Form Method Syst Des (2007) 30:199–216
DOI 10.1007/s10703-006-0027-9

From NuSMV to SPIN: Experiences with model checking flight guidance systems

Yunja Choi

Published online: 3 January 2007
© Springer Science + Business Media, LLC 2007

Abstract Model checking has become a promising technique for verifying software and hardware designs; it has been routinely used in hardware verification, and a number of case studies and industrial applications show its effectiveness in software verification as well. Nevertheless, most existing model checkers are specialized for limited aspects of a system, where each of them requires a certain level of expertise to use the tool in the right domain in the right way. Hardly any guideline is available on choosing the right model checker for a particular problem domain, which makes adopting the technique difficult in practice, especially for verifying software with high complexity.

In this work, we investigate the relative pitfalls and benefits of using the explicit model checker SPIN on commercial Flight Guidance Systems (FGSs) at Rockwell-Collins, based on the author's prior experience with the use of the symbolic model checker NuSMV on the same systems. This has been a question from the beginning of the project with Rockwell-Collins. The challenge includes the efficient use of SPIN for the complex synchronous mode logic with a large number of state variables, where SPIN is known to be not particulary efficient. We present the way the SPIN model is optimized to avoid the state space explosion problem and discuss the implication of the result. We hope our experience can be a useful reference for the future use of model checking in a similar domain.

Keywords Model checking · Flight Guidance Systems · SPIN · NuSMV

1 Introduction

Model checking [4, 17] has become a promising technique for automated verification of software and hardware systems. Motivated by a few success stories of applying the technique in practice [2, 9, 10, 13, 16, 23], the Critical Systems Research Group at the University of Minnesota has integrated the symbolic model checker NuSMV [24] into the

Y. Choi
School of Electrical Engineering and Computer Science, Kyungpook National University, Daegu, South Korea
e-mail: yuchoi76@knu.ac.kr

specification-centered system development environment Nimbus [14]. The integrated model checker has been successfully used for checking hundreds of requirements properties of specifications of the commercial Flight Guidance System (FGS), a component of the Flight Control System, at Rockwell-Collins [3, 23]. The success of the project is particulary meaningful to formal methods in practice, since all the requirements engineering activities for the FGS, from writing specifications to conducting formal verifications, are being performed by the practitioners at Rockwell-Collins.

Nevertheless, the choice of the right model checker has been a question since the beginning of the project; comparative studies on various model checkers in the application domain are quite rare, and, thus, the decision had to be based on anecdotal case studies. Despite the success of the project, we have been wondering whether other choices would have been better. The choice between symbolic and explicit model checking, in particular, is quite unclear, with only limited comparative arguments; among others, it has been argued that symbolic model checking performs better for synchronous systems with hardware-like characteristics and explicit model checking is better for asynchronous systems with a number of communicating processes. Many reports have pointed out, however, that a direct comparison of the two techniques is very difficult, if not impossible [1, 7, 18]. Some experiments report that a symbolic model checker performs better even for asynchronous systems [7], contradicting some of the arguments.

Furthermore, a system can have both synchronous and asynchronous aspects, depending on which part we are interested in and how we interpret the system behavior. For example, the FGS system consists of two FGSs, an active FGS and a passive FGS for back-up, running in parallel. The system is an asynchronous system when we consider its two communicating processes running in parallel; it can be interpreted as a synchronous system, with the complex mode logic complex enough to challenge the capability of the model checking technique, when we focus only on the mode logic of one-side FGS.

In this work, we have investigated the use of Spin [19]—a representative explicit model checker—on FGS to get a better understanding of the capabilities of explicit versus symbolic model checking in the same domain. We present the application of the Spin model checker on the specifications of FGSs and compare the result to the one with NuSMV we have reported earlier [3, 23]. This work involves the challenge of translating synchronous dataflow specifications into a modeling language designed for specifying communicating processes based on control-flow as well as converting open systems into closed systems which is required by Spin.

Our direct translation shows a disastrous performance with Spin, quickly blowing up on even a small-size FGS. The problem is mainly due to the large number of global and local variables, including history values, accessed by both the system model and its environment model. After a careful review, we encapsulated as many variables as possible within each process and used the message passing mechanism to allow other processes to access the values of the shared variables. The result is quite encouraging; Spin is able to check a property on an FGS model within a reasonable amount of time. Furthermore, the same approach makes it possible to reason about an important property regarding two communicating FGSs—the same property NuSMV fails to scale up to check.

Our optimization is systematic and performs minimal changes in order to support automation of the verification process. Considering the performance gain from this minimal optimization, we believe there is much more room for further performance improvement. In other words, Spin can be more flexible than NuSMV in handling large scale systems, but can also be more difficult to use by non-experts, as optimization can raise a couple of issues in terms of usability: (1) Property-based optimization requires different models for different

properties, which can be an issue when checking hundreds of requirements properties, and (2) more aggressive optimization may improve performance, but may result in dramatic changes of the original model, which makes it difficult for the modeler to understand. It is our belief that any optimization approach must be systematic so that it can be automated, providing traceability between the original and the optimized model. This is the focus of our future investigation.

The remainder of this paper is organized as follows: Section 2 discusses existing related work focusing on the comparison of symbolic and explicit model checking techniques. Section 3 briefly introduces the previous work related to model checking Flight Guidance Systems. Section 4 gives an overview of our approach for using Spin for checking FGSs, Section 5 describes our translation and optimization approaches for using Spin. We conclude with a discussion about the implication of the result in Section 6.

2 Related work

The pros and cons of symbolic versus explicit model checking have been the subject of a debate without a conclusive result. The difficulty is due to the fact that theoretical performance analysis is not possible for given problems; an optimal variable ordering for symbolic model checking is an NP-complete problem and so is computing the optimal reduction in applications of partial order reduction for explicit model checking [19]. Heuristics are used in both techniques, leaving us little choice but to "try and see". The difficulty of performance comparison and the importance of the empirical study on this issue are well addressed in [1].

