X \wedge \neg P9.X \\
\Psi_{12} &:= \Phi_{12} \wedge \neg P2.X \wedge \neg P4.X \wedge \neg P7.X \wedge \neg P10.X
\end{aligned}$$

Fig. 7. Activation conditions testing for active conflicting processes

$$\begin{aligned}
\Theta_7 &:= \Psi_7 \\
\Theta_8 &:= \Psi_8 \wedge \neg \Theta_7 \\
\Theta_6 &:= \Psi_6 \wedge \neg \Theta_7 \wedge \neg \Theta_8 \\
\Theta_5 &:= \Psi_5 \wedge \neg \Theta_6 \\
\Theta_1 &:= \Psi_1 \wedge \neg \Theta_5 \\
\Theta_3 &:= \Psi_3 \wedge \neg \Theta_5 \\
\Theta_2 &:= \Psi_2 \wedge \neg \Theta_1 \wedge \neg \Theta_3 \wedge \neg \Theta_5 \\
\Theta_4 &:= \Psi_4 \wedge \neg \Theta_1 \wedge \neg \Theta_3 \wedge \neg \Theta_5 \\
\Theta_9 &:= \Psi_9 \wedge \neg \Theta_8 \\
\Theta_{10} &:= \Psi_{10} \wedge \neg \Theta_7 \\
\Theta_{11} &:= \Psi_{11} \wedge \neg \Theta_1 \wedge \neg \Theta_3 \wedge \neg \Theta_8 \wedge \neg \Theta_9 \\
\Theta_{12} &:= \Psi_{12} \wedge \neg \Theta_2 \wedge \neg \Theta_4 \wedge \neg \Theta_7 \wedge \neg \Theta_{10}
\end{aligned}$$

Fig. 8. Predicates encoding the priority between conflicting processes; there are (non-circular) evaluation dependencies

that new batches will always be produced. Subsequently, we refined the models to find more sophisticated variations of the control program that are time-optimal, in the sense that the average production time of a batch is minimal. These optimal scheduling sequences were produced as counter-examples to properties stating suboptimal behaviour. The details of these exercises are given below.

3.1 Correctness of the PLC program

Both the plant as described in Section 2.1, and the informal control program description of Section 2.3 can be translated into Promela in a straightforward way, the crucial part of the modelling exercise being the real-time properties of the plant in combination with the PLC execution mechanism given in section 2.2. In this case there are two basic principles that allow us to deal with the entire system by essentially abstracting away from time (see also [12] for a more general account in the context of PLCs):

1. The control program is independent of the time that the production steps take. In the model each of the production steps P1, . . . , P12 may take some unspecified time: when activated (e.g. by opening a valve)

the control enters an ‘undefined’ state that will eventually be left to reach a final state once the corresponding postcondition holds. In this way every consistent real-time behaviour of the plant is subsumed, including the real one. Proving correctness for the general case implies correctness of the real case.

2. The execution speed of the control program is much faster than the tolerance of the plant processes, as was already mentioned above. This has two important implications:

- we can abstract away from the scan cycle time and assume that scan cycles are executed instantaneously.
- we can assume that the plant is scanned continuously so that state changes are detected without (significant) delay.

The model of the combined behaviour of the plant and the control program is obtained by putting the models of the control process and all the plant processes in parallel. Doing this, we must make sure that the continuous execution of the control program does not cause a starvation of the plant processes. This is taken care of by allowing only fair executions of our Promela model: in each execution no active process may be ignored indefinitely. We must be careful, however, not to lose the other important property, viz. that each state change of the plant is detected “immediately”. Our model takes care of this by forcing a control program execution after each potential state change of the plant.

The Promela model of this case study is too big to be part of this paper. The full version can be retrieved from the VHS document repository [2]. Here we present two excerpts, one of the plant model and one of the control program model, to illustrate the main features.

