

# DiConic Addition of Failsafe Fault-Tolerance\*

Ali Ebneenasir<sup>1</sup>

Department of Computer Science  
Michigan Technological University  
Houghton MI 49931 USA

Technical Report #CS-TR-07-03  
June 2007

## Abstract

We present a divide-and-conquer method, called DiConic, for automatic addition of failsafe fault-tolerance to distributed programs, where a failsafe program guarantees to meet its safety specification even when faults occur. Specifically, instead of adding fault-tolerance to a program as a whole, we separately revise program actions so that the entire program becomes failsafe fault-tolerant. Our DiConic algorithm has the potential to utilize the processing power of a large number of machines working in parallel, thereby enabling automatic addition of failsafe fault-tolerance to distributed programs with a large number of processes. We formulate our DiConic synthesis algorithm in terms of the satisfiability problem and apply it to several case studies using the Yices Satisfiability Modulo Theories (SMT) solver. We demonstrate our DiConic approach for the classic Byzantine Generals problem and an industrial application.

**Keywords:** Addition of fault-tolerance, Formal methods, Divide-and-conquer, Satisfiability

---

\*This work was sponsored by a grant from Michigan Technological University.

<sup>1</sup>Email: aebneenas@mtu.edu

Web: <http://www.cs.mtu.edu/~aebneenas>

Tel: +1-906-487-4372, Fax: +1-906-487-2283

# Contents

<b>1</b>	<b>Introduction</b>	<b>3</b>
<b>2</b>	<b>Preliminaries</b>	<b>5</b>
<b>3</b>	<b>Problem Statement</b>	<b>11</b>
<b>4</b>	<b>Adding Failsafe Fault-Tolerance</b>	<b>12</b>
<b>5</b>	<b>DiConic Addition</b>	<b>13</b>
5.1	Computing Fault-Span and Offending States . . . . .	14
5.2	Removing Offending Transitions . . . . .	17
5.3	Computing a New Invariant . . . . .	19
<b>6</b>	<b>Case Studies</b>	<b>21</b>
6.1	Altitude Switch Controller . . . . .	22
6.2	Cruise Control System . . . . .	26
6.3	Token Ring . . . . .	30
<b>7</b>	<b>Related Work</b>	<b>33</b>
<b>8</b>	<b>Conclusions and Future Work</b>	<b>34</b>

# 1 Introduction

Software intensive systems are inevitably subject to the occurrence of unanticipated faults as it is difficult (if not impossible) to anticipate all types of faults in early stages of development. Automatic addition of fault-tolerance to programs provides a systematic approach for equipping an existing program with necessary fault-tolerance functionalities while preserving its correctness in the absence of faults. However, the exponential complexity of automatic addition of fault-tolerance to *distributed* programs [1] urges us to devise new paradigms for the addition of fault-tolerance that scale better than existing techniques. In this paper, we present a *DIVide-and-CONquer* paradigm, called DiConic, that decomposes the problem of automatic addition of *failsafe* fault-tolerance - where a failsafe program meets its safety specification even when faults occur - into a set of sub-problems, where in each sub-problem, we focus on revising an individual program action (specified in terms of Dijkstra's guarded command language [2]) in such a way that the entire program becomes failsafe fault-tolerant.

Most existing approaches [3–13] for automated synthesis of fault-tolerant programs generate an *integrated* combinatorial model of a program (implemented as a reachability graph in enumerative techniques [10–12] or as a set of Binary Decision Diagrams (BDDs) in symbolic methods [8, 13]) which makes it difficult to scale such methods as the size<sup>2</sup> of the program increases. For example, specification-based approaches [3, 9] synthesize the abstract structure of a fault-tolerant program as a finite-state machine from the satisfiability proof of its temporal logic specification. Control-theoretic techniques [5, 7] use synchronous automata-theoretic product to generate finite state automata that represent fault-tolerant discrete-event controllers. Game-theoretic methods [4, 6] synthesize a winning strategy that captures all program behaviors. Previous work on addition of fault-tolerance to concurrent and distributed programs [10–12] takes an integrated model (represented as a finite-state machine) of an existing program and automatically revises the program model in order to generate a model of a fault-tolerant version thereof. Symbolic techniques [8, 13] enable the synthesis of larger programs, however, they still suffer from the complexity of synthesis as they deal with

---

<sup>2</sup>In this work, we determine the size of the program by the number of its processes, which has a direct relation with the total number of program actions, variables, and state space.

the synthesis problem as a whole without decomposing the problem. While all aforementioned approaches present important and useful techniques for synthesizing fault-tolerant programs and mitigating space/time complexity of addition, we believe that decomposing the synthesis problem so that each sub-problem can be solved separately is equally important and is orthogonal to complexity issues.

In order to decompose the problem of automatic addition of failsafe fault-tolerance, we focus on the following questions: Given a *fault-intolerant* program that satisfies its specification in the absence of faults, but provides no guarantees in the presence of faults, *how can we determine whether or not each program action should be revised so that the entire program becomes fault-tolerant?* If indeed a program action should be revised, *how do we independently identify the nature of the change for that action?* These questions are indeed special cases of a more general question, which is the focus of our research: Given a program that meets its specification and a new desired property raised due to the change in requirements, *how do we revise each program action so that the revised program satisfies the new property while preserving its existing specification?* To address this question for the addition of failsafe fault-tolerance, we present a sound distributed algorithm that has two components: *synthesis coordinator* and *synthesis node*. Corresponding to each program action, we instantiate a synthesis node that determines how that action should be revised towards the synthesis of a failsafe version of the input program. The synthesis coordinator synchronizes the activities of synthesis nodes. There is only one synchronization point between a synthesis node and the synthesis coordinator, which helps decreasing the communication cost of our algorithm. As the number of program actions increases one can instantiate new synthesis nodes and deploy them on different machines in a parallel platform.

Our proposed synthesis algorithm is a distributed fixpoint computation where in each round of the computation synthesis nodes report their results to the synthesis coordinator. More specifically, our algorithm takes a fault-intolerant program (as a set of actions) and a set of fault actions, and instantiates (i)  $N$  synthesis nodes, where  $N$  is the total number of program and fault actions, and (ii) one coordinator node. Each synthesis node performs a set of verification and revision operations that could be done irrespective of revisions performed on

other actions in other synthesis nodes. Afterwards, each synthesis node reports its revised action to the coordinator node and waits for coordinator’s reply. At the end of each round, the collection of the actions in all synthesis nodes comprises an *intermediate program*. The coordinator verifies some constraints on the intermediate program and broadcasts the results to all synthesis nodes that are waiting to enter the next round of synthesis. A synthesis node terminates in two possible cases. First, the synthesis node *locally* determines that its associated action should be removed in the failsafe program. Second, the synthesis node receives a termination message from the coordinator. The coordinator terminates either by finding a failsafe program or by declaring failure.

We have employed the Yices Satisfiability Modulo Theories (SMT) solver [14] (developed at SRI International) to apply our DiConic approach to four case studies among which (i) the classic problem of Byzantine Generals (BG) [15], and (ii) an altitude switch controller program that has been manually designed at the Naval Research Laboratory [16]. The Yices specifications of these case studies are available in the Appendix. The organization of this report is as follows: We present preliminary concepts in Section 2. In Section 3, we represent the problem of adding failsafe fault-tolerance. Subsequently, in Section 4, we reiterate an enumerative solution due to Kulkarni *et al.* [17] where they add failsafe fault-tolerance to distributed programs represented as finite-state machines. Then, in Section 5, we present our DiConic algorithm. We use the BG problem as a running example to demonstrate our approach. Additionally, in Section 6, we present three case studies for DiConic addition of failsafe fault-tolerance to (i) a simplified version of an altitude switch controller (from [16]), (ii) a cruise control system, and (iii) a token ring protocol. We discuss related work in Section 7. Finally, we make concluding remarks and discuss future work in Section 8.

## 2 Preliminaries

In this section, we provide formal definitions of programs, problem specifications, faults, and failsafe fault-tolerance. The definition of specifications is adapted from Alpern and Schneider [18]. The definition of programs, faults and fault-tolerance is adapted from Arora and Gouda [19] and Arora and Kulkarni [20]. The issues of modeling distributed programs

is adapted from Attie and Emerson [21], and Kulkarni and Arora [10]. To illustrate our modeling approach, we use the Byzantine Generals (BG) problem [15] as a running example and use sans serif font for the exposition of the BG example.

**Programs and processes.** A *program*  $p = \langle V_p, \Pi_p \rangle$  is a tuple of a finite set  $V_p$  of variables and a finite set  $\Pi_p$  of processes. Each variable  $v_i \in V_p$ , for  $1 \leq i \leq n$ , has a finite non-empty domain  $D_i$ . A *state*  $s$  of a program  $p$  is a valuation  $\langle d_1, d_2, \dots, d_n \rangle$  of program variables  $\langle v_1, v_2, \dots, v_n \rangle$ , where  $d_i$  is a value in  $D_i$ . The *state space*  $\mathcal{Q}_p$  is the set of all possible states of  $p$ . A *transition* of  $p$  is of the form  $(s, s')$ , where  $s$  and  $s'$  are program states. A *process*  $P_j$ ,  $1 \leq j \leq k$ , includes a finite set of actions. An action is a guarded command (due to Dijkstra [2]) of the form  $grd \rightarrow stmt$ , where  $grd$  is a Boolean expression specified over  $V_p$  and  $stmt$  is an assignment that *atomically* updates zero or more variables. An assignment always terminates once executed. The set of program actions is the union of the actions of all its processes. We represent the new values of updated variables as *primed* values. For example, if an action updates the value of an integer variable  $v_1$  from 0 to 1, then we have  $v_1 = 0$  and  $v'_1 = 1$ .

**BG example.** We consider the canonical version of the Byzantine generals problem [15] where there are 4 distributed processes  $P_g, P_j, P_k$ , and  $P_l$  such that  $P_g$  is the general and  $P_j, P_k$ , and  $P_l$  are the non-generals. (An identical explanation is applicable for arbitrary number of non-generals.) In the fault-intolerant BG program, the general sends its decision to non-generals and subsequently non-generals output their decisions. Thus, each process has a variable  $d$  to represent its decision, a Boolean variable  $b$  to represent if that process is Byzantine, and a variable  $f$  to represent if that process has finalized (output) its decision. A Byzantine process may arbitrarily change its  $d$  and/or  $f$  values. The program variables and their domains are as follows:

$$\begin{array}{ll}
 d.g : \{0, 1\} ; d.j, d.k, d.l : \{0, 1, \perp\} & // \perp \text{ denotes uninitialized decision} \\
 b.g, b.j, b.k, b.l : \{true, false\} & // b.j = true \text{ iff } P_j \text{ is Byzantine} \\
 f.j, f.k, f.l : \{false, true\} & // f.j = true \text{ iff } P_j \text{ has finalized its decision}
 \end{array}$$

We represent the actions of the non-general process  $P_j$  as follows. We label these actions with  $BG_1$  and  $BG_2$ . (The actions of other non-generals are similar.)

$$\begin{array}{ll}
 BG_1 : d.j = \perp \wedge \neg f.j & \longrightarrow d.j := d.g; \\
 BG_2 : d.j \neq \perp \wedge \neg f.j & \longrightarrow f.j := true;
 \end{array}$$

A non-general process that has not yet decided copies the decision of the general. When a non-general process decides, it can finalize its decision.

**State and transition predicates.** A *state predicate* of  $p$  is a subset of  $\mathcal{Q}_p$  specified as a Boolean expression over  $V_p$ .<sup>3</sup> An *unprimed* state predicate is specified only in terms of unprimed variables. Likewise, a *primed* state predicate includes only primed variables. A *transition predicate* (adapted from [22, 23]) is a subset of  $\mathcal{Q}_p \times \mathcal{Q}_p$  represented as a Boolean expression over both unprimed and primed variables. We say a state predicate  $X$  *holds at a state*  $s$  (respectively,  $s \in X$ ) if and only if (iff)  $X$  evaluates to true at  $s$ . Note that, a state predicate  $X$  also represents a transition predicate that includes all transitions  $(s, s')$ , where either  $s \in X$  and  $s'$  is an arbitrary state, or  $s$  is an arbitrary state and  $s' \in X$ . An action  $grd \rightarrow stmt$  is enabled at a state  $s$  iff  $grd$  holds at  $s$ . A process  $P_j \in \Pi_p$  is enabled at  $s$  iff there exists an action of  $P_j$  that is enabled at  $s$ . We define a function *Primed* (respectively, *UnPrimed*) that takes a state predicate  $X$  (respectively,  $X'$ ) and substitutes each variable in  $X$  (respectively,  $X'$ ) with its primed (respectively, unprimed) version, thereby returning a state predicate  $X'$  (respectively,  $X$ ). The function *getPrimed* (respectively, *getUnPrimed*) takes a transition predicate  $T$  and returns a primed (respectively, an unprimed) state predicate representing the set of destination (respectively, source) states of all transitions in  $T$ .

**BG example.** We define a state predicate  $\mathcal{I}_1$  that captures the set of states in which the general is not Byzantine and at most one non-general could be Byzantine. (Process variables  $p$  and  $q$  represent non-general processes in the quantifications.)

$$\begin{aligned} \mathcal{I}_1 = & \neg b.g \wedge (\neg b.j \vee \neg b.k) \wedge (\neg b.k \vee \neg b.l) \wedge (\neg b.l \vee \neg b.j) \wedge \\ & (\forall p :: \neg b.p \Rightarrow (d.p = \perp \vee d.p = d.g)) \wedge \\ & (\forall p :: (\neg b.p \wedge f.p) \Rightarrow (d.p \neq \perp)) \wedge \end{aligned}$$

In this case, a non-general non-Byzantine process is either undecided or its decision is the same as that of general. Moreover, all non-general non-Byzantine processes that are finalized have

---

<sup>3</sup>An individual state  $s$ , the empty set, and the entire state space (i.e., the universal set) are special cases of a state predicates.

decided on a non- $\perp$  value. As another example, the state predicate  $\mathcal{I}_2$  captures a set of states where the general is Byzantine and all non-generals have decided on the same value.

$$\mathcal{I}_2 = b.g \wedge \neg b.j \wedge \neg b.k \wedge \neg b.l \wedge (d.j = d.k = d.l \wedge d.j \neq \perp)$$

**Closure.** A state predicate  $X$  is *closed in an action*  $grd \rightarrow stmt$  iff executing  $stmt$  from a state  $s \in (X \wedge grd)$  results in a state in  $X$ . We say a state predicate  $X$  is *closed in a program*  $p$  iff  $X$  is closed all actions of  $p$ .