A few comparative reports are available on this issue in different domains. In [6], Spin, Smv, and Xmc were compared on a model of the i-protocol from GNU uucp version 1.04, showing that Xmc outperforms the other two model checkers. The result confirms one of the well-known arguments that an explicit model checker performs better than a symbolic model checker on asynchronous protocol verification. Interestingly, the result has been challenged by Gerard Holzmann, the inventor of Spin, and has been reversed by the careful optimization of the Spin model [18]. This occasion clearly shows the difficulty of having a fair comparison between different model checkers as well as the importance of optimization, which requires high level expertise. Their recent publication [5] presents a more comprehensive comparison among various explicit model checkers in the same problem domain.

Eisner and Peled examined another well-known argument, namely that symbolic model checking is better for hardware systems and explicit model checking is better for software systems, in verifying the software of a disk controller using the symbolic model checker Rulebase and the explicit model checker Spin [7]. Their result shows that Rulebase is able to model check a 2-process system with 10^{150} states, while Spin spaces out after checking 10^8 states with 2G of memory. This result is not favorable to the argument that explicit model checking performs better for software verification, especially for communicating processes.

There also exists a performance comparison on analyzing mode confusion on a Flight Guidance System using Murϕ, Smv, and Spin [21]. The translation is described for each model checker using an example, and the strengths and weaknesses of each tool are discussed with respect to its usability. The possibility of a state-space explosion in Spin is also pointed out, especially because *inlining* is used for all procedures in their translation. To our best knowledge, this is the only comparative study in our domain of interest, namely, aircraft control systems. Nevertheless, the model used in the case study is quite small (several thousands of states), and, thus, the performance part is considered relatively insignificant.

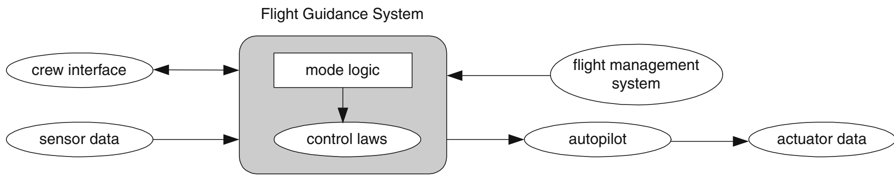

Fig. 1 Flight control system

Since we do not have the analytical data for scalability, we cannot draw a conclusion from this small case study.

3 Background

A case study has been conducted at Rockwell-Collins in cooperation with the University of Minnesota to determine if formal methods could be used to validate system requirements at a reasonable cost. A series of Flight Guidance Systems have been specified using the formal specification language RSML^{-e} (Requirements State Machine Language without events) [28], validated using the visualization and simulation facility provided by the RSML^{-e} execution environment NIMBUS [14], and verified with respect to several hundred functional and safety requirements using the NuSMV model checker and the theorem prover PVS [25] through fully automated translation, finding numerous errors in the model [23]. The project has been quite a success; the use of formal methods, including writing formal specifications and performing verification using NuSMV, was conducted by people at Rockwell-Collins with little help from the researchers at the University of Minnesota. The first phase of the project finished with the conclusion that "formal methods tools are maturing to the point where they can be profitably used on industrial sized problems" [23].

3.1 Flight guidance system

A Flight Guidance System (FGS) is a component of the Flight Control System (Fig. 1, borrowed from [21]). It compares the measured state of an aircraft to the desired state and generates pitch and roll guidance commands to minimize the difference between the measured and desired state. The desired state is determined by the crew interface and/or the flight management system together with the current state of the system. The guidance commands are calculated by control law algorithms selected by the mode logic. The mode logic determines which lateral and vertical modes are armed and active at any given time [22].

The FGS includes identical left and right sides where only one side is active and responsible for inputs and produces outputs. The inactive side simply copies its internal state from the active side, serving as a hot backup (Fig. 2). The complex mode logic of an FGS is a representative of a class of problems frequently encountered in the design of embedded control systems. For a more detailed description of the mode logic of the FGS, please refer to [22].

3.2 FGSs in RSML^{-e}

RSML^{-e} [28] is a synchronous modeling language semantically similar to Lustre [11], SpecTRM-RL [20], and SCR [16]. An RSML^{-e} specification consists of a collection of input variables, state variables, input/output interfaces, functions, macros, and constants;

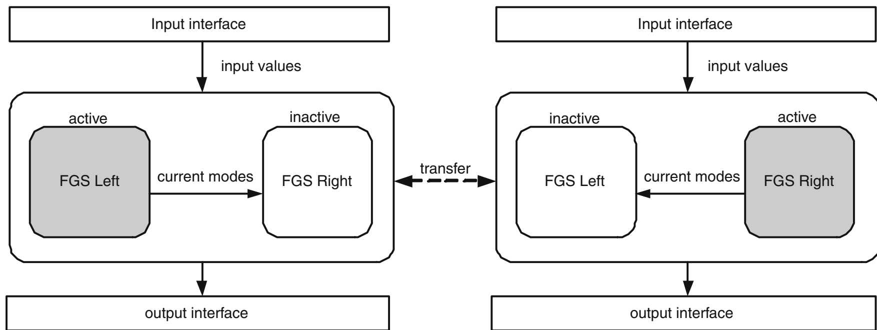

Fig. 2 Flight guidance system with two sides

input variables are used to record the values observed in the environment, *state variables* are organized in a hierarchical fashion and are used to model various states of the control model, *interfaces* act as communication gateways to the external environment, and *functions and macros* encapsulate computations providing increased readability and ease of use.

State variables represent the current state of a system at a given time where the values of state variables are computed based on the input variable values received through input interfaces and the previous values of the state variables (system configuration). The newly computed state variable values can be sent out through output interfaces via output messages. In RSML^{-e}, all variables have a global scope and an RSML^{-e} model is considered open, meaning that the system is interacting with its environment.