Figure 9 contains the Promela process that models the transfer of solution from container B1 to B3. It combines the behaviours underlying steps P1 and P3. The model consists of a do-loop that continuously tries to start the transfer of a unit of salt solution from B1 to B3. If the right conditions are fulfilled control will enter the body of the loop, and will mark the beginning of the transfer step by instantaneously (using the Promela `atomic` construct) changing the contents of both containers to undefined transitional states. At some later moment it will execute the second part of the body, instantaneously changing the transitional states to the corresponding terminal states, marking the end of the transfer.³ The `cycle` variable is a global flag that forces the execution of a scan cycle after the execution of each atomic step in the plant (flag is raised at the end of each

³ Here, and in other parts of the Promela model, we have initially added `assert` statements that express invariants that should hold at the corresponding control location (`error` being shorthand for `assert(false)`). These are used to check for the correctness of the model in initial simulation and verification trials, and are removed during the actual verification to avoid unnecessary growth of the state space.

```

proctype B1toB3()
{
  do
  :: atomic{ (cycle==0 && B1!=cempty && v8) ->
    if
    :: (B1==sol42C) -> B1=undef1
    :: (B1==sol184C) -> B1=undef2
    :: else -> error
    fi ;
    if
    :: (B3==cempty) -> B3=undef1
    :: (B3==water28C && mix) -> B3=undef2
    :: else -> error
    fi ;
    cycle=1
  } ;
  assert(v8 && (B3!=undef2 || mix)) ;
  atomic{ (cycle==0 && v8) ->
    if
    :: (B1==undef1) -> B1=cempty
    :: (B1==undef2) -> B1=sol42C
    :: else -> error
    fi ;
    if
    :: (B3==undef1) -> B3=sol42C
    :: (B3==undef2 && mix) -> B3=sol70C
    :: else -> error
    fi ;
    cycle=1
  }
}
od
}

```

Fig. 9. The Promela model of transfer between B1 and B3

```

proctype Control()
{
  int i,j ;
  do
  :: atomic{ i=1 ; j=1 ;
    do
    :: (i<15) ->
      if
      :: (theta(i,j) && !px[procnr(i)]) -> PB1(i)
      :: (result(i,j) && px[procnr(i)]) -> PB0(i)
      :: else -> skip
      fi ;
      if
      :: (j==1) -> j=2
      :: (j==2) -> j=1 ; i=i+1
      fi
      :: (i==15) -> goto endcycle
    od ;
    endcycle: cycle=0
  }
}
od
}

```

Fig. 10. The Promela model of the control process

such atomic step). After the execution of a scan cycle (also modelled as an atomic process, see below) the flag is lowered. Each atomic step in the plant is guarded by the test `cycle==0`.

The Promela process that models the control program is listed in Figure 10. This is a straightforward translation of the PLC program of Figure 4. The do loop of `Control` repeatedly executes an atomic scan cycle, in which the processes `P1`, ..., `P12` are scheduled sequentially. To deal with the symmetric subcases of each step (i.e. the disjuncts between parentheses in Figure 6) the range of the main counter is extended and a second loop counter `j` is introduced. Modulo these small adaptations the Θ -predicates of Figure 8 are captured by the

`theta(i,j)`, and the `result(i,j)` correspond to formalisations of the result conditions of the PLC program (the uninstantiated `resulti` labels of Figure 3). `PB1(i)` and `PB2(i)` correspond to the code of the `P1` part, and the `P0` part of the PLC program, respectively. The variables `px[i]` correspond to the $P_i.X$ activity predicates of the program mentioned earlier. Note that at the end of each scan cycle the global flag `cycle` is lowered, as required.

Whereas the assertions in the model served to check on our own understanding of the model, the main correctness requirement that “always eventually a new batch is produced” was verified using the SPIN facilities for model checking LTL formulas. The requirement was formalized as the following LTL property:

$$\square \diamond (B3 == \text{sol70C}) \wedge \square \diamond (B3 == \text{cempty}) \quad (1)$$

expressing that the contents of container `B3` (containing the salt solution that is considered the “production” of the plant) is infinitely often filled with the diluted salt solution, and infinitely often empty⁴. Inspection of the Promela model for processes involving `B3` shows that there are two possible execution sequences:

1. `B3` is empty, then filled with concentrated solution from `B1`, then mixed with water from `B2` giving the desired diluted solution (`sol70C`), and finally being emptied again, or
2. `B3` is empty, then filled with water from `B2`, then mixed with concentrated solution from `B1` giving the desired diluted solution (`sol70C`), and finally being emptied again.

As the two constituent properties of (1) must both be fulfilled, neither of the above scenarios is allowed to get stuck at some state, making the statement equivalent to the desired requirement that new batches of de diluted solution are always eventually produced.