**BG example.** The state predicates  $\mathcal{I}_1$  and  $\mathcal{I}_2$  are closed in actions  $BG_1$  and  $BG_2$ . We leave it to the reader to verify this claim.

**Program computations and execution semantics.** We consider a nondeterministic interleaving of all program actions generating a sequence of states. A *computation* of a program  $p = \langle V_p, \Pi_p \rangle$  is a sequence  $\sigma = \langle s_0, s_1, \dots \rangle$  of states that satisfies the following conditions: (1) for each transition  $(s_i, s_{i+1})$  ( $i \geq 0$ ) in  $\sigma$ , there exists an action  $grd \rightarrow stmt$  in some process  $P_j \in \Pi_p$  such that  $grd$  holds at  $s_i$  and the execution of  $stmt$  at  $s_i$  yields  $s_{i+1}$ ; (2)  $\sigma$  is *maximal* in that either  $\sigma$  is infinite or if it is finite, then no action is enabled in its final state, and (3)  $\sigma$  is *fair*; i.e., if a process  $P_j$  is continuously enabled in  $\sigma$ , then eventually some action of  $P_j$  will be executed.

**Distribution model.** We model distribution issues as a set of read/write restrictions imposed on program processes. More specifically, we associate with each process  $P_j$  a set of variables that it is allowed to read, denoted  $r_j$ , and a set of variables that  $P_j$  can write, denoted  $w_j$ . We assume that for each process  $P_j$ ,  $w_j \subseteq r_j$ ; i.e., if a process can write a variable, then that variable is readable too. No action in a process  $P_j$  is allowed to update a variable  $v \notin w_j$ . This constraint can be specified as the transition predicate  $rwRest \equiv (\forall v : v \notin w_j : v = v')$ ; i.e., the value of an unprimed variable  $v \notin w_j$  should be equal to the value of its primed version. Note that the transition predicate  $rwRest$  imposes a similar constraint on unreadable variables because if a variable cannot be read, then it cannot be written either. Using read/write restrictions  $rwRest$ , we formally specify an action  $grd \rightarrow stmt$  as a transition predicate  $grd \wedge Primed(stmtExpr) \wedge rwRest$ , where  $stmtExpr$  is a Boolean expression generated from the assignment  $stmt$ . For example, an



assignment  $x := 1$  can be specified as the expression  $x' = 1$ .

**BG example.** Each non-general non-Byzantine process  $P_j$  is allowed to only read  $r_j = \{b.j, d.j, f.j, d.k, d.l, d.g\}$  and it can only write  $w_j = \{d.j, f.j\}$ . Hence, we have  $w_j \subseteq r_j$ . As an example of a transition predicate, we represent the action  $BG_1$  as  $((d.j = \perp) \wedge (f.j = 0)) \wedge (d.j' = d.g') \wedge rwRest$ , where  $rwRest$  is the transition predicate  $(d.k = d.k') \wedge (d.l = d.l') \wedge (d.g = d.g') \wedge (b.g = b.g') \wedge (b.j = b.j') \wedge (b.k = b.k') \wedge (b.l = b.l') \wedge (f.k = f.k') \wedge (f.l = f.l')$ .

**Specifications.** We follow Alpern and Schneider [18] in representing the specification  $spec$  of a program as the conjunction of a *safety* specification (denoted  $\mathcal{S}$ ) and a *liveness* specification (denoted  $\mathcal{L}$ ); i.e.,  $spec = \mathcal{S} \cap \mathcal{L}$ . Intuitively,  $\mathcal{S}$  stipulates that nothing bad ever happens. Formally, we represent  $\mathcal{S}$  as a transition predicate whose transitions must not appear in program computations. An action  $grd \rightarrow stmt$  satisfies  $\mathcal{S}$  from a state predicate  $X$  iff the transition predicate  $\mathcal{X} \wedge grd \wedge Primed(stmtExpr) \wedge \mathcal{S} \wedge rwRest$  is equal to the empty set, where  $rwRest$  captures the read/write restrictions of the process that executes the action. A program  $p$  satisfies its safety specification  $\mathcal{S}$  from  $\mathcal{X}$  iff all actions of  $p$  satisfy  $\mathcal{S}$  from  $\mathcal{X}$ . In terms of program computations, a program satisfies  $\mathcal{S}$  from  $\mathcal{X}$  iff all program computations starting in  $\mathcal{X}$  satisfy  $\mathcal{S}$ . A program computation  $c = \langle s_0, s_1, \dots \rangle$  satisfies  $\mathcal{S}$  from  $\mathcal{X}$  iff  $(s_0 \in \mathcal{I}) \wedge (\forall (s_i, s_{i+1}) : i \geq 0 : (s_i, s_{i+1}) \notin \mathcal{S})$ . Otherwise, the computation  $c$  violates  $\mathcal{S}$ . The liveness specification  $\mathcal{L}$  states that something good will eventually occur. More precisely, we define  $\mathcal{L}$  as a set of infinite sequences of states. A computation is in  $\mathcal{L}$  if it contains a suffix<sup>4</sup> that is in  $\mathcal{L}$ . Note that there may exist infinitely long sequences that have no suffix in the particular set of infinite sequences specified by  $\mathcal{L}$ . For example, given a state  $s$  that does not appear in any sequence  $\langle s_0, s_1, \dots \rangle$  belonging to  $\mathcal{L}$ , the infinite sequence generated by a self-loop on  $s$  would not belong to  $\mathcal{L}$ . A computation  $c = \langle s_0, s_1, \dots \rangle$  satisfies  $\mathcal{L}$  from a state predicate  $X$  iff  $(s_0 \in \mathcal{I})$  and  $c$  has a suffix in  $\mathcal{L}$ . We say a program  $p$  satisfies its liveness specification  $\mathcal{L}$  from a state predicate  $X$  iff all computations of  $p$  satisfy  $\mathcal{L}$  from  $X$ . A program  $p$  satisfies its specification  $spec$  from a state predicate  $X$  iff  $p$  satisfies its safety and liveness specifications from  $X$ .

**BG example.** The safety specification of the BG program requires that *Validity* and *Agreement*

---

<sup>4</sup>A suffix of a sequence  $\langle s_0, s_1, \dots \rangle$  is a subsequence  $\langle s_j, s_{j+1}, \dots \rangle$  for some  $j \geq 0$ .

be satisfied. Validity stipulates that if the general is not Byzantine and a non-Byzantine non-general has finalized its decision, then the decision of that non-general process is the same as that of the general. Agreement requires that if two non-Byzantine non-generals have finalized their decisions, then their decisions are identical. Hence, the program should not execute transitions that reach the primed state predicate  $\mathcal{R}_1$ , where

$$\begin{aligned} \mathcal{R}_1 = & (\exists p, q :: \neg b.p' \wedge \neg b.q' \wedge d.p' \neq \perp \wedge d.q' \neq \perp \wedge d.p' \neq d.q' \wedge f.p' \wedge f.q') \vee \\ & (\exists p :: \neg b.g' \wedge \neg b.p' \wedge d.p' \neq \perp \wedge d.p' \neq d.g' \wedge f.p') \end{aligned}$$

Moreover, when a non-Byzantine process finalizes, it is not allowed to change its decision. Thus, the set of transitions of the following transition predicate should not be executed as well:

$$\mathcal{S}_2 = \neg b.j \wedge \neg b.j' \wedge f.j \wedge (d.j \neq d.j' \vee f.j \neq f.j')$$

Let  $\mathcal{S}_1$  be the transition predicate representing all transitions that reach a state in  $\mathcal{R}_1$ . We specify the safety specification of the BG program as the transition predicate  $\mathcal{S}_1 \vee \mathcal{S}_2$ .

**Invariants.** A state predicate  $\mathcal{I}$  is an invariant of a program  $p$  for its specification  $spec$  iff the following conditions are satisfied: (1)  $\mathcal{I}$  is closed in  $p$ , and (2)  $p$  satisfies  $spec$  from  $\mathcal{I}$ .

**BG example.** The state predicate  $\mathcal{I}_{BG} = \mathcal{I}_1 \vee \mathcal{I}_2$  is indeed an invariant of the BG program.

**Faults and fault-span.** We represent a fault-type  $F$  as a set of actions. Fault actions differ from program actions in that the program does not have execution control over fault actions. A *computation of a program*  $p = \langle V_p, \Pi_p \rangle$  *in the presence of fault*  $F$  is a sequence  $\sigma = \langle s_0, s_1, \dots \rangle$  of states that satisfies the following conditions: (1) for each transition  $(s_i, s_{i+1})$  ( $i \geq 0$ ) in  $\sigma$ , there exists either a program or a fault action  $grd \rightarrow stmt$  such that  $grd$  holds at  $s_i$  and the execution of  $stmt$  at  $s_i$  yields  $s_{i+1}$ ; (2)  $\sigma$  is *maximal*, and (3)  $\sigma$  is *fair*. A state predicate  $\mathcal{FS}$  is called a *fault-span* of a program  $p$  from its invariant  $\mathcal{I}$  for fault  $F$  (denoted  $F$ -span) iff the following conditions are satisfied: (1)  $\mathcal{I} \Rightarrow \mathcal{FS}$ , and (2)  $\mathcal{FS}$  is closed in program  $p$  and fault actions. Intuitively, the  $F$ -span of  $p$  from  $\mathcal{I}$  is a boundary around  $\mathcal{I}$  to which the state of  $p$  can be perturbed by fault and program actions.

**BG example.** The fault action  $F_1$  may cause *at most* one non-Byzantine process to become Byzantine. A Byzantine process may arbitrarily change its  $d$  and/or  $f$  values. (We include similar fault actions for  $k, l$  and  $g$ .)

$$\begin{aligned}
F_1 &: \neg b.g \wedge \neg b.j \wedge \neg b.k \wedge \neg b.l \quad \longrightarrow \quad b.j := true \\
F_2 &: b.j \quad \longrightarrow \quad d.j, f.j := 0|1, false|true \text{ }^5
\end{aligned}$$

In Section 5, we shall illustrate how to calculate the fault-span in a DiConic manner.

**Failsafe fault-tolerance.** A program  $p$  is *failsafe*  $F$ -tolerant from an invariant  $\mathcal{I}$  for its specification  $spec$  (i.e., fault-tolerant against a fault-type  $F$ ) iff there exists an  $F$ -span  $\mathcal{FS}$  such that: (1) in the absence of fault  $F$ ,  $p$  satisfies  $spec$  from  $\mathcal{I}$ , and (2) in the presence of fault  $F$ , the actions of  $p$  and  $F$  satisfy the safety of  $spec$  (i.e.,  $\mathcal{S}$ ) from  $\mathcal{FS}$ .

### 3 Problem Statement

In this section, we reiterate the problem of adding failsafe fault-tolerance to programs from [10]. Specifically, to formulate our DiConic approach in terms of the satisfiability problem, we represent the addition problem in terms of state/transition predicates.

In order to separate fault-tolerance from functional concerns, the problem of adding fault-tolerance stipulates that no new behaviors are added to programs in the *absence* of faults. More precisely, given a program  $p = \langle V_p, \Pi_p \rangle$ , its invariant  $\mathcal{I}$ , its specification  $spec = \mathcal{S} \cap \mathcal{L}$  and a fault-type  $F$ , we aim to generate a revised version of  $p$ , denoted  $p^f = \langle V_p, \Pi_p^f \rangle$  with a new invariant  $\mathcal{I}^f$  such that  $p^f$  is failsafe  $F$ -tolerant from  $\mathcal{I}^f$  for  $spec$ . If  $\mathcal{I}^f$  includes a state  $s$  that does not belong to  $\mathcal{I}$ , then the execution of  $p$  in the absence of  $F$  from  $s$  may generate new computations. Hence, we require that  $\mathcal{I}^f \Rightarrow \mathcal{I}$ . Moreover, starting in  $\mathcal{I}^f$ , the actions of  $p^f$  should not include new transitions. Otherwise,  $p^f$  may exhibit new behaviors in the absence of faults. Thus, for each action  $grd^f \rightarrow stmt^f$  in  $p^f$ , we require that there exists an action  $grd \rightarrow stmt$  in  $p$  such that the transition predicate  $\mathcal{I}^f \wedge grd^f \wedge Primed(stmt^f Expr) \wedge rwRest$  implies the transition predicate  $\mathcal{I}^f \wedge grd \wedge Primed(stmt Expr) \wedge rwRest$ . Therefore, we formally state the problem of adding failsafe fault-tolerance as follows:

---

<sup>5</sup> $d.j := 0|1$  means that  $d.j$  could be assigned either 0 or 1.

### The Problem of Adding Failsafe Fault-Tolerance.

Given  $p$ ,  $\mathcal{I}$ ,  $spec$ , and  $F$ , identify  $p^f$  and  $\mathcal{I}^f$  such that

- (1)  $\mathcal{I}^f \Rightarrow \mathcal{I}$ ,
- (2) For each action  $grd^f \rightarrow stmt^f$  in  $p^f$ , there exists some action  $grd \rightarrow stmt$  in  $p$  such that
$$(\mathcal{I}^f \wedge grd^f \wedge Primed(stmt^f Expr) \wedge rwRest) \Rightarrow (\mathcal{I}^f \wedge grd \wedge Primed(stmt Expr) \wedge rwRest)$$
where  $rwRest$  captures the read/write restrictions of the process that executes the action  $grd^f \rightarrow stmt^f$ .
- (3)  $p^f$  is failsafe  $F$ -tolerant from  $\mathcal{I}^f$  for  $spec$ .  $\square$

## 4 Adding Failsafe Fault-Tolerance

In order to simplify the presentation of our DiConic approach, in this section, we present an intuitive description of the algorithm presented in previous work [12,17] for the addition of fault-tolerance to integrated models of distributed programs. This algorithm (see Figure 1) takes a non-deterministic finite-state machine representing the fault-intolerant program and generates a state machine representing the failsafe fault-tolerant program.