The size of the FGS specification written in RSML^{-e} by people at Rockwell-Collins is around 3,500 lines (with comments) for one-sided FGS in the final stage of the project. It includes 13 input switches via crew interface, 82 enumeration variables representing the internal state of the system, and 123 macros. Note that RSML^{-e} provides macro and function constructs to improve the readability of specifications. In the FGS specifications, macros are extensively used as an alternative representation of synchronous events which, in the end, provide a concise way to specify system properties in temporal logic.[1] Since macros encapsulate complex mode logic, the use of macros in temporal logic specification tremendously simplifies logic expression, which plays a crucial role in enhancing the usability of model checking.

3.3 Model checking FGSs using NuSMV

The model checker NuSMV has been used as a primary verification tool in the project with Rockwell-Collins. Due to the space limitation, we discuss only performance and issues. Details can be found elsewhere, including the translation of RSML^{-e} into the NuSMV input language [3] and the process of using NuSMV for verification purposes [23].

3.3.1 Performance

A total of 298 properties are verified in batch mode using the model checker NuSMV (version 2.1) after the RSML^{-e} specification is automatically translated into the NuSMV

[1] The system properties need to be specified in temporal logic [26] in order to be model checked.

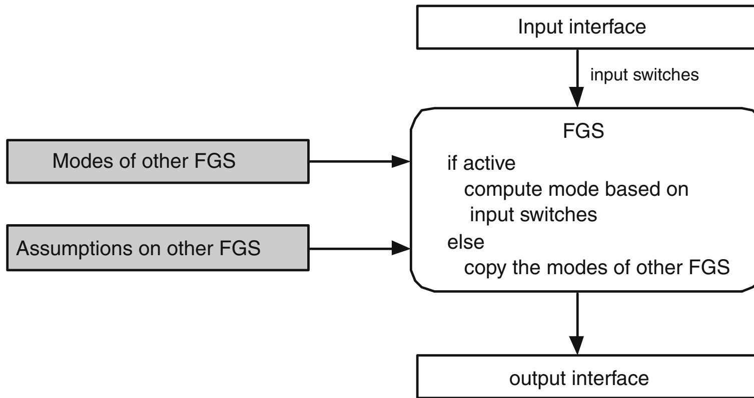

Fig. 3 Flight guidance system with one side

input language. The verification time for each property varies from a couple of minutes to two hours depending on the property and on whether counter examples are generated or not. In the final stage, when all the errors in the model have been corrected,[2] NuSMV verifies all 298 properties in about 1.5 hours on a 800MHz Linux machine with 512M of memory. The translation and the model checking process are completely automated except for the requirements properties being manually translated to temporal logic and attached to the generated NuSMV file.

3.3.2 Issues

Though successful, verification has been focused on one side of the FGS using assumptions about the other side as invariants. These invariants are manually identified by the engineers and imposed on the NuSMV model directly. As shown in Fig. 3, the FGS model contains only one side that is active initially and is activated/inactivated whenever the transfer switch is pressed. When active, it computes the mode of the FGS based on the current mode and the input values. When inactive, it copies the modes of the other side received as random input with some invariants.

This approach is taken because the identified requirements concern mostly the mode logic in the FGS functions that can be checked by assume-guarantee reasoning under the synchrony hypothesis[3] using one-sided FGS. The validation and verification activities have been incremental, from a very abstract FGS to a refined full-size FGS, mainly due to the lack of information on the model checking scalability. NuSMV successfully scales up to the full-size, one-sided FGS, but has not succeeded in scaling up to two-sided communicating FGSs; it appears that the size of the one-sided FGS has almost reached the limit of the NuSMV.

[2] We do not have detailed information on types of errors identified from the 298 properties, since all the verification activities are performed by people at Rockwell-Collins who contacted us only when interpretation of the counter examples was not straightforward. For more details, please refer to [23].

[3] The synchrony hypothesis says that the underlying machine is infinitely fast, and, hence, the reaction of the system to an input event is instantaneous [8].

4 Motivations and challenges for using SPIN

From the experience, it becomes clear that symbolic model checking can be very powerful and also usable, but only up to a certain point. The technique is based on exhaustive state-space search and does not provide alternative options when it reaches its limitation. On the other hand, explicit model checking usually provides more flexibility in dealing with a large state-space, sometimes trading off exhaustiveness for efficiency. Our hope is that the flexibility of explicit model checking may be able to provide us with a certain level of verification capability even for larger systems that symbolic model checking is not capable of handling. We set up two goals of the investigation to realize (or nullify) our hope;

(i) To check the possibility of using the explicit model checker SPIN for verifying FGSs
(ii) To come up with a systematic approach of translating RSML^{-e} into the input language of SPIN to support full automation and better usability, if the first goal turns out to be achievable.

4.1 The SPIN model checker

SPIN is one of the most widely used model checkers for software mainly because it is designed for checking distributed asynchronous systems, a characteristic suitable for a wide range of software systems.

PROMELA, the modeling language of SPIN, is designed to facilitate the construction of models of distributed systems, supporting the specification of non-deterministic control structures, process creation, and inter-process communication. A PROMELA model is constructed from three basic constructs: processes, data objects, and message channels. Processes are instantiations of *proctype*s where each *proctype* contains zero or more data declarations and one or more statements. PROMELA supports primitive data types, such as Boolean, bit, byte, and integer. More complex data types can be defined based on these primitive types. Channels are used to transfer messages between two active processes.

The SPIN verifier is based on explicit model checking that explores system state-space explicitly constructed as a state transition graph, and supports verification of safety properties, liveness properties, deadlock-freeness, and general LTL properties. SPIN provides three major verification options, exhaustive, hash-compact, and bit-state hashing, to be chosen depending on the verification complexity and available resources. Exhaustive verification is sound but requires the most time and memory, bitstate hashing is an option that allows possibly leaving some states unexplored, trading off the soundness of the verification result for scalability, and the hash compact option tries to reduce the possibility of leaving states unexplored to the minimum while reducing the memory requirements [19].