Verification results. It turned out to be feasible to apply the model checker sequentially to the different initialisations of our model with material for 0 up to 8 batches (including the intermediate different possibilities for half batches; 30 runs in total). In order to avoid the explosion of the more than 8100 possible initial configurations that are in principle possible, we considered only configurations filling the plant “from the top”, i.e. filling tanks in the order `B1`, ..., `B7`. The other initializations are reachable from these by normal operation of the plant. As satisfaction of correctness property (1) for our initial configurations implies its satisfaction for all reachable configurations, this is sufficient. Using simulations of our model we convinced ourselves that our model did indeed include the required normal operation steps.

⁴ the constant `cempty` was chosen to be different from the Promela reserved word `empty`.

After initial simulations and model checking runs had been used to remove small mistakes from our model (identifier overloading, initialisation errors), the model was systematically checked for property (1) for the 30 initializations with different batch volumes described above. No errors were reported, except for initializations with batch volumes 0, 0.5, 7.5 and 8, as should be the case, since the plant can only be productive if at least one batch volume can be circulated. The model checking was done using SPIN version 3.3.7 on a SUN Enterprise E3500-server (6 SPARC CPUs with 3.0 GB main memory). The model checking was run in exhaustive state space search mode with fair scheduling. The error states reported unreachable in all runs. The shortest runs were completed in the order of seconds and consumed in the order of 20MB memory; the longest run required in the order of 40 minutes and 100MB.

3.2 Deriving optimal schedules

The control schedule of Figure 3 that we have shown to be correct by the procedure sketched in the previous subsection, follows an essentially crude strategy. After each scan cycle it enables *all* non-conflicting processes in the plant whose preconditions it has evaluated to hold true. It is not a priori clear that this strategy would also lead to a plant operation that is optimal in the sense that the average time to produce a batch is minimal.

To determine optimal schedules for the various batch loads of the plant we have refined the models of the previous section as follows:

1. We added a notion of time to the model. To avoid an unnecessary blow-up of the state space due to irrelevant points in time, i.e. times at which nothing interesting can happen, we have borrowed an idea from discrete event simulation, viz. that of *variable time advance procedures* [15].
2. We refined the plant model using the information from Table 1, such that each process in the plant will take precisely the amount of time specified.
3. We refined the model of the control program such that after each scan cycle any subset of the maximal set of allowed non-conflicting processes determined by the original control program can be selected for execution.

The search for optimal schedules was conducted by finding counterexamples to the claim:

$$\square(\text{batches} < N) \quad (2)$$

where `batches` is a global variable that counts the number of times that a brine solution is transferred from B3 to B4. This property is checked for increasing values of `N` in the context of a given maximal clock value `maxtime`, i.e. no events will be possible when time exceeds the value of `maxtime`. The assumption is that for `maxtime`

large enough such counterexamples will display regular scheduling patterns. Below, we elaborate on each of the above points and the search procedure.

A variable time advance procedure. In real-time discrete event systems events have associated clocks that can be set by the occurrence of other events. An event occurs when its clock expires. Such systems can be simulated by calculating at each event occurrence the point in time at which the *next* event will occur, and then jumping to that point in time. This is known as *variable time advance* [15].

We wish to apply this idea to our model because it will not litter the global state space with states whose time component is uninteresting, in the sense that there is no process in the plant that begins or ends. As we can only calculate when plant processes will end once they have started, we can only use this time advance procedure if we assume that processes will always be started when others end (or at time 0). It is not difficult to see, however, that we will not lose schedules this way that are strictly faster than what we can obtain using this policy.

Claim For each schedule s there exists a schedule s' in which no scheduled event occurs later and processes are exclusively started at the beginning of the schedule or when other processes end.

Proof sketch Assume that s is a schedule. If s has the desired format take $s' = s$, otherwise there exists in s a first scheduling of a process p with a start event b_p and a termination event e_p such that the occurrence of b_p does not coincide with a termination event of another process p' . Let t be the latest point in time of a process termination occurrence before b_p , or, if no such occurrence exists, the start time of the schedule. Any process that conflicts with p must have terminated at or before t , or started at or after $s(e_p)$ to avoid overlapping with p , where $s(e)$ denotes the time at which event e is scheduled according to s . But then we can produce a new schedule ns by putting $ns(b_p) = t$, $ns(e_p) = s(e_p) - (s(b_p) - ns(b_p))$ and $ns(e) = s(e)$ for all other events. If ns has the desired format take $s' = ns$, otherwise repeat the same procedure starting with ns . Because processes do not have arbitrarily small processing times (i.e. they are non-Zeno), the desired schedule s' is obtained in the limit.