The algorithm in Figure 1 first identifies a *valid* fault-span of the program that is a superset of the invariant and is closed in program and fault transitions. Subsequently, the algorithm computes and removes the set of *offending states* from where a sequence of fault transitions can directly violate safety. The removal of offending states is accomplished by making them unreachable. Then the algorithm identifies and removes a set of offending transitions that either directly violate safety or reach an offending state. However, since the offending states are made unreachable in Step 2, offending transitions that start in an offending state need not be removed. While removing offending states/transitions, the addition algorithm may introduce states that have no outgoing transitions; i.e., *deadlock* states. The exclusion of deadlock states from the invariant may leave transitions that start in the new invariant and reach a recently removed state, thereby violating the closure of the new invariant. These transitions are removed in Step 5. The algorithm iterates through Steps 4 and 5 until either all invariant states are removed or a closed non-empty invariant is found. Existing implementations [12,13,17] of the algorithm in Figure 1 execute all steps on an integrated model

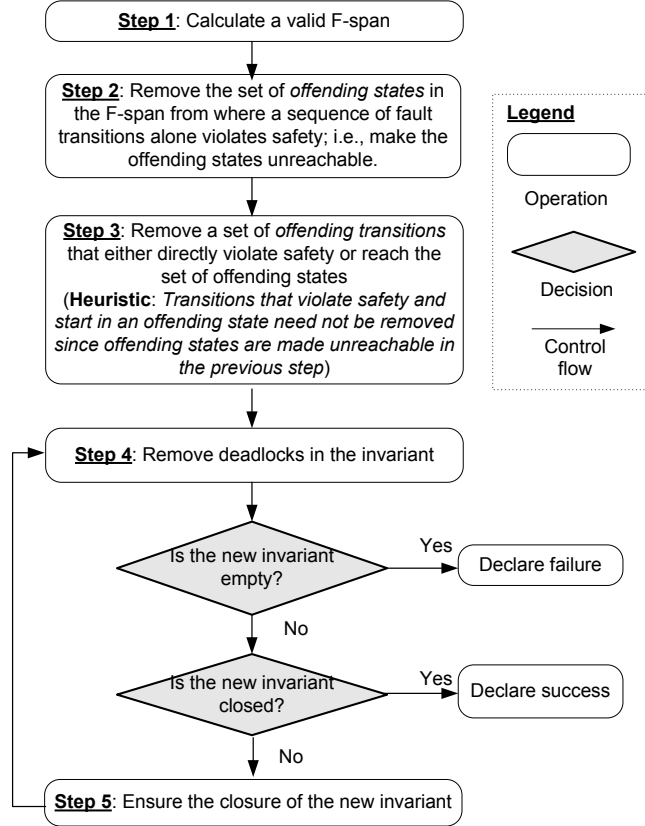


Figure 1: Adding failsafe fault-tolerance to non-deterministic finite state automata.

represented either as a non-deterministic finite-state machine [12] or as a set of symbolic entities (e.g., BDDs) [13], thereby making it difficult to scale these algorithms as the number of program processes increases.

## 5 DiConic Addition

In this section, we present our DiConic algorithm. Specifically, in Section 5.1, we illustrate how to compute the fault-span and the set of offending states (Steps 1 and 2 in Figure 1). In Section 5.2, we present a DiConic method for removing offending transitions (Step 3 in Figure 1). Finally, in Section 5.3, we calculate a new invariant for the synthesized failsafe program in a DiConic fashion. This step corresponds to the loop that includes Steps 4 and 5 in Figure 1.

## 5.1 Computing Fault-Span and Offending States

In this section, we decompose the problem of calculating a fault-span and the set of offending states by introducing a distributed forward/backward reachability algorithm. Specifically, in calculating the fault-span using forward reachability, the integrated algorithms [12, 13, 17] implement a fixpoint computation that explicitly explores the set of states reachable from the invariant of the fault-intolerant program by all program and fault transitions. In each iteration, the forward reachability algorithm computes a new set of states reachable from previously calculated states. Continuing thus, the forward reachability algorithm reaches a point where no more states are reachable. The union of the sets of all states computed in all iterations comprises the fault-span of the fault-intolerant program. Instead of calculating a set of states reachable from a given state predicate  $X$ , in each iteration, the backward reachability algorithm computes a set of states from where  $X$  is reached.

Our proposed DiConic approach for the calculation of the fault-span (respectively, the set of offending states) is a divide-and-conquer fixpoint computation with two kinds of components, namely the *coordinator* and the *reachability* nodes. A reachability node contains a single program action and computes the set of states reached by the execution of that action from a state predicate  $X$  (respectively, the set of states from where the execution of that action reaches a state in  $X$ ). The coordinator node manages the reachability nodes. Figures 2 and 3 illustrate the two components of our DiConic approach for forward reachability. (We omit the DiConic backward reachability algorithm as it is similar to the forward reachability.) The coordinator starts by sending out a *Base* state predicate to all nodes. Each node computes the set of states outside *Base* that are reachable from *Base* by that node's action (see Figure 2). The coordinator takes the union of all reachable states, denoted *reachedStates*, calculated by all nodes. If this union is empty, then the fixpoint computation is terminated. Otherwise, another iteration starts with a new *Base*, which is the union of the old *Base* and the state predicate *reachedStates*. For the calculation of the fault-span, the *Base* predicate is the invariant of the fault-intolerant program, and corresponding to each program and fault action, a reachability node is instantiated. Note that for the calculation of the set of offending states using DiConic backward reachability, only fault actions are considered since

we want to compute the set of states from where safety is violated by a sequence of fault transitions *alone*. Specifically, for each fault action, we first identify the set of states from where that action directly violates safety. The union of all such states comprises the base set for a backward reachability computation using only fault actions. Therefore, in the rest of this section, we present our DiConic approach assuming that the fault-span, denoted  $FS$ , and the set of offending states, denoted  $OS$ , have already been calculated.

```

ForwardReachability_Node( $grd \rightarrow stmt$ : action;  $X$ : state predicate;
                         $rwRest$ : transition predicate)
{
- Wait for  $X$  from the coordinator;
- repeat {
  -  $reachedStates := \emptyset$ ; // set of states reached outside  $X$  by
    // the execution of  $grd \rightarrow stmt$  from  $X$ ;
  -  $transPred := (grd \wedge X) \wedge \text{Primed}(stmtExpr) \wedge$ 
     $rwRest \wedge \text{Primed}(\neg X)$ ;
  - If ( $transPred$  is satisfiable) then
    -  $reachedStates := \text{getPrimed}(transPred)$ ;
  - Send  $reachedStates$  to coordinator;
  - Wait for  $X$  from the coordinator;
- } until(termination signal is received);
}

```

Figure 2: Reachability node in DiConic forward reachability.

BG example. In order to calculate the fault-span of the canonical BG program, denoted  $\mathcal{FS}_{BG}$ , we instantiate 14 forward reachability nodes as we have 2 program actions and 2 fault actions corresponding to each non-general process and 2 fault actions corresponding to the general process. We also create a reachability coordinator. The base predicate is equal to the invariant  $\mathcal{I}_{BG}$ . We consider only the program/fault actions of process  $P_j$  as they are structurally similar to other non-general processes. Clearly, in the first iteration of the fixpoint computation, starting from the invariant, program actions  $BG_1$  and  $BG_2$  would reach to a state in the invariant (due to the closure of the invariant in program actions).

The reachability node corresponding to the fault action  $F_1$  returns  $\mathcal{I}_1 \vee (b.g \wedge \neg b.j \wedge \neg b.k \wedge \neg b.l)$  as the set of states reachable from the invariant by  $F_1$ . Specifically, starting from the state predicate  $\mathcal{I}_1$  (see Section 2) in the invariant, if the fault action  $F_1$  does not cause the general

process  $P_g$  to become Byzantine, then the set of states reached from  $\mathcal{I}_1$  is the same as  $\mathcal{I}_1$ . Nonetheless, if  $P_g$  becomes Byzantine, then the  $b$  values of non-general processes would remain false as at most one process could be Byzantine. Thus, the set of states reached from  $\mathcal{I}_1$  by faults is  $(\mathcal{I}_1 \vee (b.g \wedge \neg b.j \wedge \neg b.k \wedge \neg b.l))$ . If the BG program is in a state in  $\mathcal{I}_2$ , then the fault action  $F_1$  will not be executed as  $P_g$  is already Byzantine.

The reachability node corresponding to the fault action  $F_2$  returns the same predicate as that of  $F_1$ . If  $P_j$  is Byzantine, then the execution of the action  $F_2$  from  $\mathcal{I}_1$  would not violate the closure of  $\mathcal{I}_1$  since the second and the third conjuncts of  $\mathcal{I}_1$  are specified for non-Byzantine processes. Moreover,  $F_2$  will not be enabled from  $\mathcal{I}_2$ . If  $P_g$  is Byzantine, then  $F_2$  will not be enabled from  $\mathcal{I}_1$ , thereby preserving the closure of  $\mathcal{I}_1$ . In this case, a similar reasoning as that for  $F_1$  would yield  $\mathcal{I}_2$  as the set of states reachable from  $\mathcal{I}_2$  by  $F_2$ .

Since the state predicate  $\mathcal{I}_1 \vee (b.g \wedge \neg b.j \wedge \neg b.k \wedge \neg b.l)$  is closed in all program and fault actions, the fault-span of the BG program is equal to  $\mathcal{FS}_{BG}$ , where

$$\mathcal{FS}_{BG} = \mathcal{I}_1 \vee (b.g \wedge \neg b.j \wedge \neg b.k \wedge \neg b.l).$$

```

Reachability_Coordinator(Base: state predicate)
{ //  $N$  denotes the total number of program and fault actions.
  -  $X := Base$ ;
  - repeat {
    - Send  $X$  to all nodes;
    - Wait for  $reachedStates_i$  from Node $_i$ ;
    -  $reachedStates := reachedStates_1 \vee \dots \vee reachedStates_N$ ;
    -  $X := X \vee reachedStates$ ;
  - } until( $reachedStates = \emptyset$ );
  - Send termination signal to all nodes;
  - return  $X$ ;
}

```

Figure 3: Reachability coordinator in DiConic forward reachability.

Both fault actions may violate safety if they start in a state in  $UnPrimed(\mathcal{R}_1)$  (see Section 2); i.e., the base set for backward reachability is  $\mathcal{R}_1$ . Subsequently, we compute the set of all states from where the execution of fault actions would yield a state in  $\mathcal{R}_1$ . The fault actions  $F_1$  and  $F_2$  yield a state in  $\mathcal{R}_1$  only if they are enabled in  $UnPrimed(\mathcal{R}_1)$ . Therefore, the second iteration of the fixpoint computation is the last iteration, returning  $UnPrimed(\mathcal{R}_1)$  as the set of offending



states.

## 5.2 Removing Offending Transitions

We decompose the removal of the offending transitions into the elimination of such transitions in each action without actually creating an integrated model of the program. Specifically, each synthesis node investigates two cases for the existence of offending transitions in the set of transitions represented by its associated action. Steps 1 and 2 in Figure 4 identify and exclude a subclass of the offending transitions that start outside the set of offending states and reach an offending state. The motivation behind these steps is that such transitions take the program to states from where the occurrence of faults may violate safety. Thus, in order to preserve safety even when faults occur, the program must not take such transitions. Steps 3 and 4 in Figure 4 determine and remove another subclass of the offending transitions that start outside the invariant in the fault-span  $\mathcal{FS}$  and directly violate the safety specification. In the next subsection, we describe how to divide the task of resolving deadlock states that may be created due to the elimination of offending transitions.

**BG example.** We create two synthesis nodes corresponding to the actions  $BG_1$  and  $BG_2$ . The execution of Steps 1 and 2 in Figure 4 for the action  $BG_1$  returns an empty set since the only way the execution of  $BG_1$  generates a state in the set of offending states is that  $BG_1$  is enabled in an offending state. In other words, the action  $BG_1$  may violate safety if it is enabled in a state in the fault-span where the general is Byzantine and the other two non-generals have finalized with different decisions. However, since such states have already been included in the set of offending states and Steps 1 and 2 of the Synthesis\_Node ensure that the offending states are not reached, we do not have to exclude such transition from the action  $BG_1$ . Therefore, the first two steps of the Synthesis\_Node algorithm do not change action  $BG_1$ .

The synthesis node corresponding to  $BG_2$  revises  $BG_2$  in the following way. Specifically, if the action  $BG_2$  is enabled from a state outside the set of offending states where a non-general non-Byzantine process has finalized with a different decision from that of  $P_j$ , then the safety of specification will be violated if  $P_j$  finalizes its decision. Thus, a set of transitions starting in  $(d.k \neq \perp \wedge d.j \neq d.k) \vee (d.l \neq \perp \wedge d.j \neq d.l)$  would be excluded from  $BG_2$  in Steps 1 and

```

Synthesis_Node(grd → stmt: action; OS, Inv, FS: state predicate;
               S, rwRest: transition predicate)
/* S denotes the safety specification and FS denotes the fault-span. */
/* OS denotes the set of offending states. */
{
  /* Does grd → stmt start outside OS and reach OS? */
  - transPred := (grd ∧ ¬OS) ∧
                 (Primed(stmtExpr ∧ OS)) ∧ rwRest; (1)
  - If (transPred is satisfiable) then
    - Exclude(transPred, grd → stmt); (2)

  /* Does grd → stmt directly violate safety in FS – Inv? */
  - transPred := (FS ∧ ¬Inv ∧ grd ∧ Primed(stmtExpr) ∧
                 rwRest ∧ S); (3)
  - If (transPred is satisfiable) then
    - Exclude(transPred, grd → stmt); (4)

  /* Synchronize this node with the synthesis coordinator. */
  repeat {
    - Send grd to the synthesis coordinator; (5)
    - If (grd is unsatisfiable) then declare that
      grd → stmt is removed; exit(); (6)
    - Wait to receive a new invariant Invnew and DeadlockStates; (7)

    /* Does the revised action contain closure-violating transitions? */
    - closureViolatingTrans := grd ∧ Invnew ∧ Primed(stmtExpr) ∧
                               Primed(¬Invnew) ∧ rwRest; (8)
    - If (closureViolatingTrans is satisfiable) then
      - Exclude(closureViolatingTrans, grd → stmt); (9)
    until ((Invnew is unsatisfiable) ∨ (DeadlockStates is unsatisfiable));
    - If (Invnew is satisfiable) then return grd → stmt; (10)
    - declare failure in synthesizing a failsafe program; (11)
  }
}

ExcludeTransitions(transPred: transition predicate, grd → stmt: action)
// exclude the set of transitions represented by transPred
// from the action grd → stmt.
{
  - transPred := grd ∧ Primed(stmtExpr) ∧ ¬transPred;
  - grd := getUnPrimed(transPred); }

```

Figure 4: Synthesis node.