4.2 Challenges

Using SPIN for verifying FGSs specified in RSML^{-e} involves four major challenges;

(i) Translation of synchronous dataflow specifications (RSML^{-e}) to asynchronous control-flow based specifications (PROMELA).
(ii) Translation of open systems into closed systems.
(iii) Imposing invariants

RSML^{-e} computes the next values of the system state variables based on their current values and input variable values, assuming that the computation takes no time (synchrony

hypothesis) and those input and state variable values do not change during the computation. On the other hand, PROMEAL semantics is based on control-flow, and those values can change during the computation. For example, the result of computing $\{x = x + 1; y = x\}$ with the initial value $x = 0$ would be $x = 1,\ \ y = 0$ in RSML^{-e} since both x and y refer to the same value $x = 0$ during the computation. The result would be $x = 1,\ \ y = 1$ in PROMELA, however, since the value of x changes before the value of y is computed. Compared to that of NuSMV, which is also dataflow-based, the translation of RSML^{-e} into PROMELA is much more involved.

Both RSML^{-e} and NuSMV are open systems, assuming that systems are interacting with arbitrary environments, whereas SPIN requires closed systems with an explicitly modeled system environment. This means that we need to explicitly model the behavior of the system environment in terms of possible sequences of values for those 13 input variables of the FGS. Moreover, SPIN does not have a built-in mechanism to support imposing invariants. For example, the invariant "a and b are not true at the same time" can be easily imposed in NuSMV as "INVAR $\neg(a \wedge b)$", which is not supported by SPIN.

Note that NuSMV does not have any of these issues and the translation of RSML^{-e} specifications into the NuSMV input language is relatively straightforward.

5 Model checking FGSs using SPIN

For the investigation, we start from a translation of the RSML^{-e} model of one-sided FGSs with invariants into the input language of SPIN, PROMELA. Our intuitive hypothesis is that SPIN must be able to handle the synchronous one-sided FGSs in order to be scaled to the two-sided asynchronous FGSs. Our first attempt at a direct translation turns out to be too inefficient for model checking. To achieve better performance, we adopt modularization and encapsulation by utilizing the SPIN message passing mechanism and by modifying macros to cope with the structural change. The result of the change is quite promising; it enables SPIN to model check one-sided synchronous FGSs with tolerable cost in terms of time and memory compared to that of NuSMV. The same structuring and modularization approach enables SPIN to scale up to the two-sided asynchronous FGS, which was not feasible with NuSMV; this will be presented in the next section.

5.1 Basic translation

The first challenge of using SPIN for one-sided FGS is to model the synchronous system with a large number of variables and complex mode logic in PROMELA, the modeling language of SPIN, which is designed for modeling asynchronous processes. As noted in the previous section, the FGS specification includes 13 input variables, 82 state variables, and 123 macros. Furthermore, most of the macros refer to the history values of state variables, and, thus, the model needs to remember the values of the variable history. Though the depth of the history is limited to one (to the previous value), it doubles the number of variables required to keep track of the system state. In PROMELA, it requires at least $2 \times (13 + 82)$ variables to keep track of all the necessary values.

For an initial approach, we performed a faithful translation of the RSML^{-e} specifications to PROMELA aiming at automated translation, as we did for the NuSMV translation. Figure 4 illustrates the structural change from RSML^{-e}to PROMELA; we have abstracted out RSML^{-e} input/output messages, which are used to pass input/output values from/to the system environment according to the interface behavior. RSML^{-e} input variables, which are the internal

representation of the input messages, are translated into PROMELA global input variables; an input history variable is also declared for each input variable.

Note that input variables and message channels are declared as global variables in PROMELA,[4] since those are to be accessed by both the FGS model and the environment model. FGS state variables as well as their corresponding history variables are translated into local variables in the FGS process, and their transition relation is translated with the PROMELA *if* statement. User defined enumeration types are translated into the PROMELA *mtype*. Macros are declared as C-style macros in PROMELA; SPIN performs pre-processing for C-style macros, and, thus, does not introduce additional variables for macros.

5.2 From dataflow-based to controlflow-based

The computation logic is converted from dataflow-based to controlflow-based by performing a dependency analysis on state variables and by organizing the control flow based on the variable dependency; all the variables dependent on other variables are computed before the other variables are updated. For example, the specification $\{x = x + 1,\ y = x\}$ in RSML^{-e} is converted into $\{y = x,\ x = x + 1\}$ in PROMELA considering the data dependency.

5.3 Modeling of closed environment

SPIN requires a closed system where all the system interaction with its environment must be explicitly specified. Luckily, the modeling of the pure non-deterministic input is relatively simple for Boolean or enumeration variables, since the PROMELA *if* statement supports non-determinism. For example, for the heading switch HDG, the non-deterministic input value assignment can be done as follows;

```
if
:: 1 -> HDG = On;
:: 1 -> HDG = Off;
fi;
```

As shown in Fig. 4, we explicitly model the system environment by introducing another process. As a result, the translated model has two processes, one for the FGS itself, the other one for the environment. The synchrony of the one-sided FGS is ensured by using the hand-shaking message passing mechanism between the environment process and the actual FGS process.

5.4 Imposing invariants

In order to model the FGS with only one side, we also have to explicitly model the possible inputs from the other FGS with constraints. Unfortunately, PROMELA does not support imposing invariants on the model itself, and, thus, imposing constraints on the non-deterministic input values from the *other side FGS* can be quite tricky. One way is to restrict the non-deterministic value assignment with the constraints and check that the constraints are actually satisfied by the model using the assertion statement or the LTL verifier. For example, the

[4] All the variables in RSML^{-e} have a global scope, whereas PROMELA supports both global and local variable scopes.

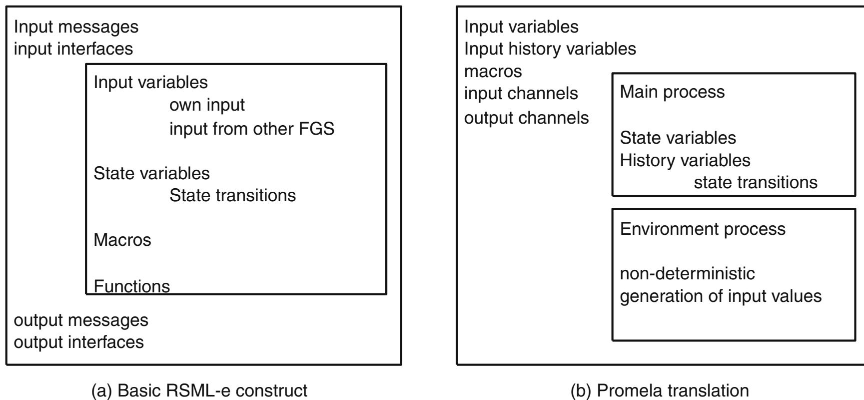

Fig. 4 Basic RSML^{-e} construct and its translation to PROMELA

following is an assumption to be satisfied in the input values from the *other side FGS* and the way it is imposed in NuSMV as an invariant.