The variable time advance procedure is implemented by the Promela process `Advance` given in Figure 11. The basic idea of `Advance` is quite simple: when it becomes active it will calculate the next point in time when a plant process will terminate. To do so it uses the global array `ptime(i)` containing the termination times of the processes `i`, whose values are calculated as part of the Promela processes modelling the plant, and the global time variable `time`, which is controlled


```

proctype Advance()
{ int i ; short minstep ;
  do
    :: atomic{promptcondition} ->
      minstep=maxstep ; i=1 ;
      do
        :: (i<13) ->
          if
            :: (px[i] && ((ptime(i)-time)<minstep)) ->
              minstep=(ptime(i)-time)
            :: else -> skip
          fi ;
          i=i+1
        :: (i==13) -> goto step
      od ;
      step: time=time+minstep
    od
}

```

Fig. 11. The Promela model of the time advance process

by `Advance`. The activation of `Advance` is controlled by the predicate `promptcondition`. This predicate is true if and only if all processes that have been enabled by the control program have indeed become active and none have yet terminated.

The refined plant model The refined model of the plant differs from the original model by adding some simple timing information to the plant processes. The (atomic) start event of each plant process is used to calculate the termination time of that process. The termination event of each plant process is then guarded with the additional condition that the global time `time` must equal the calculated termination time.

The refined control model. To allow the new model of the control program to enable any subset of the permissible plant process events at any time, we split the loop of the original model of Figure 10. The first of the two loops scans only for termination conditions of plant processes and executes the corresponding control fragments `PB0(i)`. Subsequently, any of the conditions Ψ_i of Figure 7 can be set. The best nondeterministically the value `false` before the second loop is entered. This loop scans the remaining valid preconditions of the plant processes of which the corresponding control fragments `PB1(i)` are then executed. All guards in both loops of the new version contain tests to monitor the progress of time and will stop control if `time` exceeds `maxtime`. This will cause the combined plant and control model to terminate.

The refinement as described above causes an exponential blow-up of the branching structure of the control process, as maximal sets of non-conflicting events are replaced by a choice between all their subsets. This leads to a corresponding blow-up of the exploration time for the state space, even if no new states are reached in this way. To control this phenomenon a global system parameter `cuts` is introduced that specifies the maximal number of events per branching point that may be postponed when enabled. Without the use of this parameter the state space exceeded our resources in practically all

load	simtime	batches	states	steps	period
1	3476	10	711	94253	380
2	4968	25	1715	239277	206
3	4732	25	1652	230236	206
4	4774	25	1736	241261	206
5	3947	20	1423	196627	206
6	4120	20	1459	201685	206
7	3320	10	1935	2568780	314

Table 2. SPIN schedule measurements

experiments (> 1GB memory, > 40 minutes response time).

Finding optimal schedules. Looking for optimal schedules we restricted ourselves to the interesting cases involving initial plant loads of 1 through 7 batches. For our initial experiments we fixed `maxtime` to be 5000 time units (50,000 s). For each initial load we needed two or three runs to determine the maximal number of batches for which counterexamples could be produced in a very short time (in the order of seconds system time). It turned out that all counterexamples produced contained schedules that rapidly (i.e. within 700 time units) converged to a repeating pattern with a fixed duration. None of the counterexamples required a `cuts` value greater than 2.

The best measurements in terms of shortest periods detected, are collected in Table 2. The interpretation of the columns is as follows:

- load:** indicates the number of batches with which the plant is initialised,
- simtime:** indicates the the duration (in simulated time units) of the counterexample traces,
- batches:** the number of batches produced in that trace,
- states:** the number of states visited to produce the trace,
- steps:** the number of atomic steps executed,
- period:** period of the periodic behaviour in time units.

A first analysis of Table 2 shows the state space that needs to be searched to produce the counterexamples is very small, and could make one suspicious of the quality of the results that are obtained. Here, one should realise that the state space of the model has been kept small by the consequent use of `atomic` constructs for large Promela code chunks, the use of variable time advance, and the fact that the reachable part of the plant model state space is not so large (e.g. it does not have complicated data structures or queues). As the column with the atomic step counts indicates the small number of states does not make it a trivial model checking problem. This is caused by the substantial nondeterministic branching in the model, where many alternatives must be evaluated before a counterexample is produced.