2 of Synthesis\_Node. Steps 3 and 4 do not exclude any additional transitions from  $BG_2$  as the only transitions that directly violate safety upon finalization are the same transitions excluded in Steps 1 and 2. Therefore, the revised action  $BG_2$  is as follows:

$$\begin{aligned}
 BG_2 : & (d.j \neq \perp \wedge f.j = 0) \wedge (d.k = \perp \vee d.j = d.k) \wedge (d.l = \perp \vee d.j = d.l) \\
 & \longrightarrow f.j := true
 \end{aligned}$$

In the case of the BG program, the synthesis nodes associated with the actions of  $P_k$  and  $P_l$  perform revisions similar to that of synthesis nodes associated with  $BG_1$  and  $BG_2$ .

### 5.3 Computing a New Invariant

In this section, we illustrate how our DiConic approach decomposes the problem of calculating a new invariant for the failsafe program by the collaboration of the `Synthesis_Node` and the `Synthesis_Coordinator` algorithms (see Figures 4 and 5). First, the `Synthesis_Coordinator` algorithm ensures that the removal of offending states does not lead to the removal of all invariant states (see Steps 1 and 2 in Figure 5). Then the `Synthesis_Coordinator` enters a loop which synchronizes the activities of all synthesis nodes once they start participating in the calculation of the new invariant (see the loop in Steps 5-9 in Figure 4). These two loops are in fact a DiConic implementation of the loop that includes Steps 4 and 5 in Figure 1. The `Synthesis_Coordinator` waits to receive the revised guards from all synthesis nodes after the removal of offending transitions in Steps 1-4 of the `Synthesis_Node` algorithm. The synthesis nodes enter a waiting state after sending their revised guard to the coordinator. The `Synthesis_Coordinator` calculates the set of all invariant states from where no action is enabled (see Steps 4 and 5 in Figure 5); i.e., deadlock states. Afterwards, the coordinator computes a new invariant by removing the deadlock states. The `Synthesis_Coordinator` sends the new invariant and a state predicate representing deadlock states to all synthesis nodes. If the new invariant becomes empty, then `Synthesis_Coordinator` declares failure in the addition of failsafe fault-tolerance. After receiving the new invariant, each synthesis node determines whether or not its associated action includes transitions that reach a state outside the invariant  $Inv_{new}$  (Steps 8 and 9 in Figure 4); i.e., violates the closure of the invariant. An action violates the closure of the invariant if it contains transitions that reach states that have been excluded by other actions in other synthesis nodes. Afterwards, each synthesis node starts another iteration by sending its revised guard  $grd$  to the coordinator. The collaboration between the synthesis coordinator and synthesis nodes continues until either a valid new invariant is found or the synthesis fails.

**BG example.** Since the invariant  $\mathcal{I}_{BG}$  does not include any offending states (i.e.,  $\mathcal{I}_{BG} \cap UnPrimed(\mathcal{R}_1) = \emptyset$ ), the revision of the action  $BG_2$  does not create new deadlock states in the invariant. Thus, the predicate *DeadlockStates* computed in Step 5 of Figure 5 is equal to the empty set. As a result, the synthesis nodes exit the repeat-until loop in the first iteration.

```

Synthesis_Coordinator(Inv, OS: state predicate)
{ /* OS denotes the set of offending states. */
  /* n denotes the number of program actions. */
  -  $Inv_{new} := Inv \wedge \neg OS$ ; (1)
  - If ( $Inv_{new}$  is unsatisfiable) then
    - declare that a failsafe program cannot be synthesized; exit(); (2)
  repeat {
    - Wait for the revised guards  $grd_1, \dots, grd_n$ 
      from synthesis nodes  $Node_1, \dots, Node_n$ ; (3)
    -  $Guards := grd_1 \vee \dots \vee grd_n$ ; (4)
    -  $DeadlockStates := (Inv_{new} \wedge \neg Guards)$ ; (5)
    -  $Inv_{new} := Inv_{new} \wedge \neg DeadlockStates$ ; (6)
    - Send  $Inv_{new}$  and  $DeadlockStates$  to nodes  $Node_1, \dots, Node_n$ ; (7)
  } until ( $(Inv_{new}$  is unsatisfiable)  $\vee$  ( $DeadlockStates$  is unsatisfiable));
  - If ( $Inv_{new}$  is unsatisfiable) then
    - declare failure in synthesizing a failsafe program; exit(); (8)
  - Notify nodes of successful termination;
}

```

Figure 5: Synthesis coordinator.

Therefore, the invariant of the failsafe program is equal to  $\mathcal{I}_{BG}$ .

**Theorem 5.1** DiConic addition of failsafe fault-tolerance is sound.

**Proof.** Let  $p = \langle V_p, \Pi_p \rangle$  be the input fault-intolerant program with its invariant  $\mathcal{I}$ , its specification  $spec$ , and a fault-type  $F$ . Also, let  $p^f = \langle V_p, \Pi_p^f \rangle$  be the program synthesized by our DiConic method (i.e., algorithms `Synthesis_Coordinator` and `Synthesis_Node`) and  $\mathcal{I}^f$  be the invariant of  $p^f$ . Since the `Synthesis_Coordinator` algorithm only removes states from  $\mathcal{I}$ ,  $\mathcal{I}^f$  is a subset of  $\mathcal{I}$ . Also, by construction, we have  $\mathcal{I}^f \wedge OS \equiv false$ . Likewise, the `Synthesis_Node` algorithm only removes offending and closure-violating transitions from an action. Thus, by construction, no action in  $p^f$  includes new transitions, and all actions of  $p^f$  satisfy the safety of  $spec$ . Therefore,  $p^f$  meets the first and second requirements of the addition problem (see Section 3).

Moreover, if the synthesis is successful, then  $\mathcal{I}^f$  is non-empty and does not have any deadlock states. Since no new transitions are included in the actions of  $p^f$ , starting from  $\mathcal{I}^f$ , the set of computations of  $p^f$  in the absence of faults is a subset of the set of computations of  $p$  starting in  $\mathcal{I}^f$ . Thus, in the absence of faults,  $p^f$  satisfies  $spec$  from  $\mathcal{I}^f$ .

Now, let  $c = \langle s_0, s_1, \dots \rangle$  be a computation of  $p^f$  in the presence of  $F$  in which a transition

$(s_i, s_{i+1})$ , for  $i \geq 0$ , violates the safety of *spec*. If  $(s_i, s_{i+1})$  is a program transition, then there must be an action in  $p^f$  that includes  $(s_i, s_{i+1})$  and violates safety. This is a contradiction with the fact that all offending transitions have been removed from  $p^f$ . If  $(s_i, s_{i+1})$  is a fault transition, then  $s_i$  is a reachable offending state. The transition  $(s_{i-1}, s_i)$  cannot be a program transition because it would be an offending transition. Thus,  $(s_{i-1}, s_i)$  is a fault transition and  $s_{i-1}$  is an offending state. By induction, it follows that  $s_0$  is an offending state. However, since all computations start in the invariant, we have  $s_0 \in \mathcal{I}^f$ , which is a contradiction as all offending states have been removed from  $\mathcal{I}^f$ . Therefore,  $p^f$  is failsafe  $F$ -tolerant from  $\mathcal{I}^f$  for *spec*.  $\square$

**Comment on the completeness of our approach.** The elimination of closure-violating transitions in each synthesis node may create new deadlock states in the invariant. The number of such deadlock states depends on how appropriate closure-violating transitions are selected for elimination among all program actions. The current strategy for eliminating closure-violating transitions may result in the removal of an inappropriate set of such transitions, thereby leading to the failure of the addition of fault-tolerance to the program at hand; i.e., our algorithm may fail to find a failsafe program while there exists one.

## 6 Case Studies

In this section, we present three additional case studies of the application of our DiConic approach in adding failsafe fault-tolerance. Specifically, in Subsection 6.1, we add failsafe fault-tolerance to a simplified version of an altitude switch (ASW) controller. We have chosen this example to compare the failsafe fault-tolerant program synthesized in a divide-and-conquer approach with the ASW program that has been manually designed in [16]. In Subsection 6.2, we illustrate how to add failsafe fault-tolerance in a DiConic way to the controlling program of a Cruise Control System (CCS) (adapted from [24]). Finally, in Subsection 6.3, we demonstrate how DiConic addition of fault-tolerance simplifies the synthesis of a token ring program that tolerates state-corruption faults without violating its safety specification. We note that we have formulated these case studies in terms of the satisfiability problem and have employed the Yices SMT solver to validate our approach.

The Yices specifications of these case studies are available in the Appendix.

## 6.1 Altitude Switch Controller

In this section, we demonstrate the application of our DiConic approach in adding fail-safe fault-tolerance to a simplified version of an altitude switch (ASW) controller program (adapted from [16]).

**The fault-intolerant altitude switch (ASW).** The ASW program monitors a set of input variables and generates an output. Also, the program has a set of internal variables as follows: (i) *AltBelow* is equal to 1 if the altitude is below a specific threshold, otherwise, it is equal to 0; (ii) *ActuatorStatus* is equal to 1 if the actuator is powered on, otherwise, it is equal to 0; (iii) *Inhibit* is equal to 1 when the actuator power-on is inhibited, otherwise, it is equal to 0, and (iv) *Reset* is equal to 1 if the system is being reset.

The ASW program has a mode variable *Mode* and can be in three different modes: (i) the *Initialization* mode when the ASW system is initializing, denoted  $Mode = I$ ; (ii) the *Await-Actuator* mode if the system is waiting for the actuator to power on, denoted  $Mode = AA$ , and (iii) the *Standby* mode, denoted  $Mode = SB$ .

Moreover, we model the signals that come from the input sensors to indicate the occurrence of faults using the following variables: (i) if the system fails in the initialization mode, then the variable *InitFail* will be set to 1, otherwise, *InitFailed* remains 0; (ii) if the altitude sensors fail and do not recover in a certain number of built-in reset attempts, then the variable *AltFail* will be equal to 1, otherwise, *AltFail* remains 0, and (iii) if the Actuator fails in the Await-Actuator mode, then the variable *ActFail* will be equal to 1, otherwise, *ActFail* remains 0. The domain of all variables except *Mode* is equal to  $\{0, 1\}$ . The fault-intolerant program consists of only one process with the following actions:

$$\begin{aligned}
A_1 : (Mode = I) & \longrightarrow Mode := SB; \\
A_2 : (Mode = SB) \wedge (Reset = 1) & \longrightarrow Mode := I; Reset := 0; \\
A_3 : (Mode = SB) \wedge (AltBelow = 1) \wedge (Inhibit = 0) \wedge \\
& (ActuatorStatus = 0) \longrightarrow Mode := AA; \\
A_4 : (Mode = AA) \wedge (ActuatorStatus = 0) \wedge (Inhibit = 0) & \longrightarrow Mode := SB; ActuatorStatus := 1; \\
A_5 : (Mode = AA) \wedge (Reset = 1) & \longrightarrow Mode := I; Reset := 0;
\end{aligned}$$

The program changes its mode from Initialization to Standby. The program goes to the Initialization mode when it is either in Standby or in Await-Actuator mode and the reset signal is received. If the program is in the Standby mode, the altitude is below a pre-determined threshold, the actuator power-on is not inhibited and the actuator is not powered on, then the program changes its mode to Await-Actuator. In the Await-Actuator mode, the program either powers on the actuator and goes to the standby mode, or enters the initialization mode upon receiving the reset signal.

**Read/Write restrictions.** All variables can be read. However, the program cannot write the variables *InitFail*, *AltFail*, *ActFail*, *AltBelow* and *Inhibit*.

**Faults.** If the altitude sensors incur malfunction then the state of the program will be perturbed to a faulty state. We represent the fault actions as follows:

$$\begin{aligned}
F_1 : (InitFail = 0) & \longrightarrow InitFail := 1; \\
F_2 : (AltFail = 0) & \longrightarrow AltFail = 1; \\
F_3 : (ActFail = 0) & \longrightarrow ActFail = 1;
\end{aligned}$$

The guards of the above actions represent conditions under which the program *detects* the occurrence of faults.

**Safety specification.** The problem specification requires that the program does not change its mode from Standby to Await-Actuator if the altitude sensors are failed; i.e., *AltFail* is equal to 1. Moreover, if the initialization has failed, then the ASW program must not enter the Standby mode. Finally, if the actuator fails in the AA mode, then the ASW program should not power on the actuator. Thus, we represent the safety specification of the ASW program by the transition predicate  $\mathcal{S}_{ASW}$ , where

$$\begin{aligned} \mathcal{S}_{ASW} = & ((AltFail = 1) \wedge (Mode = SB) \wedge (Mode' = AA)) \vee \\ & ((InitFail = 1) \wedge (Mode = I) \wedge (Mode' = SB)) \vee \\ & ((ActFail = 1) \wedge (Mode = AA) \wedge (ActuatorStatus = 0) \wedge (ActuatorStatus' = 1)) \end{aligned}$$

This transition predicate specifies the set of transitions that must not appear in any computation of the ASW program even when faults occur.

**Invariant.** If the *Inhibit* is activated, then the actuators should not be powered on. Moreover, the program should not be in a faulty mode. Thus, the invariant of the ASW program is equal to  $\mathcal{I}_{ASW}$ , where

$$\begin{aligned} \mathcal{I}_{ASW} = & ((Inhibit = 1) \Rightarrow (ActuatorStatus = 0)) \wedge \\ & ((InitFail = 0) \wedge (AltFail = 0) \wedge (ActFail = 0)) \end{aligned}$$

**DiConic calculation of the fault-span and offending states.** In DiConic forward reachability on fault actions, in the first iteration, the fault action  $F_1$  reaches the state predicate  $reachedStates_1 \equiv ((Inhibit = 1) \Rightarrow (ActuatorStatus = 0)) \wedge (InitFail = 1)$  from the invariant  $\mathcal{I}_{ASW}$ . Likewise, the reachability nodes corresponding to  $F_2$  and  $F_3$  respectively send the following set of reached states to the reachability coordinator:  $reachedStates_2 \equiv ((Inhibit = 1) \Rightarrow (ActuatorStatus = 0)) \wedge (AltFail = 1)$  and  $reachedStates_3 \equiv ((Inhibit = 1) \Rightarrow (ActuatorStatus = 0)) \wedge (ActFail = 1)$ . Thus, since program actions are closed in the invariant, the set of states reachable by fault and program actions in the first iteration of the repeat-until loop in Figure 3 is equal to  $\mathcal{FS}_{ASW} = ((Inhibit = 1) \Rightarrow (ActuatorStatus = 0)) \wedge ((AltFail = 1) \vee (InitFail = 1) \vee (ActFail = 1))$ . The set of states reachable from  $\mathcal{FS}_{ASW}$  by fault and program actions would yield no new states outside  $\mathcal{FS}_{ASW}$ , thereby returning  $\mathcal{FS}_{ASW}$  as the fault-span of the ASW program for fault actions  $F_1, F_2$  and  $F_3$ . Since fault actions do not directly violate safety specification, the base set of states for backward reachability is empty. Therefore, the set of offending states is empty.