- If the mode of the `other side FGS` is `On` then either `ROLL` is selected or `HDG` is selected or `NAV` is selected in the `other side FGS`.
- `INVAR Other_FGS_Mode = On` → `Other_FGS_ROLL = Selected | Other_FGS_HDG = Selected | Other_FGS_NAV = Selected`

In PROMELA, this invariant is manually enforced in the environment model so that the environment does not generate input values that do not obey the constraint. This process can be automated for a given set of invariants. Alternatively, we can use SPIN to check the correctness of invariants after manually imposing them.

Figure 5 shows a small part of translated PROMELA code for a version of the FGS; the upper part of the left side of the figure contains samples of macro declarations, message type declarations for the values of variables, and input variable declarations. The lower part of the left side shows a part of the environment process that generates random input values with constraints. The upper part of the right side of the figure shows how state variables and their history variables are declared in the *FGS* process. Sample mode logic specifications of the *FGS* are illustrated in the lower part.

5.5 Initial performance

The performance of SPIN on the one-sided FGS after applying the basic translation is very poor; SPIN quickly spaces out of memory when checking a property (published in [23]) on a 800 MHz Linux machine with 768 M of memory;

P1. If this side is active and the mode annunciations are off, the mode annunciations shall be turned on when the onside FD is turned on.

With the hash-compact option with compression, the verification process terminates using up 625M of actual memory after 24 minutes. The situation is similar or worse with other SPIN verification options.

```
// macros  (total 123 of macros)
#define When_FD_Switch_Pressed  (FD_Switch == On
&& PREV_STEP_FD_Switch != On)
#define  When_Turn_FD_Off
(When_FD_Switch_Pressed_Seen && Overspeed != 1)
#define When_FD_Switch_Pressed_Seen
(When_FD_Switch_Pressed &&
No_Higher_Event_Than_FD_Switch_Pressed)
          :
// enumeration type declaration
mtype = {Off, On, Cleared, Selected, Armed, Active,
Undefined, done, LEFT, RIGHT, Disengaged, Engaged,
Capture, Track};

// input variables (total 13 of input variables,
// 13 of input history variables)
mtype FD_Switch = Undefined;
mtype HDG_Switch = Undefined;
          :
bool Offside_FD_On;
bool Offside_Modes_On;
bool Offside_Roll_Selected;
          :
proctype env(chan  input1, out){
 do
 :: 1 ->
    // initial non-deterministic value assignments
   if
    :: 1 -> FD_Switch = On;
    ::  1 -> FD_Switch = Off;
   fi;
          :
   // imposing invariants on other side FGS
   if
   :: 1 ->  Offside_Modes_On = 0;
            Offside_Roll_Selected =0;
            Offside_Hdg_Selected=0;
            Offside_Nav_Active =0;
   :: 1 ->  Offside_Modes_On = 1;
            Offside_Roll_Selected =1;
            Offside_Hdg_Selected=0;
            Offside_Nav_Active =0;
   :: 1 ->  Offside_Modes_On = 1;
            Offside_Roll_Selected =0;
            Offside_Hdg_Selected=1;
            Offside_Nav_Active =0;
   :: 1 ->  Offside_Modes_On = 1;
            Offside_Roll_Selected =0;
            Offside_Hdg_Selected=0;
            Offside_Nav_Active =1;
            Offside_Nav_Selected =1;
   fi;
          :
}
```

```
proctype FGS(chan input, out1)
{
// FGS state variables (total 82 of state variables)
mtype Onside_FD = Off;
bool Onside_FD_On = 0;
mtype Modes = Off;
mtype ROLL  = Undefined;
bool  FD_Cues_On = 0;
          :

// history variables (total 82 of history variables)
mtype PREV_STEP_Modes = Undefined;
mtype PREV_STEP_ROLL = Undefined;
          :

// mode logic
          :
  if
   :: ! Is_This_Side_Active ->
               Onside_FD = Offside_FD;
   :: (Onside_FD == Off) &&
      When_Turn_FD_On
                ->      Onside_FD = On;
   :: (Onside_FD == On) &&
      When_Turn_FD_Off
               ->     Onside_FD = Off;
   :: else -> skip;
  fi;

 Onside_FD_On = (Onside_FD == On);

  if
   :: Is_This_Side_Active != 1 ->
               Modes = Offside_Modes;
   :: (Modes ==Off) &&
      When_Turn_Modes_On
      &&   Is_This_Side_Active ->
                        Modes = On;
   :: (Modes == On) &&
      When_Turn_Modes_Off
      &&   Is_This_Side_Active ->
                        Modes = Off;
   :: else -> skip;
  fi;

          :
}
```

Fig. 5 Direct translation of FGS from RSML^{-e} to PROMELA

5.6 Optimization through modularization

The major source of the state-space explosion of the directly translated version of the FGS is the large number of global variables, mainly the input variables and their history variables, in addition to the large number of local state variables in the FGS process. Global variables, especially, are expensive in SPIN verification since SPIN needs to keep track of all their values