Six of the measured periods can easily be shown to be optimal! For a plant with a batch load of 1 this can be readily checked by hand by moving a single batch

through the plant and measuring the total duration of the critical branches of the path. In the earlier experiments [5] (which use slightly different plant process durations), we thought that we had made a mistake when we measured the same period for most initialisation loads of the plant. Closer analysis of the schedules, however, revealed that this is the result of the fact that the plant has one process that clearly dominates the time consumption during the production of batches, viz. the heating of container B5 (147 time units). Since filling B5, heating it, and emptying B5 must be part of every production cycle, the average production time of a batch must be greater or equal than $33+147+26=206$ time units. This makes the schedules for loads 2 to 6 optimal as well.

The schedule for the load of 7 batches is almost certainly optimal. It was found by imposing a very tight `maxtime` upperbound for the production of 10 batches, after a more relaxed upperbound had already produced a period of 346 time units (cf. UPPAAL results below). Further tightening of the upperbound lead to loss of counterexamples, which means that with a load of 7 batches, 10 batches cannot be produced within 3320 time units. By itself, this does not prove that there is no shorter period, however, as this might require an initial behaviour that is less productive than that of the schedule that we found.⁵

It must be concluded that the plant can be scheduled in the overall optimal time of 2060 seconds for all loads, except for the extreme loads of 1 and 7. Because of our analysis above, these are not only time optimal schedules, but also resource optimal ones, in the sense that the (expensive) heating and distillation equipment B5 is in continuous use. From the energy perspective, probably the schedule for load 2 is optimal, as this involves the circulation, heating and cooling of the smallest volume.

4 Optimisation with UPPAAL

4.1 The UPPAAL model

The translation from the plant and control models into a UPPAAL model is quite straightforward, and in many points similar to the philosophy of the Promela model. The main differences between the models are the representation of time, naturally, as this is a built-in feature of UPPAAL, and the scheduling of the concurrent processes. In the Promela model the latter is restricted by fairness conditions and the use of the `cycle` flag variable. The UPPAAL model enforces this by another, more explicit restriction that is explained below.

The UPPAAL model is a parallel composition of a plant automaton and a control automaton. The plant

```

process P13{
clock x12;
state S0, S1{x12<=32}, S2, S3;
commit S2, S3;
init S0;
trans S0 -> S1 {guard Pump1==0, V8==1,
                V9==0, V11==0, Mixer==1;
                sync urgon?;
                assign x12:=0; },
S1 -> S2 {guard x12==32, B1==1, B3==0;
          assign B1:=0, B3:=2; },
S1 -> S2 {guard x12==32, B1==2, B3==0;
          assign B1:=1, B3:=2; },
S1 -> S2 {guard x12==32, B1==1, B3==1;
          assign B1:=0, B3:=3; },
S1 -> S2 {guard x12==32, B1==2, B3==1;
          assign B1:=1, B3:=3; },
S2 -> S3 {sync compute!; },
S3 -> S0 {sync compute!; };
}

```

Fig. 12. The UPPAAL model of transfer between B1 and B3

automaton is given as a parallel composition of 12 subautomata, representing the basic plant processes. Each of these plant automata is equipped with a clock that measures the duration of the process after it has started. As an example, the UPPAAL model for the transport of concentrated salt solution from container B1 to container B3 is given in Figure 12. Location S0 represents the passive control state of the process. As soon as the control process has opened and closed the appropriate valves, etc., the process starts. This immediate activation is achieved by a synchronisation on the *urgent* channel *urgon* (an UPPAAL feature: a channel on which enabled synchronisations cannot be delayed). After the time of the duration of the process has passed, the container volumes are updated depending on their previous values. After a process finishes the control program is activated twice (by synchronisation on the channel *compute*), for reasons given below.