**Removing the set of offending transitions.** We create five synthesis nodes corresponding to the actions  $A_1$ - $A_5$ . In addition to an action  $A_i$  ( $1 \leq i \leq 5$ ), each synthesis node  $Node_i$  takes the invariant  $\mathcal{I}_{ASW}$ , the fault-span  $\mathcal{FS}_{ASW}$ , the safety specification  $\mathcal{S}_{ASW}$ , the set of offending states  $OS$  and the write restrictions as its inputs. Since  $OS = \emptyset$ , the first two steps



in Figure 4 result in an unsatisfiable transition predicate. Next, we describe the results of executing Steps 3 and 4 in each  $Node_i$ .

- **Action**  $A_1$ . Step 3 calculates the transition predicate  $(Mode = I) \wedge (InitFail = 1) \wedge (Mode' = SB)$  as the set of transitions that violate the safety specification and should be excluded from action  $A_1$ . Therefore, action  $A_1$  would be revised as follows:

$$(Mode = I) \wedge (\mathbf{InitFail} = \mathbf{0}) \longrightarrow Mode := SB;$$

- **Action**  $A_2$ . Since the safety specification  $\mathcal{S}_{ASW}$  does not stipulate anything about changing the program mode to initialization, this action does not include any offending transitions. Therefore,  $A_2$  remains unchanged.
- **Action**  $A_3$ . The safety specification  $\mathcal{S}_{ASW}$  states that the ASW program must not enter the AA mode if the altitude sensors have failed (i.e.,  $AltFail = 1$ ). Thus, the synthesis node corresponding to action  $A_3$  revises  $A_3$  as follows:

$$(Mode = SB) \wedge (AltBelow = 1) \wedge (Inhibit = 0) \wedge (ActuatorStatus = 0) \wedge (\mathbf{AltFail} = \mathbf{0}) \longrightarrow Mode := AA;$$

- **Action**  $A_4$ . If faults cause the actuator to fail, then the program should not power on the actuator. Thus, due to the execution of the fault action  $F_3$ , action  $A_4$  may be enabled in states where  $ActFail = 1$  holds, thereby leading to the violation of safety. The elimination of such transitions originated in the fault-span outside the invariant results in the following revised action:

$$(Mode = AA) \wedge (ActuatorStatus = 0) \wedge (Inhibit = 0) \wedge (\mathbf{ActFail} = \mathbf{0}) \longrightarrow Mode := SB; ActuatorStatus := 1;$$

- **Action**  $A_5$ . This action remains unchanged for the same reason as that of action  $A_2$ .

Since the set of offending states is empty, no state is removed from the invariant. Therefore, the invariant of the failsafe program is equal to  $\mathcal{I}_{ASW}$ . We note that the failsafe ASW program synthesized in this section is similar to the ASW program that has been manually designed in [16].

## 6.2 Cruise Control System

In this section, we present a case study on the application of our DiConic approach in adding failsafe fault-tolerance to the controlling program of a Cruise Control System (CCS). The cruise control example in this section is a simplified version of the example in [24].

**The fault-intolerant CCS program.** The CCS program has 4 input variables *Ignition*, *EngState*, *Brake*, and *Lever* that represent the values of the input signals. (To distinguish variables from their values, system variables start with capitalized letters.) The domain of the variable *Ignition* contains values of *on* and *off* that represent the state of the ignition switch. The variable *EngState* represents the working state of the engine with the domain  $\{running, off\}$ . The variable *Brake* illustrates the state of an input signal from the brakes that shows whether or not the driver has applied the brakes. The domain of *Brake* is equal to  $\{notApplied, applied, unknown\}$ . The *Lever* variable represents the position of the cruise control lever set by the driver. The cruise control lever can be in three positions *off*, *constant*, and *resume*. The CCS program also has a variable *SysMode* that stores its operating mode. The CCS program can be in the following modes: *off*, *inactive*, *cruise*, and *override*. If the CCS program is in none of the above-mentioned modes then its mode is *unknown*. We represent the fault-intolerant CCS program by the actions  $A_1$ - $A_9$ .

$$\begin{aligned}
A_1 &: ((SysMode = off) \wedge (Ignition = on)) &\longrightarrow& SysMode := inactive; \\
A_2 &: ((SysMode = inactive) \wedge (Ignition = off)) &\longrightarrow& SysMode := off; \\
A_3 &: ((SysMode = inactive) \wedge (Lever = constant) \wedge (Ignition = on) \wedge \\
&\quad (EngState = running)) &\longrightarrow& SysMode := cruise; \\
A_4 &: ((SysMode = cruise) \wedge (Ignition = off)) &\longrightarrow& SysMode := off; \\
A_5 &: ((SysMode = cruise) \wedge (EngState = off)) &\longrightarrow& SysMode := inactive; \\
A_6 &: ((SysMode = cruise) \wedge ((Brake = applied) \vee (Lever = off))) &\longrightarrow& SysMode := override; \\
A_7 &: ((SysMode = override) \wedge (Ignition = off)) &\longrightarrow& SysMode := off; \\
A_8 &: ((SysMode = override) \wedge (EngState = off)) &\longrightarrow& SysMode := inactive; \\
A_9 &: ((SysMode = override) \wedge (Ignition = on) \wedge (EngState = running) \wedge \\
&\quad ((Lever = constant) \vee (Lever = resume))) &\longrightarrow& SysMode := cruise;
\end{aligned}$$

If the program is in the *off* mode and the ignition is on then the program transitions to the *inactive* mode (Action  $A_1$ ). The program transitions to the *off* mode if it is in the *inactive*

mode and the ignition turns off (Action  $A_2$ ). In the *inactive* mode, if (i) the lever is set on *constant*, (ii) the ignition is on, and (iii) the engine is running, then the program transitions to the *Cruise* mode (Action  $A_3$ ). In the *cruise* mode, the program transitions to (i) the *off* mode if the ignition turns off (Action  $A_4$ ), or (ii) the *inactive* mode when the engine turns off (Action  $A_5$ ). If the lever is off or the driver has applied the brakes, then the program moves to the *override* mode from *cruise* mode (Actions  $A_6$ ). In the *override* mode, the program goes to the *off* mode if the ignition turns off (Action  $A_7$ ). If the engine turns off then the program will transition to the *inactive* mode from the *override* mode (Action  $A_8$ ). Finally, in the *override* mode, the program transitions to the *cruise* mode if (i) the ignition is on, (ii) the engine is running, and (iii) the lever is on *constant* or *resume* (Action  $A_9$ ).

**The invariant of the CCS program.** In general, the CCS program should be in one of the modes *off*, *inactive*, *cruise*, or *override*, and the brakes subsystem should also be working properly. Also, if the program is in the *off* mode then the ignition should be off. If the program is in the *inactive* mode then the ignition should be on. If the program is in the *cruise* mode, then the ignition should be on, the engine should be running, the lever should not be off, and the brakes must not be applied. In the *override* mode, the ignition is on and the engine should be running. Hence, we represent the invariant of the CCS program by the state predicate  $\mathcal{I}_{CCS}$ , where

$$\mathcal{I}_{CCS} = ((SysMode \neq unknown) \wedge (Brake \neq unknown)) \wedge I_1, \text{ where}$$

$$\begin{aligned} \mathcal{I}_1 = & ((SysMode = off) \vee (SysMode = inactive) \vee \\ & (SysMode = cruise) \vee (SysMode = override)) \wedge \\ & ((SysMode = off) \Rightarrow (Ignition = off)) \wedge \\ & ((SysMode = inactive) \Rightarrow (Ignition = on)) \wedge \\ & ((SysMode = cruise) \Rightarrow ((Ignition = on) \wedge (EngState = running) \wedge \\ & (Brake \neq applied) \wedge (Lever \neq off))) \wedge \\ & ((SysMode = override) \Rightarrow ((Ignition = on) \wedge (EngState = running))) \end{aligned}$$

**Read/Write restrictions.** The CCS program can read all variables, however, it cannot write the variables *Ignition*, *EngState*, *Brake*, and *Lever* as they represent the values of the input signals.

**Fault actions.** The malfunction of the brake subsystem may corrupt the value of *Brake* to an unknown value. In addition, faults may perturb the cruise control system to an unknown operating mode. We represent the set of fault transitions by the following actions:

$$\begin{aligned} F_1 : (Brake \neq unknown) &\longrightarrow Brake := unknown; \\ F_2 : (SysMode \neq unknown) &\longrightarrow SysMode := unknown; \end{aligned}$$

**The safety specification of the CCS program.** The specification of the CCS program stipulates that as long as the driver has applied the brakes the CCS program must never transition to the cruise mode. Also, in cases where the signal coming from the brake subsystem is corrupted then it is not safe for the program to transition to the cruise mode. Finally, the program must not transition to the cruise mode in cases where it is in an unknown mode. Hence, we represent the safety specification of the CCS program by the following transition predicate:

$$\begin{aligned} \mathcal{S}_{CCS} = & ((Brake = applied) \wedge (Brake' = applied) \wedge (SysMode' = cruise)) \vee \\ & ((Brake = unknown) \wedge (SysMode' = cruise)) \vee \\ & ((SysMode = unknown) \wedge (SysMode' = cruise)) \end{aligned}$$

**DiConic calculation of the fault-span and offending states.** We instantiate two forward reachability nodes corresponding to the fault actions  $F_1$  and  $F_2$  and nine forward reachability nodes corresponding to program actions  $A_1$ - $A_9$ . The state predicates computed by the reachability nodes of  $F_1$  and  $F_2$  are as follows:

- **Action  $F_1$ .** This action may set the *Brake* to an unknown value, thereby falsifying the state predicate  $(SysMode \neq unknown) \wedge (Brake \neq unknown)$ . However,  $F_1$  does not violate the closure of  $\mathcal{I}_1$ . Therefore, the program may be perturbed to  $\mathcal{I}_1$  by  $F_1$ .
- **Action  $F_2$ .** The execution of the fault action  $F_2$  violates the closure of both state predicates  $\mathcal{I}_1$  and  $(SysMode \neq unknown) \wedge (Brake \neq unknown)$ .

Since the invariant is closed in program actions, in the first round of the forward reachability the set of reachable states would be equal to  $\mathcal{FS}_{CCS} = \mathcal{I}_1 \vee (SysMode = unknown) \vee (Brake = unknown)$ . Subsequent iterations of forward reachability would yield the same predicate  $\mathcal{FS}_{CCS}$  as the reachable state predicate from  $\mathcal{FS}_{CCS}$  by fault and program actions.

This is because if only  $(SysMode = unknown)$  holds, then no program actions is enabled, thereby no new states are reached. If only  $(Brake = unknown)$  holds, then  $\mathcal{I}_1$  is the set of reachable states from  $\mathcal{FS}_{CCS}$ . For this reason, the fault-span of the CCS program is  $\mathcal{FS}_{CCS}$ . Since fault actions do not directly violate the safety of specification, the base set of states for backward reachability is empty. Therefore, the set of offending states is empty.

**Removing the set of offending transitions.** We create nine synthesis nodes corresponding to the actions  $A_1$ - $A_9$ . Since the only actions that cause the CCS program to transition to the cruise mode are actions  $A_3$  and  $A_9$ , the only actions that may directly violate safety would be  $A_3$  and  $A_9$ . Thus, only the corresponding synthesis nodes of these actions perform revisions and the rest of the actions remain unchanged in the failsafe program. We present the revised actions  $A_3$  and  $A_9$  as follows:

- **Action  $A_3$ .** Step 3 of the corresponding synthesis node calculates the transition predicate  $((Brake = applied) \wedge (Brake' = applied)) \vee (Brake = unknown) \wedge (SysMode' = cruise)$  as the set of transitions that violate the safety specification and should be excluded from action  $A_3$ . Therefore, action  $A_3$  would be revised as follows:

$$A_3 : ((SysMode = inactive) \wedge (Lever = constant) \wedge (Ignition = on) \wedge (EngState = running) \wedge (\mathbf{Brake} = \mathbf{notApplied})) \longrightarrow SysMode := cruise;$$

- **Action  $A_9$ .** For a similar reason, the synthesis node corresponding to the action  $A_9$  returns the following revised action:

$$A_9 : ((SysMode = override) \wedge (Ignition = on) \wedge (EngState = running) \wedge (\mathbf{Brake} = \mathbf{notApplied}) \wedge ((Lever = constant) \vee (Lever = resume))) \longrightarrow SysMode := cruise;$$

Since the set of offending states is empty, no state is removed from the invariant. Therefore, the invariant of the failsafe program is equal to  $\mathcal{I}_{CCS}$ . We would like to emphasize that our DiConic approach significantly simplifies the addition of failsafe fault-tolerance by focusing on each of the above nine actions separately instead of adding fault-tolerance to all of them at once.

### 6.3 Token Ring

In this section, we synthesize a token ring program that is failsafe fault-tolerant to process restart faults. The fault-intolerant Token Ring (TR) program consists of four processes  $P_0, P_1, P_2$ , and  $P_3$  arranged in a ring. Each process  $P_i$ ,  $0 \leq i \leq 3$ , has a variable  $x_i$  with the domain  $\{-1, 0, 1\}$ . We say that process  $P_i$ ,  $1 \leq i \leq 3$ , has the token if and only if  $(x_i \neq x_{i-1})$  and fault transitions have not corrupted  $P_i$  and  $P_{i-1}$ . And,  $P_0$  has the token if  $(x_3 = x_0)$  and fault transitions have not corrupted  $P_0$  and  $P_3$ . Process  $P_i$ ,  $1 \leq i \leq 3$ , copies  $x_{i-1}$  to  $x_i$  if the value of  $x_i$  is different from  $x_{i-1}$ . This action passes the token to the next process. Also, if  $(x_0 = x_3)$  holds then process  $P_0$  copies the value of  $(x_3 + 1) \bmod 2$  to  $x_0$ . Thus, if we initialize every  $x_i$ ,  $0 \leq i \leq 3$ , with 0, then process  $P_0$  has the token and the token circulates along the ring. We specify the action of  $P_0$  as follows.

$$A_0 : (x_0 = x_3) \quad \longrightarrow \quad x_0 := (x_3 + 1) \bmod 2;$$

Since processes  $P_1, P_2$ , and  $P_3$  are similar, we present their actions in a parameterized format, where  $1 \leq i \leq 3$ .

$$A_i : (x_i \neq x_{i-1}) \quad \longrightarrow \quad x_i := x_{i-1};$$

**Read/Write restrictions.** Each process  $P_i$ ,  $1 \leq i \leq 3$ , is only allowed to read  $x_{i-1}$  and  $x_i$ , and allowed to write  $x_i$ . Process  $P_0$  is allowed to read  $x_3$  and  $x_0$ , and write  $x_0$ .