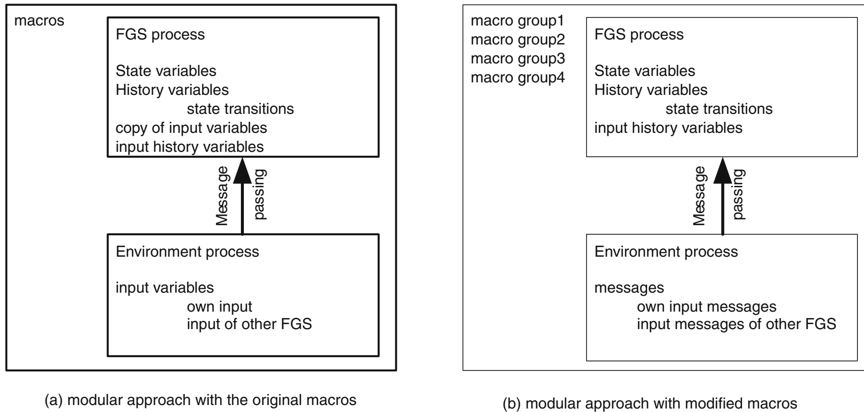

Fig. 6 Modular translation approaches

in the state transition graph. To optimize the performance, we look into mechanical ways to minimize the number of variables. Note that one of our goals in this investigation is to support a usable verification process in practice, and, thus, any modification of the original model for the purpose of optimization needs to be systematic so that it can be automated in future work.

Figure 6 illustrates two possible modularization approaches that convert all the global variables into local variables. Both are based on localizing input variables in the environment process, but differ in the methods of sharing their values with the FGS process; one maintains copies of the input variable values in the FGS process, the other converts input variable values into messages to be passed to the FGS process based on the classification of macros.

5.6.1 Localizing variables

Conceptually, input variables can be considered either a part of the FGS state machine or a part of the environment model that generates the values of input variables. In that sense, we should be able to declare input variables as local variables in the environment process and pass them to the other process through message channels, eliminating all the global variables by modularizing each process. Unfortunately, this is not as straightforward as it sounds, since the input variables and their history values are referred by macros, which are again referred by *FGS* state transitions; *env* has to be able to update the input variables, and the *FGS* process has to be able to access those variable values via macros. For example, `When_Turn_FD_On` macro is referred in the transition condition of the state variable `Onside_FD` in the *FGS* process (Fig. 5). The macro refers to the macro `When_FD_Switch_Pressed_Seen`, which again refers to the macros `When_FD_Switch_Pressed` and `No_Higher_Event_Than_FD_Switch_Pressed`. These final macros refer to input variable values such as `FD_Switch`.

One possible way of dealing with such a usage of macros is to pre-process all the macros by replacing all the references to a macro in the specification with their value expression. The macros, however, are specified by practitioners in a way to match their languages used for communicating with pilots, and many of the properties (98 out of 298) to be verified are specified in terms of macro names. Therefore, we do not want to change them if it is not really necessary.

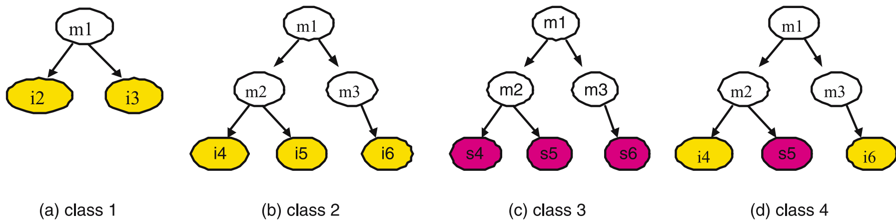

Fig. 7 Four classifications of macros used in FGSs

Alternatively, we can use the PROMELA message passing mechanism in order to make modularization possible with such a usage of macros; all the input variables can be declared as local variables in the *env* process, and all the values can be passed to the *FGS* process through message channels between *env* and *FGS* processes. We can separately declare input history variables as local variables in the *FGS* process to reduce the number of messages to be passed through the message channel. Nevertheless, in order to make use of the macros, we need to preserve the names of the passed message values, which is possible in PROMELA either by creating a local copy of the messages in the *FGS* process, or by creating one message channel per variable (Fig. 6). Both cases are still expensive, since creating a local copy of the messages increases the number of variables, and message channels are always treated as global variables in PROMELA.

5.6.2 Macro classification and optimization strategy

To achieve better optimization, we classify the macros according to the following four categories as shown in Fig. 7: (1) macros whose values are determined only by input variable values with the depth of the reference tree being one, (2) macros whose values are determined only by input variable values with the depth of the reference tree being more than one, (3) macros whose values are determined by state variables locally declared in the *FGS* process, and (4) macros whose values are determined by both input variable and state variable values. Figure 7 (a) shows the first case that the macro m_1 directly refers to the values of input variables i_2 and i_3. Figure 7 (b) shows the second case with multi-level macros that refer to only input variables in the end, whereas Fig. 7 (c) shows the same multi-level macro referring to only state variables in the end. The last reference tree of the figure illustrates the multi-level macro referring to both input and state variables in the end.

Based on this classification, we perform three types of optimization as follows.

Category 1: We change the input variable names referred in the macro to the corresponding message field names passed from the *env* process to the *FGS* process.

Category 2: We create a local variable in the *env* process for each macro in this category so that it can be evaluated in the same process and passed to the *FGS* via a message channel.

Category 3: No change.

Category 4: We modify the names referred in the macros in this category to reflect the change of names of the input variable via a message channel and conversion of macros in category 2.

Through the approach for the macros in category 1, we do not need to keep a local copy for the input variable values, reducing the number of local variables. For example, the macro definition `When_FD_Switch_Pressed` is changed from `(FD_Switch == On) && PREV_STEP_FD_Switch != On)` to `(values.msg[0]== On && PREV_STEP_FD_Switch != On)`, since `FD_Switch` is an input variable declared in the *env* process whose values are passed to the *FGS* process through a message channel (Fig. 8). *values* is the name of the message channel that passes a user-defined message type *msg*, which is an array of *mtype*. The second approach is to treat complex macros that cannot be handled by the simple approach as for those in category 1. For example, the `When_Turn_FD_On` macro described on the previous page has a reference tree of depth 3. We could traverse the names referenced in the reference tree and change them to reflect the changes in the first approach. Instead, we convert them into local variables in the *env* process for simplicity. In this way, all the references are resolved in the *env* module and only the resulting value will be passed to the *FGS* process (see inside *proctype env* in Fig. 8). Macros in category 3 do not need to be changed, since all the macros are referenced in the *FGS* process where all the state variables are declared as local variables. The approach for macros in category 4 combines the approaches taken for the macros in category 1 and category 2.