As with Promela, the execution of the control automaton is modelled based on the assumption that the execution of the PLC control program is instantaneous. Technically, this can be implemented using the *committed location* feature of UPPAAL. Once reached, committed locations have to be left without time delay or interleaving with other transitions. In our model all locations of the control automaton are modelled as committed locations, apart from the first one that represents the waiting state of the control process. Without restrictions, repeated instantaneous execution of the control program would lead to Zeno-behaviour. In the Promela model such behaviour is prevented by imposing a fairness requirement. These are not available in UPPAAL. In the UPPAAL model execution of control steps is restricted to those points in time when the conditions in the plant change, as these are the only moments when the control

⁵ Our earlier results [5], which used different process durations, showed easily produced optimal schedules for loads 1,2,3,4 and 7, with optimal schedules obtained for loads 5 and 6 only after substantial tightening of upperbounds.

```

A1 -> A2{assign bonus:=bonus+(T11==1?1:0);},
A1 -> A2{guard T11==1; assign T11:=0;},
A2 -> A3{assign bonus:=bonus+(T12==1 ?1:0);},
A2 -> A3{guard T12==1; assign T12:=0;},

```

Fig. 13. Rewarding the selection of enabled processes (selection variable is not set to 0, if it was 1 before)

process can change also its state. More precisely, it is the case that each change in the state of the plant requires two control program executions: one to finish some process (close valves etc.), and one to start up new ones. In general, starting up new processes could be delayed for some time, but this is not needed. As was already argued for the Promela model, no (time-)optimal schedules are lost by assuming that new processes only start when some other process finishes.

The control automaton itself consists of two sequential parts, again similar to the Promela model. In the first part the activation conditions for processes are evaluated, including a nondeterministic choice of a subset of the processes that can be activated. As in the Promela case, this nondeterministic step is prerequisite for finding optimal schedules. In the second part the control sets the actuator variables for valves, pumps, heater and mixer.

4.2 Optimal schedules with Cost-Optimal UPPAAL

Behrmann et al. [4] proposed an extension of UPPAAL that includes concepts from branch and bound algorithms to provide generic support for optimal search strategies in a real-time context. This extension makes it possible to derive optimal solutions for problems that can be modelled in terms of so-called *Uniformly Priced Timed Automata*. In this model a cost function increases with a fixed rate as time elapses or with a specified amount if a transition is taken. This cost is treated as a special clock with extra operations, but such that the currently implemented data structures in UPPAAL suffice.

Among the techniques that were adopted in cost-optimal UPPAAL are heuristic search orders. Heuristic search orders are obtained by assigning priorities to symbolic states. The symbolic state is extended with a *priority* field and the model checker selects the next state with the highest priority. To change priorities one has to decorate the transitions with an assignment to the priority field.

In our cost-optimal UPPAAL model of the case study we have defined the `cost` to be identical to the time elapsed, i.e. the cost rate equals 1, and added the priority variable `heur` that is calculated according to the expression $1 * \text{bonus} + 100 * \text{depth} - \text{cost}$. The variable `bonus` is used to reward selecting larger rather than smaller subsets of enabled process start events, as is shown in Figure 13. This heuristic directs the exploration such that the controller tries first to start all permissible plant

```

process P13{
clock x12;
state S0, S1{x12<=32}, S2, S3;
commit S2, S3;
init S0;
trans S0 -> S1{guard Pump1==0, V8==1,
                V9==0, V11==0, Mixer==1;
                sync urgon?;
                assign x12:=0, depth:=depth+1; },
S1 -> S2{guard x12==32, B1==1, B3==0;
        assign B1:=0, B3:=2; },
S1 -> S2{guard x12==32, B1==2, B3==0;
        assign B1:=1, B3:=2; },
S1 -> S2{guard x12==32, B1==1, B3==1;
        assign B1:=0, B3:=3; },
S1 -> S2{guard x12==32, B1==2, B3==1;
        assign B1:=1, B3:=3; },
S2 -> S3{sync compute!; },
S3 -> S0{sync compute!; },
}

```

Fig. 14. The depth is increased if a process starts; with depth the heuristic gives preference to depth-first search

load	batches	heuristic search		depth first search	
		simtime	states	simtime	states
1	10	3476	14063	4178	17547
2	25	4968	35952	11048	45781
3	25	4732	34584	10992	45385
4	25	4774	35344	11288	46181
5	20	3947	28808	9294	37716
6	20	4120	29568	9294	37716
7	10	3320	14377	4418	17957

Table 3. This table shows the first solution found by cost-optimal UPPAAL, either with heuristic search or depth-first

processes. The bonus is made extra rewarding for the selection of the evaporation process, which should be in (almost) continuous use for an optimal exploitation of the plant resources.

The variable `depth` is used to reward depth-first over breadth-first search by rewarding the starting of new enabled plant processes substantially higher than transitions of processes that are active already. This is done as good scheduling solutions are likely have this property. Transition `S0 -> S1` in figure 14 models the beginning of the process P13. We increase the priority if this transition is taken to reward the start of P13. Transitions that start other processes and transitions of the controller that enable processes to start are rewarded similarly.