**State-corruption faults.** The faults may corrupt the value of a process if that process and its predecessor (i.e., its locality) are not already corrupted. We model the effect of faults on a process  $P_i$  by setting  $x_i = -1$ . Thus, the fault actions are as follows:

$$\begin{aligned} F_0 : (x_0 \neq -1) \wedge (x_3 \neq -1) & \longrightarrow x_0 := -1; \\ F_1 : (x_1 \neq -1) \wedge (x_0 \neq -1) & \longrightarrow x_1 := -1; \\ F_2 : (x_2 \neq -1) \wedge (x_1 \neq -1) & \longrightarrow x_2 := -1; \\ F_3 : (x_3 \neq -1) \wedge (x_2 \neq -1) & \longrightarrow x_3 := -1; \end{aligned}$$

Note that there exist no read/write restrictions for the fault actions because we assume that fault actions can read and write arbitrary program variables.

**Safety specification.** Based on the problem specification, the fault-tolerant program is not allowed to take a transition where a non-corrupted process copies a corrupted value from its predecessor.

$$\begin{aligned} \mathcal{S}_{TR} : & ((x_0 = -1) \wedge (x_1 \neq -1) \wedge (x'_1 = -1)) \vee \\ & ((x_1 = -1) \wedge (x_2 \neq -1) \wedge (x'_2 = -1)) \vee \\ & ((x_2 = -1) \wedge (x_3 \neq -1) \wedge (x'_3 = -1)) \vee \\ & ((x_3 = -1) \wedge (x_0 \neq x'_0)) \end{aligned}$$

**Invariant.** The invariant of the program consists of the states where no process is corrupted and there exists only one token in the ring. The state predicate  $\mathcal{I}_{TR}$  represents the invariant of the token ring program, where

$$\begin{aligned} \mathcal{I}_{TR} : & ((x_0 = 0) \wedge (x_1 = 0) \wedge (x_2 = 0) \wedge (x_2 = 0)) \vee \\ & ((x_0 = 1) \wedge (x_1 = 0) \wedge (x_2 = 0) \wedge (x_2 = 0)) \vee \\ & ((x_0 = 1) \wedge (x_1 = 1) \wedge (x_2 = 0) \wedge (x_2 = 0)) \vee \\ & ((x_0 = 1) \wedge (x_1 = 1) \wedge (x_2 = 1) \wedge (x_2 = 0)) \vee \\ & ((x_0 = 1) \wedge (x_1 = 1) \wedge (x_2 = 1) \wedge (x_2 = 1)) \vee \\ & ((x_0 = 0) \wedge (x_1 = 1) \wedge (x_2 = 1) \wedge (x_2 = 1)) \vee \\ & ((x_0 = 0) \wedge (x_1 = 0) \wedge (x_2 = 1) \wedge (x_2 = 1)) \vee \\ & ((x_0 = 0) \wedge (x_1 = 0) \wedge (x_2 = 0) \wedge (x_2 = 1)) \vee \end{aligned}$$

**DiConic calculation of the fault-span and offending states.** In DiConic forward reachability on fault actions, in the first iteration, the fault action  $F_0$  reaches the state predicate  $RS_0$  from the invariant  $\mathcal{I}_{TR}$ , where  $RS_0 \equiv ((x'_0 = -1) \wedge (((x'_1 = 0) \wedge (x'_2 = 0)) \vee ((x'_2 = 1) \wedge (x'_3 = 1)) \vee ((x'_1 = 1) \wedge (x'_3 = 0))))$  from the invariant  $\mathcal{I}_{TR}$ . Likewise, the reachability nodes corresponding to  $F_1, F_2$  and  $F_3$  respectively send the state predicate  $RS_1, RS_2,$  and  $RS_3$  to the reachability coordinator, where

$$\begin{aligned} RS_1 & \equiv (x'_1 = -1) \wedge \\ & ((x'_0 = 0) \wedge (x'_2 = 0)) \vee ((x'_2 = 1) \wedge (x'_3 = 1)) \vee ((x'_0 = 1) \wedge (x'_3 = 0)), \\ RS_2 & \equiv (x'_2 = -1) \wedge \\ & (((x'_0 = 0) \wedge (x'_1 = 0)) \vee ((x'_1 = 1) \wedge (x'_3 = 1)) \vee ((x'_0 = 1) \wedge (x'_3 = 0))), \text{ and} \\ RS_3 & \equiv (x'_3 = -1) \wedge \\ & (((x'_0 = 0) \wedge (x'_1 = 0)) \vee ((x'_1 = 1) \wedge (x'_2 = 1)) \vee ((x'_0 = 1) \wedge (x'_2 = 0))). \end{aligned}$$

At the end of the first round of the forward reachability computation, the reachability coordinator computes the state predicate *reachedStates* (see Figure 3), which is equal to

$RS_0 \vee RS_1 \vee RS_2 \vee RS_3$ . Note that, in the first round of forward reachability, the execution of program actions from the invariant  $\mathcal{I}_{TR}$  does not yield states outside the invariant. In the second round, the actions of the fault-intolerant program that start in a state in *reachedStates* may propagate the corrupted value and result in states in which more than one process is corrupted. Thus, all possible states in the state space may be reached due to the occurrence of faults and execution of program actions outside the invariant. Therefore, the fault-span of the token ring program for fault actions  $F_0$ - $F_3$  from the invariant  $\mathcal{I}_{TR}$  is equal to true (i.e., the universal set). Since in this case fault actions do not directly violate safety specification, the base set of states for backward reachability is empty. Therefore, the set of offending states is empty.

**Removing the set of offending transitions.** We create four synthesis nodes corresponding to the actions of  $A_0$ - $A_3$ . Since the set of offending states is empty, Steps 1 and 2 in *Synthesis\_Node* (see Figure 4) do not result in the exclusion of any transitions. Nonetheless, program actions violate the safety of specification when they propagate a corrupted value outside the invariant. The synthesis nodes eliminate such transitions in Steps 3 and 4, thereby revising program actions as follows:

- **Action  $A_0$ .** Step 3 calculates the transition predicate  $(x_3 = -1) \wedge (x_0 = x_3) \wedge (x'_0 = 0)$  as the set of transitions that violate the safety specification and should be excluded from action  $A_0$ . Therefore, the revised action  $A_0$  is as follows:

$$(x_0 = x_3) \wedge (\mathbf{x}_3 \neq -\mathbf{1}) \quad \longrightarrow \quad x_0 := (x_3 + 1) \text{ mod } 2;$$

- **Actions  $A_1$ - $A_3$ .** We present the revised actions  $A_1$ - $A_3$  as follows, for  $1 \leq i \leq 3$ :

$$(x_i \neq x_{i-1}) \wedge (\mathbf{x}_{i-1} \neq -\mathbf{1}) \quad \longrightarrow \quad x_i := x_{i-1};$$

Since the set of offending states is empty, no state is removed from the invariant. Therefore, the invariant of the failsafe program is equal to  $\mathcal{I}_{TR}$ .



## 7 Related Work

In this section, we discuss related work on existing divide-and-conquer synthesis methods [25, 26] and techniques for reducing the time/space complexity of automatic addition of fault-tolerance [11, 13, 17, 23, 27]. Specifically, Smith [25] presents a divide-and-conquer approach for synthesizing programs from their algebraic specifications, where he decomposes the program specification into the specifications of sub-problems and combines the results of synthesizing solutions for sub-problems. Puri and Gu [26] propose a divide-and-conquer synthesis method for asynchronous digital circuits, where they decompose circuit specifications and satisfy design constraints for each sub-specification. While the aforementioned approaches inspire our work, they are essentially specification-based approaches and do not directly focus on adding failsafe fault-tolerance to existing programs.

Previous work on automatic addition of fault-tolerance [11, 13, 17, 23, 27] mostly focuses on techniques for reducing time/space complexity of synthesis. For example, Kulkarni *et al.* [17] present a set of heuristics based on which they reduce the time complexity of adding fault-tolerance to integrated models of distributed programs. Kulkarni and Ebneenasir [11] present a technique for reusing the computations of an existing fault-tolerant program in order to enhance its level of tolerance. They also present a set of pre-synthesized fault-tolerance components [27] that can be reused during the addition of fault-tolerance to different programs. We have presented a SAT-based technique [23] where we employ SAT solvers to solve some verification problems during the addition of fault-tolerance to integrated program models. Bonakdarpour and Kulkarni [13] present a symbolic implementation of the heuristics in [17] where they use BDDs to model distributed programs.

Our DiConic approach significantly differs from the previous work in that we present a new paradigm for decomposing the problem of adding failsafe fault-tolerance instead of adding failsafe fault-tolerance to an integrated combinatorial model of the entire program. More specifically, our DiConic approach is orthogonal to time/space complexity issues since one can benefit from existing techniques to mitigate the time/space complexity of revising individual program actions. While we have used the Yices SMT solver to implement our

DiConic approach, the distributed nature of our algorithm enables us to implement different techniques and data structures for the revision of different program actions depending on their granularity (e.g., single variable updates, nested iterative instructions, etc.).

## 8 Conclusions and Future Work

We presented a divide-and-conquer approach, called DiConic, for the addition of failsafe fault-tolerance to (distributed) programs, where in the presence of faults a failsafe program guarantees to satisfy at least its safety specification. We specifically focused on the following question: *Given a fault-intolerant program and a specific fault-type, how can we revise each program action separately so that the entire program becomes failsafe fault-tolerant?* To address this question, we decomposed the problem of adding failsafe fault-tolerance into sub-problems. In each sub-problem, we concentrated on one program action and determined how that action should be revised so that the entire program would be failsafe fault-tolerant. We validated our approach for the classic problem of Byzantine Generals problem and for a simplified version of an altitude switch controller.

As an extension of this work, we are investigating the DiConic addition of recovery to programs, where we revise program instructions in isolation so that the entire program would eventually recover to its invariant after faults stop occurring. Another extension of this work is to enhance the efficiency of the synthesis coordinator by developing its distributed version. DiConic addition of fault-tolerance is a special case of a more general problem in which new Temporal Logic [28] properties are incrementally added to an existing program (as defined and addressed in our previous work [29]). Towards this end, we plan to develop new DiConic algorithms for automatic addition of new properties to (concurrent and distributed) programs.

## References

- [1] Sandeep S. Kulkarni and Ali Ebnesnasir. Complexity issues in automated synthesis of failsafe fault-tolerance. *IEEE Transactions on Dependable and Secure Computing*, 2(3):201–215, 2005.
- [2] E. W. Dijkstra. *A Discipline of Programming*. Prentice Hall, 1976.
- [3] E. A. Emerson and E. M. Clarke. Using branching time temporal logic to synchronize synchronization skeletons. *Science of Computer Programming*, 2:241–266, 1982.
- [4] A. Pnueli and R. Rosner. On the synthesis of a reactive module. In *ACM Symposium on Principles of Programming Languages*, pages 179–190, Austin, Texas, 1989.
- [5] S. Lafortune and F. Lin. On tolerable and desirable behaviors in supervisory control of discrete event systems. *Discrete Event Dynamic Systems: Theory and Applications*, 1(1):61–92, 1992.
- [6] Wolfgang Thomas. On the synthesis of strategies in infinite games. In *STACS*, pages 1–13, 1995.
- [7] K.H. Cho and J.T. Lim. Synthesis of fault-tolerant supervisor for automated manufacturing systems: A case study on photolithography process. *IEEE Transactions on Robotics and Automation*, 14(2):348–351, April 1998.
- [8] N. Wallmeier, P. Hütten, and Wolfgang Thomas. Symbolic synthesis of finite-state controllers for request-response specifications. In *CIAA, LNCS, Vol. 2759*, pages 11–22, 2003.
- [9] P. C. Attie, A. Arora, and E. A. Emerson. Synthesis of fault-tolerant concurrent programs. *ACM Transactions on Programming Languages and Systems (TOPLAS)*. (A preliminary version of this paper appeared in *Proceedings of PODC '98.*), 26(1):125 – 185, 2004.

- [10] S. S. Kulkarni and A. Arora. Automating the addition of fault-tolerance. In *In Proceedings of the 6th International Symposium on Formal Techniques in Real-Time and Fault-Tolerant Systems*, pages 82–93, 2000.
- [11] Sandeep S. Kulkarni and A. Ebzenasir. Enhancing the fault-tolerance of nonmasking programs. In *Proceedings of the 23rd International Conference on Distributed Computing Systems*, pages 441–449, 2003.
- [12] Ali Ebzenasir and Sandeep S. Kulkarni. FTSyn: A framework for automatic synthesis of fault-tolerance. <http://www.cs.mtu.edu/~aebzenas/research/tools/ftsyn.htm>.
- [13] B. Bonakdarpour and Sandeep S. Kulkarni. Exploiting symbolic techniques in automated synthesis of distributed programs. In *IEEE International Conference on Distributed Computing Systems (to appear)*, 2007.
- [14] Bruno Dutertre and Leonardo de Moura. A Fast Linear-Arithmetic Solver for DPLL(T). In *Proceedings of the 18th Computer-Aided Verification conference*, volume 4144 of *LNCS*, pages 81–94. Springer-Verlag, 2006.
- [15] Leslie Lamport, Robert Shostak, and Marshall Pease. The byzantine generals problem. *ACM Transactions on Programming Languages and Systems*, 4(3):382–401, 1982.
- [16] R. Bharadwaj and C. Heitmeyer. Developing high assurance avionics systems with the SCR requirements method. In *Proceedings of the 19th Digital Avionics Systems Conference, Philadelphia, PA*, October 2000.
- [17] S. S. Kulkarni, A. Arora, and A. Chippada. Polynomial time synthesis of Byzantine agreement. In *Symposium on Reliable Distributed Systems*, pages 130 – 139, 2001.
- [18] B. Alpern and F. B. Schneider. Defining liveness. *Information Processing Letters*, 21:181–185, 1985.
- [19] A. Arora and M. G. Gouda. Closure and convergence: A foundation of fault-tolerant computing. *IEEE Transactions on Software Engineering*, 19(11):1015–1027, 1993.