Figure 8 is a fraction of the modular translation of the FGS based on the approach, which shows the same portion of the model illustrated in Fig. 5 highlighting the changed part with bigger font.

5.6.3 Property-based model slicing

To further reduce the verification complexity, we apply slicing techniques [12, 15, 27] whenever applicable; for a given property to be verified, we perform a dependency analysis for variables appearing in the property and remove statements in the PROMELA model that do not have a dependency relationship with those variables. NuSMV supports a similar approach named *Cone of Influence Reduction* built in its model checking algorithm that has been used for checking all the FGS properties. SPIN, however, does not support built-in property-based slicing, and, thus, we have applied it manually to the PROMELA models.

5.6.4 Performance improvement after optimization

These optimization approaches enable SPIN to check the property **P1** within 13 minutes with the bit-state hashing option; the search depth it has explored is 3,632,455, the total number of states and transitions it has explored are 3.4×10^7 and 4.33×10^7, respectively. A total of 412.3 M of memory is consumed. Partial order reduction is used by default for all our experiments. The use of an exhaustive search algorithm, however, is still not feasible for this model. Since the bit-state hashing algorithm performs a partial search of the state space, the result can be unsound.

In comparison, the same property is verified using NuSMV within 10 seconds with 24 M of memory consumption.

5.7 Model checking two-sided FGSs using SPIN

Once we manage to check one-sided FGSs using SPIN, an extension to two-sided FGSs is fairly straightforward. We create two identical FGS processes with a communication channel between them without imposing any synchrony hypothesis. Instead of having variables for

```
// macros  (total 123 macros)
#define When_FD_Switch_Pressed
    ( values.msg[0]== On    &&
                        PREV_STEP_FD_Switch != On)

#define  When_Turn_FD_Off
                (When_FD_Switch_Pressed_Seen &&
                                values.msg[10] != 1   )

#define When_FD_Switch_Pressed_Seen
                (When_FD_Switch_Pressed &&
            No_Higher_Event_Than_FD_Switch_Pressed)
                        :

// enumeration type declaration
mtype = {Off, On, Cleared, Selected, Armed, Active,
Undefined, done, LEFT, RIGHT, Disengaged, Engaged,
Capture, Track};

// for input variable values and macros converted to local
// variables
typedef Array1 {
    mtype msg[33]  };

// for input variable values from the other side of FGS
typedef Array2 {
    mtype msg[13]   };

proctype env(chan out2fgs, otherinput)
{
  Array1 values;
  Array2 otherFGS;

do
 :: 1 ->
    // initial non-deterministic value assignments
   if
    :: 1 -> values.msg[0] = On;
    ::  1 -> values.msg[0] = Off;
   fi;

  // assignment of a macro value to corresponding
  // message field (category 2).

  values.message[26] = When_Turn_FD_Off;

   if
    :: 1 -> otherFGS.msg[0] = On;
    :: 1 -> otherFGS.msg[0] = Off;
   fi;
          :
}
```

```
proctype FGS (chan input_from_otherFGS,
                                input_from_env)
{
  Array2  otherFGS;
  Array1  env;

// FGS state variables (total 82 number of state
// variables)
mtype Onside_FD = Off;
bool Onside_FD_On = 0;
mtype Modes = Off;

 input_from_env?env;

// mode logic
   if
          :: !Is_This_Side_Active  ->
                input_from_otherFGS?otherFGS;
          :: else -> skip;
   fi;
          :
   if
    :: !Is_This_Side_Active  ->
                    Onside_FD = otherFGS.msg[0];
    :: (Onside_FD == Off) && When_Turn_FD_On ->
                                Onside_FD = On;
    :: (Onside_FD == On) && (      env.msg[26]==1   ) ->
                                Onside_FD = Off;
    :: else -> skip;
  fi;

 Onside_FD_On = (Onside_FD == On);

  if
    :: Is_This_Side_Active != 1 ->
                          Modes =otherFGS.msg[1];
    :: (Modes ==Off) && When_Turn_Modes_On  &&
                Is_This_Side_Active -> Modes = On;
    :: (Modes == On) && When_Turn_Modes_Off &&
                Is_This_Side_Active -> Modes = Off;
    :: else -> skip;
  fi;

          :
}
```

Fig. 8 Modular translation of FGS from RSML^{-e} to PROMELA

the *other side FGS* generated by the environment process with constraints, we can directly wire the two FGS processes so that a passive FGS can get the actual mode values from the other active FGS. There is no need to change the FGS process from the one-sided model, and all the random value generation and constraints for the values of *other side FGS* are removed from the *env*.

SPIN explored this two-sided FGS model for the same property **P1** using the bit-state hashing option focusing on one side FGS at a time. The **P1** is expressed in a general term "this side", since it was specified assuming one-sided FGSs, which can be rephrased as follows;

P1.1 If the left side is active and the mode annunciations are off, the mode annunciations shall be turned on when the onside FD is turned on.
P1.2 If the right side is active and the mode annunciations are off, the mode annunciations shall be turned on when the onside FD is turned on.

SPIN successfully checked **P1.1** within 8 minutes consuming a total of 349.3 M of actual memory and reaching the search depth of 1,076,816. The total number of states and transitions explored are 1.5×10^7 and 2.0×10^7, respectively. Though the result of the verification is based on a partial search and can be unsound, its expected coverage of the state-space is reported to be over 98% by SPIN.

6 Discussion

We have presented our initial experience in using SPIN on commercial avionics specifications. Our primary goal was to answer the question whether explicit model checking would have been more suitable for our specific problem domain in terms of its scalability and usability in comparison to those of symbolic model checking we have used for the same domain.