Using this heuristic we could produce optimal schedules for the batch plant without resorting to an initial upper bound `maxtime` on time, as was done using SPIN. The heuristic has a similar effect on the reduction of the search space as the `cuts` parameter in the SPIN case, but once defined, needs no further user control for different initialisations. For the experiments we asked UPPAAL to

produce schedules for the same initial load and the same number of batches as SPIN did. Table 3 gives the results of these experiments. For loads 1 to 6 with the heuristic UPPAAL found an initial solution that has the same period and same simtime as the best solution found by SPIN. The solution for load 7 converges to a schedule with period 346. For the same set of problems depth-first search finds worse solutions and explores more states. Other experiments show that depth-first search cannot find a solution at all if we set a reasonable upperbound `maxtime` on the time, as we did for SPIN.

Cost-optimal UPPAAL offers the option to start backtracking as soon as it finds an initial solution. It uses the cost of this solution to prune states that are more expensive. The bound on the cost is lowered each time it finds a better solution. UPPAAL was not able to find better schedules within reasonable CPU-time and memory for initial loads 1 to 6. For load 7 it was able to find a schedule with simtime 3288. Backtracking yielded a schedule that produces the last batch within 314 time units; this is the period SPIN found for its best schedules. To do so UPPAAL explores 74240 states, which takes 16.5 seconds on a Pentium III 500MHz.

5 Conclusion

In this paper we have shown how the Promela/SPIN environment can be used to verify and optimize control schedules for a small-size PLC controlled batch plant. The approach in this paper relies quite heavily on the structured design of an initial control program in our previous work [13], and on the analysis of formal approaches to PLCs [12].

It is interesting to see that we succeeded in dealing with this real-time embedded system using standard Promela and SPIN. For the verification of the initial control program this was due to a property of the plant, viz. that we could assume instantaneous and immediate scanning of all state changes of the plant. This is a consequence of the tolerance of the plant processes for much slower reaction times than those realised by the PLC control. This makes us conclude that this abstraction can be used for checking non-timed correctness criteria in all process control problems that have this property.

The original task we set ourselves was just to check the correctness of the plant control in the sense that the designed program would in principle always be capable of producing more batches for any reasonable initial load. Having achieved that task we wondered how the model might be used to also look at the optimality of the schedules. As we wanted to treat this in terms of small modifications of the model only, we added time in the form of an explicit time advancing process. This is very close in spirit to the real-time Promela/SPIN extension DTSpin [1]. Given the particular properties of the plant, however, viz. that without loss of optimality

plant processes can be assumed to start when others terminate, we could do this by only generating those points in time in which plant events could take place. From the schedules that we obtained we can conclude that in this case study this variable time advance procedure reduced the generated state space by approximately a factor of 20.

Although more experiments are certainly needed, we believe that variable time advance procedures can be useful for this kind of application. One way to think of them is as an explicitly programmed analogue of the notion of time regions as in timed automata [3]. Taking advantage of specific properties of systems such as ours an explicit approach can sometimes yield better results. On the basis of our modified model we could find optimal schedules in an interactive mode: up to 5 runs with different parameter values, each taking between seconds and a few minutes of user time on a SUN Enterprise E3500-server (most in the seconds range). This is certainly due to the particular characteristics of the given plant, with its very critical heating process. Also, we have been lucky in the sense that the optimal schedules often were found in those parts of the search tree that were explored earlier. Given the small numbers of explored states it might seem that the gain by using variable time advance over the explicit addition of a clock variable to the model (as in DTSpin) is immaterial (going from $O(10^3)$ states to $O(10^{4-5})$). One should keep in mind, however, that the number of atomic steps required to produce a counterexample would also grow by a factor 20, taking it to $O(10^{6-7})$, which most likely would interfere with the use of the model-checker in an interactive mode to find (almost) optimal schedules quickly.

Another approach to the optimal scheduling for the VHS case study 1 is reported by Niebert and Yovine [14]. Here the problem is analysed using the tools OpenKronos and SMI. It is difficult to compare the results of this approach directly with ours, as they include also the production of the initial loads into their schedules, which we just assume to be present. The more general findings seem to be consistent with ours, however. OpenKronos could be used successfully to produce optimal schedules for loads of up to 3 batches before falling victim to the state explosion problem. The symbolic model checker SMI produced results for 6 batches and more, with a computation time of approximately 17 minutes per batch.