- [20] A. Arora and Sandeep S. Kulkarni. Detectors and correctors: A theory of fault-tolerance components. In *International Conference on Distributed Computing Systems*, pages 436–443, May 1998.
- [21] P. Attie and A. Emerson. Synthesis of concurrent programs for an atomic read/write model of computation. *ACM TOPLAS (a preliminary version appeared in PODC '96)*, 23(2), March 2001.
- [22] Ali Ebneenasir. *Automatic Synthesis of Fault-Tolerance*. PhD thesis, Michigan State University, May 2005.
- [23] Ali Ebneenasir and Sandeep S. Kulkarni. SAT-based synthesis of fault-tolerance. *Fast Abstracts of the International Conference on Dependable Systems and Networks, Palazzo dei Congressi, Florence, Italy, June 28 - July 1, 2004*.
- [24] R. Jeffords and C. Heitmeyer. Automatic generation of state invariants from requirements specifications. *Proceedings of the 6th ACM SIGSOFT international symposium on Foundations of software engineering*, pages 56 – 69, 1998.
- [25] D.R. Smith. A problem reduction approach to program synthesis. *Proceedings of the Eighth International Joint Conference on Artificial Intelligence*, pages 32–36, 1983.
- [26] R. Puri and J. Gu. A divide-and-conquer approach for asynchronous interface synthesis. In *ISSS '94: Proceedings of the 7th international symposium on High-level synthesis*, pages 118–125, 1994.
- [27] Sandeep S. Kulkarni and Ali Ebneenasir. Adding fault-tolerance using pre-synthesized components. *Fifth European Dependable Computing Conference (EDCC-5), LNCS, Vol. 3463, p. 72*, 2005.
- [28] E. A. Emerson. *Handbook of Theoretical Computer Science*, volume B, chapter 16: Temporal and Modal Logics, pages 995–1067. Elsevier Science Publishers B. V., 1990.
- [29] A. Ebneenasir, S. S. Kulkarni, and B. Bonakdarpour. Revising UNITY programs: Possibilities and limitations. In *International Conference on Principles of Distributed Systems (OPODIS)*, pages 275–290, 2005.

## Appendix A: The Yices Specifications

In this section, we illustrate how we employ the Yices SMT solver to apply the algorithms `Synthesis_Node` and `Synthesis_Coordinator` (see Figures 4 and 5) to each case study presented in Section 6. Specifically, for each case study, we define a logical context that specifies a static structure for the fault-intolerant program based on the program and fault actions. The synthesis coordinator distributes a copy of this logical context to each synthesis node. A synthesis node uses this context to interact with an instance of the Yices SMT solver (running locally) towards revising its associated action. In Subsection A.1, we present the Yices specification of the Byzantine Generals problem and the assertions used to verify the conditions in algorithms `Synthesis_Node` and `Synthesis_Coordinator`. Then, in Subsection A.2, we demonstrate how to add failsafe fault-tolerance to the altitude switch program. Subsequently, in Subsection A.3, we present the Yices specification of the cruise control system. Finally, in Subsection A.4, we discuss the specification of the token ring program.

### Appendix A.1: Byzantine Generals Problem

In this section, we present the logical context in which we specified the Byzantine Generals problem.

```
1 (define-type decision (scalar undecided attack retreat))
2 ;; defining the unprimed variables
3 (define bg::bool)
4 (define dg:: (subtype (gendec::decision) (or (= gendec attack) (= gendec retreat))))
5 (define b1::bool)
6 (define d1::decision)
7 (define f1::bool)
8 (define b2::bool)
9 (define d2::decision)
10 (define f2::bool)
11 (define b3::bool)
12 (define d3::decision)
13 (define f3::bool)
14
15 ;; defining the primed variables
16 (define bgp::bool)
17 (define dgp:: (subtype (gendec::decision) (or (= gendec attack) (= gendec retreat)) ))
18 (define b1p::bool)
19 (define d1p::decision)
20 (define f1p::bool)
```

```

21(define b2p::bool)
22(define d2p::decision)
23(define f2p::bool)
24(define b3p::bool)
25(define d3p::decision)
26(define f3p::bool)
27
28;; defining an invariant
29(define Inv1::bool ( and
30(and (not bg) (or (not b1)(not b2)) (or (not b1)(not b3)) (or (not b2)(not b3)) )
31( => (not b1) (or (= d1 undecided) (= d1 dg) ) )
32( => (not b2) (or (= d2 undecided) (= d2 dg) ) )
33( => (not b3) (or (= d3 undecided) (= d3 dg) ) )
34(=> (and (not b1) (= f1 true)) (/= d1 undecided) )
35(=> (and (not b2) (= f2 true)) (/= d2 undecided) )
36(=> (and (not b3) (= f3 true)) (/= d3 undecided) ) ))
37
38(define Inv2::bool
39(and (= bg true) (not b1) (not b2) (not b3) (= d1 d2) (= d1 d3)
40(/= d1 undecided)))
41
42(define Inv::bool (or Inv1 Inv2))
43
44;; defining a primed version of the invariant
45(define Inv1p::bool ( and
46(and (not bgp) (or (not b1p)(not b2p)) (or (not b1p)(not b3p)) (or (not b2p)(not b3p)) )
47( => (not b1p) (or (= d1p undecided) (= d1p dgp) ) )
48( => (not b2p) (or (= d2p undecided) (= d2p dgp) ) )
49( => (not b3p) (or (= d3p undecided) (= d3p dgp) ) )
50(=> (and (not b1p) (= f1p true)) (/= d1p undecided) )
51(=> (and (not b2p) (= f2p true)) (/= d2p undecided) )
52(=> (and (not b3p) (= f3p true)) (/= d3p undecided) ) ))
53
54(define Inv2p::bool
55(and (= bgp true) (not b1p) (not b2p) (not b3p) (= d1p d2p)
56(= d1p d3p) (/= d1p undecided)))
57
58(define Invp::bool (or Inv1p Inv2p))
59
60;; defining the safety of specification
61(define safety::bool (or
62(and (not b1p) (not b2p) (/= d1p undecided) (/= d2p undecided) (= f1p true)
63(= f2p true) (/= d1p d2p) )
64(and (not b1p) (not b3p) (/= d1p undecided) (/= d3p undecided) (= f1p true)
65(= f3p true) (/= d1p d3p) )
66(and (not b3p) (not b2p) (/= d3p undecided) (/= d2p undecided) (= f3p true)
67(= f2p true) (/= d3p d2p) )
68(and (not bgp) (not b1p) (/= d1p undecided) (= f1p true) (/= d1p dgp))
69(and (not bgp) (not b2p) (/= d2p undecided) (= f2p true) (/= d2p dgp))
70(and (not bgp) (not b3p) (/= d3p undecided) (= f3p true) (/= d3p dgp))
71(and (not b1) (not b1p) (= f1 true) (or (/= d1 d1p) (/= f1 f1p)) )
72(and (not b2) (not b2p) (= f2 true) (or (/= d2 d2p) (/= f2 f2p)) )
73(and (not b3) (not b3p) (= f3 true) (or (/= d3 d3p) (/= f3 f3p)) )

```

```

74)    )
75
76;; defining fault actions as transition predicates
77(define faultAC1::bool (and
78((not bg) (not b1) (not b2) (not b3))
79(or
80(and (= bgp true) (not b1p) (not b2p) (not b3p))
81(and (not bgp) (= b1p true) (not b2p) (not b3p))
82(and (not bgp) (not b1p) (= b2p true) (not b3p))
83(and (not bgp) (not b1p) (not b2p) (= b3p true)) )
84))
85
86(define faultAC2::bool (or
87(and (= b1 true) (= d1 attack) (= d1p retreat) (= f1 false) (= f1p true) )
88(and (= b1 true) (= d1p attack) (= d1 retreat) (= f1 false) (= f1p true) )
89(and (= b1 true) (= d1 attack) (= d1p retreat) (= f1p false) (= f1 true) )
90(and (= b1 true) (= d1p attack) (= d1 retreat) (= f1p false) (= f1 true) )
91))
92
93(define faultAC3::bool (or
94(and (= b2 true) (= d2 attack) (= d2p retreat) (= f2 false) (= f2p true) )
95(and (= b2 true) (= d2p attack) (= d2 retreat) (= f2 false) (= f2p true) )
96(and (= b2 true) (= d2 attack) (= d2p retreat) (= f2p false) (= f2 true) )
97(and (= b2 true) (= d2p attack) (= d2 retreat) (= f2p false) (= f2 true) )
98))
99
100(define faultAC4::bool (or
101(and (= b3 true) (= d3 attack) (= d3p retreat) (= f3 false) (= f3p true) )
102(and (= b3 true) (= d3p attack) (= d3 retreat) (= f3 false) (= f3p true) )
103(and (= b3 true) (= d3 attack) (= d3p retreat) (= f3p false) (= f3 true) )
104(and (= b3 true) (= d3p attack) (= d3 retreat) (= f3p false) (= f3 true) )
105))
106
107;; defining read/write restrictions
108(define rwRest1::bool
109(and (= b1 b1p) (= b2 b2p) (= b3 b3p) (= bg bgp) (= f2 f2p) (= f3 f3p)
110      (= d2 d2p) (= d3 d3p) (= dg dgp)))
111
112(define rwRest2::bool
113(and (= b1 b1p) (= b2 b2p) (= b3 b3p) (= bg bgp) (= f1 f1p) (= f3 f3p)
114      (= d1 d1p) (= d3 d3p) (= dg dgp) ) )
115
116(define rwRest3::bool
117(and (= b2 b2p) (= b1 b1p) (= b3 b3p) (= bg bgp) (= f2 f2p) (= f1 f1p)
118      (= d1 d1p) (= d2 d2p) (= dg dgp) ) )
119
120;; defining the fault-span
121(define fs::bool (or Inv1 (and bg (not b1) (not b2) (not b3))))
122
123;; defining program actions as transition predicates
124(define grd11::bool (and (= f1 false) (= d1 undecided)) )
125(define st11::bool (and (= d1p dgp) (= f1p false) (= d2 d2p) (= d3 d3p) (= dg dgp)))
126(define AC11::bool (and grd11 st11 rwRest1))

```



```

127
128(define grd12::bool
129(and (not b1) (= f1 false) (/= d1 undecided)
130;; uncomment these two constraints and you will get the failsafe program
131;; (or (= d2 undecided) (= d1 d2))
132;; (or (= d3 undecided) (= d1 d3))
133)
134)
135
136(define st12::bool (and (= d1 d1p) (not b1p) (= f1p true)
137                        (= d2 d2p) (= d3 d3p) (= dg dgp)))
138(define AC12::bool (and grd12 st12 rwRest2))
139
140;; defining the set of offending states
141(define os::bool (or
142(and (not b1) (not b2) (/= d1 undecided) (/= d2 undecided)
143      (= f1 true) (= f2 true) (/= d1 d2) )
144(and (not b1) (not b3) (/= d1 undecided) (/= d3 undecided)
145      (= f1 true) (= f3 true) (/= d1 d3) )
146(and (not b3) (not b2) (/= d3 undecided) (/= d2 undecided)
147      (= f3 true) (= f2 true) (/= d3 d2) )
148(and (not bg) (not b1) (/= d1 undecided) (= f1 true) (/= d1 dg))
149(and (not bg) (not b2) (/= d2 undecided) (= f2 true) (/= d2 dg))
150(and (not bg) (not b3) (/= d3 undecided) (= f3 true) (/= d3 dg))
151))
152
153;; defining a primed version of the set of offending states
154(define osp::bool (or
155(and (not b1p) (not b2p) (/= d1p undecided) (/= d2p undecided)
156      (= f1p true) (= f2p true) (/= d1p d2p) )
157(and (not b1p) (not b3p) (/= d1p undecided) (/= d3p undecided)
158      (= f1p true) (= f3p true) (/= d1p d3p) )
159(and (not b3p) (not b2p) (/= d3p undecided) (/= d2p undecided)
160      (= f3p true) (= f2p true) (/= d3p d2p) )
161(and (not bgp) (not b1p) (/= d1p undecided) (= f1p true) (/= d1p dgp))
162(and (not bgp) (not b2p) (/= d2p undecided) (= f2p true) (/= d2p dgp))
163(and (not bgp) (not b3p) (/= d3p undecided) (= f3p true) (/= d3p dgp))
164))
165
166;; defining the Guards predicate that is computed by the coordinator
167(define revisedGuards::bool
168(or
169(or
170(and (not b1) (= f1 false) (= d1 undecided) )
171(and (not b1) (= f1 false) (/= d1 undecided)
172 (or (= d2 undecided) (= d1 d2))
173 (or (= d3 undecided) (= d1 d3))  ))
174(or
175(and (not b2) (= f2 false) (= d2 undecided) )
176(and (not b2) (= f2 false) (/= d2 undecided)
177 (or (= d1 undecided) (= d1 d2))
178 (or (= d3 undecided) (= d2 d3))  ))
179(or

```

```

180 (and (not b3) (= f3 false) (= d3 undecided) )
181 (and (not b3) (= f3 false) (/= d3 undecided)
182 (or (= d2 undecided) (= d3 d2))
183 (or (= d1 undecided) (= d1 d3))  ))
184))
185
186 (define guards::bool
187 (and (not b1) (not b2) (not b3) (not f1) (not f2) (not f3)
188 (or
189 ( or
190 (and (not b1) (= f1 false) (= d1 undecided) )
191 (and (not b1) (= f1 false) (/= d1 undecided)
192 )
193 )
194 ( or
195 (and (not b2) (= f2 false) (= d2 undecided) )
196 (and (not b2) (= f2 false) (/= d2 undecided)
197 )
198 )
199 ( or
200 (and (not b3) (= f3 false) (= d3 undecided) )
201 (and (not b3) (= f3 false) (/= d3 undecided)
202 )
203 )
204 ))
205 )
206
207 (define removedStates::bool
208 (or
209 (not (and
210 (or (= d2 undecided) (= d1 d2))
211 (or (= d3 undecided) (= d1 d3))  ))
212 (not (and
213 (or (= d1 undecided) (= d1 d2))
214 (or (= d3 undecided) (= d2 d3)) ))
215 (not (and
216 (or (= d2 undecided) (= d3 d2))
217 (or (= d1 undecided) (= d1 d3))  ))
218 )
219 )
220
221 (define revisedG1::bool
222 ( or
223 (and (not b1) (= f1 false) (= d1 undecided) )
224 (and (not b1) (= f1 false) (/= d1 undecided)
225 (or (= d2 undecided) (= d1 d2))
226 (or (= d3 undecided) (= d1 d3))  ))
227 )
228
229 ;; The list of queries exchanged with Yices by the reachability nodes,
230 ;; reachability coordinator, synthesis nodes, and the synthesis coordinator.
231
232 ;; Calculating the f-span