Our investigation reveals that SPIN performs poorer than the symbolic model checker NuSMV on the one-sided synchronous FGS model, but scales better to asynchronous two-sided FGSs once we manage to handle the one-sided FGS using SPIN. This is mainly because SPIN allows partial search to cope with high verification complexity, whereas NuSMV is designed mainly for exhaustive verification. SAT-based model checking supported by NuSMV can be considered as an alternative, trading incompleteness for scalability. Nevertheless, even the SAT-based model checking has failed to cope with the complexity of two-sided FGSs. Though the verification result might be unsound, SPIN may be a better choice when the size of the system is too big to use exhaustive model checking. We also observe that SPIN can be usable in practice in a sense that the optimization required in this work is systematic, and, thus, can be automated.

Nevertheless, we do not intend to put an emphasis on the performance differences because of three major reasons; (1) the result of using SPIN can be unsound because of the use of bit-state hashing that allows SPIN to perform a partial search of the system with the possibility of leaving states unexplored, (2) performance can be highly dependent on the level of optimization performed by the modeler, and (3) performance can vary depending on properties to be checked. We believe that our initial optimization approach is just a starting point, and far better optimization is possible as we gain better knowledge of the model checker. We also do not rule out the possibility of using NuSMV for the two-sided FGSs by using aggressive optimization and/or abstraction in the future.

There are a number of issues to be considered; first, our experience shows that SPIN requires a more aggressive optimization approach to make it work on FGSs, and can be more sensitive to the slight differences between models. Second, the modularization and restructuring of the system model for performance improvement is based on an understanding of the techniques and implementations of SPIN. This means that these verification tools still highly depend on the expertise of the user. Third, the result of the optimization can be difficult to understand. As we showed in Fig. 8, we have changed the names of the macros,

which are designed to improve readability of the specification, to incomprehensible message field names. This optimization approach may require support to help users interpret the optimized system model. Finally, SPIN supports verification of one property at a time and all optimization approaches need to be property-based. Considering that 298 properties need to be verified on the FGS routinely as the system evolves, we may need 298 different optimized models, one for each property, in the worst case. Nevertheless, we believe that property-based optimization is not a critical issue as long as we can automate the translation and the optimization process. We leave this claim to a future investigation.

References

1. Avrunin GS, Corbett JC, Dwyer MB (2000) Benchmarking finite-state verifiers. Softw Tools for Technol Transf 2(4):317–320
2. Chan W, Anderson RJ, Beame P, Burns S, Modugno F, Notkin D, Reese JD (1998) Model checking large software specifications. IEEE Trans Softw Eng 24(7):498–520
3. Choi Y, Heimdahl M (2002) Model checking RSML^{-e} requirements. In: Proceedings of the 7th IEEE/IEICE international symposium on high assurance systems engineering, pp 109–118
4. Clarke EM, Grumberg O, Peled D (1999) Model checking. MIT Press
5. Dong Y, Du X, Holzmann GJ, Smolka SA (2003) Fighting livelock in the GNU i-protocol: a case study in explicit-state model checking. Int J Softw Tools Technol Transf (4)
6. Dong Y, Du X, et al (1999) Fighting livelock in the i-protocol: a comparative study of verification tools. In: Proceedings of TACAS
7. Eisner C, Peled D (2002) Comparing symbolic and explicit model checking of a software system. In: SPIN Workshop
8. Berry G, Gonthier G (1992) The ESTEREL synchronous programming language: design, semantics, implementation. Sci Comput Programm 19(2):87–152
9. Giannakopoulou D, Pasareanu CS, Barringer H (2002) Assumption generation for software component verification. In: 17th IEEE international conference on automated software engineering, pp 3–12
10. Goering R (1997) Model checking expands verification's scope. Electron Eng Today
11. Halbwachs N, Caspi P, Raymond P, Pilaud D (1991) The synchronous dataflow programming language lustre. Proc IEEE 79(9):1305–1320
12. Hatcliff J, Dwyer MB, Zheng H (2000) Slicing software for model construction. Higher-Order and Symbolic Comput 13(4):315–353
13. Havelund K, Pressburger T (2000) Model checking Java programs using Java pathfinder. Int J Softw Tools Technol Transf 366–381
14. Heimdahl MPE, Whalen M, Thompson J (2003) NIMBUS: A tool for specification centered development. Presented at the 11th IEEE international requirements engineering conference
15. Heimdahl MPE, Whalen MW (1997) Reduction and slicing of hierarchical state machines. In: Proceedings of the 5th ACM SIGSOFT symposium on the foundations of software engineering
16. Heitmeyer C, Kirby J Jr, Labaw B, Archer M, Bharadwaj R (1998) Using abstraction and model checking to detect safety violations in requirements specifications. IEEE Trans Softw Eng 24(11):927–948
17. Holzmann GJ (1997) The model checker spin. IEEE Trans Softw Eng 23(5)
18. Holzmann GJ (1999) The engineering of a model checker: the gnu i-protocol case study revisited. In: SPIN workshop
19. Holzmann GJ (2003) The SPIN model checker: primer and reference manual. Addison-Wesley Publishing Company
20. Leveson NG, Heimdahl MPE, Reese JD (1999) Designing specification languages for process control systems: Lessons learned and steps to the future. In: Proceedings of the 7th ACM SIGSOFT symposium on the foundations on software engineering, LNCS, vol 1687, pp 127–145
21. Luettgen G, Carreno V (1999) Analyzing mode confusion via model checking. In: SPIN workshop
22. Miller S, Tribble A (2002) A methodology for improving mode awareness in flight guidance design. In: Digital avionics systems conference
23. Miller SP, Tribble AC, Heimdahl MPE (2003) Proving the shalls. In: Formal methods Europe
24. NuSMV: A New Symbolic Model Checking. Available at http://nusmv.irst.itc.it/
25. Owre S, Shankar N, Rushby JM (1993) User guide for the PVS specification and verification system. Computer Science Laboratory; SRI International, Menlo Park, CA 94025, beta release edition

26. Pnueli A (1977) The temporal logic of programs. In: Proceedings of the 18th IEEE symposium foundations of computer science, pp 46–57
27. Weiser M (1984) Program slicing. IEEE Trans Softw Eng SE-10(4):352–357
28. Whalen MW (2000) A formal semantics for RSML^{-e}. Master's thesis, University of Minnesota