Since our initial experiments with Promela and SPIN an implementation of cost-optimal UPPAAL has become available. We have used this tool to redo the optimisation part of the case study and compared it to our SPIN results. The results obtained with this variant of UPPAAL did confirm correctness of our results. Compared to SPIN, cost-optimal UPPAAL certainly offers a more convenient interface for handling guided state-space explorations. For SPIN the latter can only be done indirectly by repeated verification runs with different parameter

settings under control of the user. In cost-optimal UPPAAL the user can control everything directly by defining heuristic cost functions, so that fewer multiple runs are required, although some additional experimentation is needed to fine-tune the heuristic used.

On the basis of the limited experience gained by this case study, the situation is best characterised as follows. Given the current situation, cost-optimal UPPAAL is probably the more convenient tool to search for optimal schedules for this kind of application if an UPPAAL model is available or can be developed. In view of the very reasonable performance of SPIN using the techniques described in this paper, however, developing a UPPAAL model may not pay off if a good Promela model is available, as for this case study. In our case such a model was in fact produced because it proved much easier to develop a Promela model for the (time-abstracted) verification of the plant, then to do so with UPPAAL. Given this situation it can be concluded that it would be very interesting to look also into the possible extension of SPIN with cost-guided state space exploration features. This could prove to be a powerful combination for the application of SPIN to optimisation problems. Recent work by Edelkamp et al. [6] combining SPIN with heuristic directed search algorithms also shows the viability of such a combination.

References

1. DTSpin homepage. <http://www.win.tue.nl/~dragan/DTSpin.html>.
2. VHS: sources of case study 1. <http://www.cs.kun.nl/~mader/vhs/cs1.html>.
3. R. Alur and D.L. Dill. A theory of timed automata. *Theoretical Computer Science*, (138):183–335, 1994.
4. G. Behrmann, A. Fehnker, T.S. Hune, K.G. Larsen, P. Petterson, and J. Romijn. Efficient guiding towards cost-optimality in UPPAAL. In *7th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS 2001)*, 2001.
5. E. Brinksma and A. Mader. Verification and optimization of a PLC control schedule. In *Proceedings of SPIN2000*, volume 1885 of *Lecture Notes in Computer Science*. Springer, 2000.
6. S. Edelkamp, A.L. Lafuente, and S. Leue. Directed explicit model checking with HSF-SPIN. In *Model Checking Software*, volume 2057 of *Lecture Notes in Computer Science*. Springer, 2001.
7. T.A. Henzinger, P.-H. Ho, and H. Wong-Toi. Hytech: a model checker for hybrid systems. *Software Tools for Technology Transfer*, (1):110–123, 1997.
8. G.J. Holzmann. The model checker SPIN. *IEEE Transactions on Software Engineering*, 23(5):279–295, May 1997.
9. International Electrotechnical Commission. *IEC International Standard 1131-3, Programmable Controllers, Part 3, Programming Languages*, 1993.
10. S. Kowalewski. Description of case study CS1 “experimental batch plant”. <http://www-verimag.imag.fr/VHS/main.html>, July 1998.
11. K.G. Larsen, P. Petterson, and W. Yi. Uppaal in a nutshell. *Software Tools for Technology Transfer*, (1):134–153, 1997.
12. A. Mader. A classification of PLC models and applications. In R. Boel and G. Stremersch, editors, *Discrete Event Systems - Analysis and Control*. Kluwer Academic Publishers, 2000. Proceedings of WODES2000.
13. A. Mader, E. Brinksma, H. Wupper, and N. Bauer. Design of a PLC control program for a batch plant - VHS case study 1. *European Journal of Control*, 7:416–439, 2001.
14. P. Niebert and S. Yovine. Computing optimal operation schemes for multi batch operation of chemical plants. VHS deliverable, May 1999. <http://www-verimag.imag.fr/VHS/main.html>.
15. G.S. Shedler. *Regenerative Stochastic Simulation*. Academic Press, 1993.
16. F.W. Vaandrager and J.H. van Schuppen. *Hybrid Systems: Computation and Control*, volume 1569 of *Lecture Notes in Computer Science*. Springer-Verlag, 1999.
17. S. Yovine. Kronos: a verification tool for real-time systems. *Software Tools for Technology Transfer*, (1):123–134, 1997.