```

```

233;; (assert (and Inv faultAC1 (not Invp) ) )
234
235;; the first action can only reach os if it is enabled in os
236;; (assert (and fs grd11 (not os) rwRest1 st11 osp) )
237
238
239;; Does the first action directly violate safety?
240;;(assert (and fs (not Inv) grd11 st11 rwRest1 safety) )
241
242
243;; Does action 2 reach os from states outside mo?
244;; (assert (and fs grd12 (not os) rwRest1 st12 osp) )
245
246;; Does the second action directly violate safety?
247;; (assert (and fs (not Inv) grd12 st12 rwRest1 safety) )
248
249;; Remove os states from the invariant
250;; (assert (and Inv os))
251;; no deadlocks are introduced in the invariant.
252(check)

```

## Appendix A.2: Altitude Switch Controller

In this section, we present the logical context in which we specified the altitude switch program.

```

1 (define-type mode(scalar Init Standby AwaitActuator))
2
3;; defining unprimed variables
4 (define m::mode)
5 (define rst::bool) ;; reset
6 (define inhibit::bool)
7 (define altBelow::bool)
8 (define actStatus::bool)
9 (define initFail::bool)
10 (define altFail::bool)
11 (define actFail::bool)
12
13;; defining the primed variables
14 (define mp::mode)
15 (define rstp::bool)
16 (define inhibitp::bool)
17 (define altBelowp::bool)
18 (define actStatusp::bool)
19 (define initFailp::bool)
20 (define altFailp::bool)
21 (define actFailp::bool)
22

```

```

23;; defining invariant
24(define Inv::bool
25  (and
26    (=> inhibit (not actStatus) )
27    (and (not initFail) (not altFail) (not actFail) )
28  ))
29
30;; defining the safety of specification
31(define safety::bool
32  (or
33    (and altFail (= m Standby) (= mp AwaitActuator))
34    (and initFail (= m Init) (= mp Standby))
35    (and actFail (= m AwaitActuator) (not actStatus) actStatusp)
36  ))
37
38
39;; defining the fault-span
40(define fs::bool
41  (and
42    (=> inhibit (not actStatus) )
43    (or initFail altFail actFail )
44  ))
45
46;; defining read/write restrictions
47(define rwRest::bool
48  (and (= inhibit inhibitp) (= altBelow altBelowp)
49    (= initFail initFailp)
50    (= altFail altFailp)
51    (= actFail actFailp)
52  ))
53
54;; Defining the first program action
55(define grd1::bool (and (= m Init) (not initFail)))
56(define st1::bool (= mp Standby))
57(define A1::bool (and grd1 st1 rwRest))
58
59
60;; Defining the second program action
61(define grd2::bool (and (= m Standby) rst))
62(define st2::bool (and (= mp Init) (not rstp)))
63(define A2::bool (and grd2 st2 rwRest))
64
65;; Defining the third program action
66(define grd3::bool
67  (and (= m Standby) altBelow (not inhibit) (not actStatus) (not altFail)))
68(define st3::bool (= mp AwaitActuator))
69(define A3::bool (and grd3 st3 rwRest))
70
71;; Defining the fourth program action
72(define grd4::bool
73  (and (= m AwaitActuator) (not inhibit) (not actStatus)
74    (not actFail)))
75(define st4::bool (and (= mp Standby) actStatusp) )

```

```

76(define A4::bool (and grd4 st4 rwRest))
77
78
79;; Defining the fifth program action
80(define grd5::bool (and (= m AwaitActuator) rst ) )
81(define st5::bool (and (= mp Init) (not rstp)) )
82(define A5::bool (and grd5 st5 rwRest (= actStatus actStatusp)))
83
84;; Defining the fault actions
85(define F1::bool (and (not initFail) initFailp))
86(define F2::bool (and (not altFail) altFailp))
87(define F3::bool (and (not actFail) actFailp))
88
89
90;; Calculating the f-span
91;; check whether fault actions violate the closure of the invariant
92;; uncomment the following actions one by one
93
94;; (assert (and Inv F1 (not Invp) ))
95;; (assert (and Inv F2 (not Invp) ))
96;; (assert (and Inv F3 (not Invp) ))
97
98;; Does any fault action (starting in the f-span) directly violate safety?
99;; uncomment the following actions one by one
100
101;; (assert (and fs F1 safety ) ) ;; this action does not!
102;; (assert (and fs F2 safety ) ) ;; this action does not!
103;; (assert (and fs F3 safety ) ) ;; this action does not!
104
105;; Therefore, the set of offending state is empty.
106;; No need to execute Steps 1 and 2 of the synthesis node
107
108;; Does each action directly violate safety from f-span outside the invariant?
109
110;; (assert (and fs (not Inv) A1 safety) )
111;; (assert (and fs (not Inv) A3 safety) )
112;; (assert (and fs (not Inv) A4 safety) )
113
114;; since the set of offending states is empty,
115;; no deadlock states are created in the invariant.
116;; the invariant remains unchanged.
117
118(check)

```

## Appendix A.3: Cruise Control System

In this section, we present the Yices specification of the Cruise Control System.

```

1 (define-type mode(scalar Off Inactive Cruise Override Unknown))
2 (define-type leverState(scalar leverOff Constant Resume))
3 (define-type brakeState(scalar brakeUnknown Applied notApplied))
4
5 ;; defining unprimed variables
6 (define SysMode::mode)
7 (define Ignition::bool)
8 (define Lever::leverState)
9 (define EngState::bool)
10 (define Brake::brakeState)
11
12 ;; defining the primed variables
13 (define SysModep::mode)
14 (define Ignitionp::bool)
15 (define Leverp::leverState)
16 (define EngStatep::bool)
17 (define Brakep::brakeState)
18
19 ;; Defining the invariant
20 (define I1::bool
21 (and (/= SysMode Unknown)
22 (=> (/= SysMode Off) (= Ignition false))
23 (=> (/= SysMode Inactive) (= Ignition true))
24 (=> (/= SysMode Cruise) (and (= Ignition true) (= EngState true)
25 (/= Brake Applied) (/= Lever leverOff)))
26 (=> (/= SysMode Override) (and (= Ignition true) (= EngState true)))
27))
28
29 (define Inv::bool
30 (and (/= SysMode Unknown) (/= Brake brakeUnknown) I1 ))
31
32 ;; Defining the primed version of the invariant
33 (define I1p::bool
34 (and (/= SysModep Unknown)
35 (=> (/= SysModep Off) (= Ignitionp false))
36 (=> (/= SysModep Inactive) (= Ignitionp true))
37 (=> (/= SysModep Cruise) (and (= Ignitionp true) (= EngStatep true)
38 (/= Brakep Applied) (/= Leverp leverOff)))
39 (=> (/= SysModep Override) (and (= Ignitionp true) (= EngStatep true)))
40))
41
42 (define Invp::bool
43 (and (/= SysModep Unknown) (/= Brakep brakeUnknown) I1p ))
44
45
46 ;; Defining the safety of specification
47 (define safety::bool
48 (or
49 (and (= Brake Applied) (= Brakep Applied) (= SysModep Cruise))
50 (and (= Brake brakeUnknown) (= SysModep Cruise))
51 (and (= SysMode Unknown) (= SysModep Cruise)) ))
52
53 ;; Defining the f-span

```

```

54(define fs::bool
55(or I1 (= SysMode Unknown) (= Brake brakeUnknown) ))
56
57;; Defining read/write restrictions
58(define rwRest::bool
59(and (= Ignition Ignitionp) (= EngState EngStatep)
60(= Brake Brakep) (= Lever Leverp) ))
61
62;; Defining the first action
63(define grd1::bool (and (= SysMode Off) Ignition ))
64(define st1::bool (= SysModep Inactive))
65(define A1::bool (and grd1 st1 rwRest))
66
67;; Defining the second action
68(define grd2::bool (and (= SysMode Inactive) (not Ignition) ) )
69(define st2::bool (= SysModep Off))
70(define A2::bool (and grd2 st2 rwRest))
71
72;; Defining the third action
73(define grd3::bool
74(and (= SysMode Inactive) (= Lever Constant) Ignition EngState
75;; this additional constraint is added for the sake of failsafe fault-tolerance.
76;; (= Brake notApplied)
77))
78(define st3::bool (= SysModep Cruise))
79(define A3::bool (and grd3 st3 rwRest))
80
81;; Defining the fourth action
82(define grd4::bool (and (= SysMode Cruise) (not Ignition) ))
83(define st4::bool (= SysModep Off) )
84(define A4::bool (and grd4 st4 rwRest))
85
86;; Defining the fifth action
87(define grd5::bool (and (= SysMode Cruise) (not EngState) ))
88(define st5::bool (= SysModep Inactive) )
89(define A5::bool (and grd5 st5 rwRest))
90
91;; Defining the sixth action
92(define grd6::bool (and (= SysMode Cruise)
93      (or (= Brake Applied) (= Lever leverOff)) ))
94(define st6::bool (= SysModep Override))
95(define A6::bool (and grd6 st6 rwRest))
96
97;; Defining the seventh action
98(define grd7::bool (and (= SysMode Override) (not Ignition) ))
99(define st7::bool (= SysModep Off) )
100(define A7::bool (and grd7 st7 rwRest))
101
102;; Defining the Eight action
103(define grd8::bool (and (= SysMode Override) (not EngState) ))
104(define st8::bool (= SysModep Inactive) )
105(define A8::bool (and grd8 st8 rwRest))
106

```

```

107;; Defining the ninth action
108(define grd9::bool (and (= SysMode Override) Ignition EngState
109                        (/= Lever leverOff)
110;; this additional constraint is added for the sake of failsafe fault-tolerance.
111(= Brake notApplied)
112))
113(define st9::bool (= SysModep Cruise) )
114(define A9::bool (and grd9 st9 rwRest))
115
116
117;; Defining the first fault action
118(define F1::bool (and (/= Brake brakeUnknown) (= Brakep brakeUnknown) rwRest))
119
120;; The reason why we need to have rwRest here is to codify
121;; the fact that this fault action only updates the Brake variable;
122;; not that fault actions have read/write restrictions!
123
124;; Defining the second fault action
125(define F2::bool (and (/= SysMode Unknown) (= SysModep Unknown) rwRest))
126
127;; Calculating the f-span
128;; check whether fault actions violate the closure of the invariant
129;; uncomment the following actions one by one
130
131;; (assert (and Inv F1 (not Invp) ) )
132
133;; (assert (and Inv F2 (not Invp) ) )
134
135;; Does any fault action (starting in the f-span) directly violate safety?
136;; uncomment the following actions one by one
137
138;; (assert (and fs F1 safety) ) ;; this action does not!
139;; (assert (and fs F2 safety) ) ;; this action does not!
140
141;; Therefore, the set of offending state is empty.
142;; No need to execute Steps 1 and 2 of the synthesis node
143
144
145;; Does each action directly violate safety from f-span outside the invariant?
146
147;; (assert (and fs (not Inv) A3 safety) )
148
149;; (assert (and fs (not Inv) A9 safety) )
150
151;; since the set of offending states is empty,
152;; no deadlock states are created in the invariant.
153;; the invariant remains unchanged.
154
155(check)

```



## Appendix A.4: Token Ring

In this section, we present the logical context in which we specified the token ring program.

```
1;; defining unprimed variables
2(define x0::(subtype (n::int) (and (> n -2) (< n 2) )) )
3(define x1::(subtype (n::int) (and (> n -2) (< n 2) )) )
4(define x2::(subtype (n::int) (and (> n -2) (< n 2) )) )
5(define x3::(subtype (n::int) (and (> n -2) (< n 2) )) )
6
7;; defining the primed variables
8(define x0p::(subtype (n::int) (and (> n -2) (< n 2) )) )
9(define x1p::(subtype (n::int) (and (> n -2) (< n 2) )) )
10(define x2p::(subtype (n::int) (and (> n -2) (< n 2) )) )
11(define x3p::(subtype (n::int) (and (> n -2) (< n 2) )) )
12
13;; defining the invariant
14(define Inv::bool
15(or
16(and (= x0 0) (= x1 0) (= x2 0) (= x3 0))
17(and (= x0 1) (= x1 0) (= x2 0) (= x3 0))
18(and (= x0 1) (= x1 1) (= x2 0) (= x3 0))
19(and (= x0 1) (= x1 1) (= x2 1) (= x3 0))
20(and (= x0 1) (= x1 1) (= x2 1) (= x3 1))
21(and (= x0 0) (= x1 1) (= x2 1) (= x3 1))
22(and (= x0 0) (= x1 0) (= x2 1) (= x3 1))
23(and (= x0 0) (= x1 0) (= x2 0) (= x3 1)) ))
24
25(define Invp::bool
26(or
27(and (= x0p 0) (= x1p 0) (= x2p 0) (= x3p 0))
28(and (= x0p 1) (= x1p 0) (= x2p 0) (= x3p 0))
29(and (= x0p 1) (= x1p 1) (= x2p 0) (= x3p 0))
30(and (= x0p 1) (= x1p 1) (= x2p 1) (= x3p 0))
31(and (= x0p 1) (= x1p 1) (= x2p 1) (= x3p 1))
32(and (= x0p 0) (= x1p 1) (= x2p 1) (= x3p 1))
33(and (= x0p 0) (= x1p 0) (= x2p 1) (= x3p 1))
34(and (= x0p 0) (= x1p 0) (= x2p 0) (= x3p 1)) ))
35
36;; defining the safety of specification
37(define safety::bool
38(or
39(and (= x0 -1) (/= x1 -1) (= x1p -1))
40(and (= x1 -1) (/= x2 -1) (= x2p -1))
41(and (= x2 -1) (/= x3 -1) (= x3p -1))
42(and (= x3 -1) (/= x0 x0p)) ))
43
44;; defining read/write restrictions
45(define rwRest0::bool
46(and (= x1 x1p) (= x2 x2p)
47(= x3 x3p) ;; we need this to represent the atomicity of the action
48))
```

```

49
50(define rwRest1::bool
51(and (= x3 x3p) (= x2 x2p) (= x0 x0p) ))
52
53(define rwRest2::bool
54(and (= x0 x0p) (= x1 x1p) (= x3 x3p) ))
55
56(define rwRest3::bool
57(and (= x1 x1p) (= x0 x0p) (= x2 x2p) ))
58
59;; Defining the first action
60(define grd0::bool
61(and (/= x3 -1) (= x0 x3)))
62(define st0::bool (= x0p (mod (+ x3 1) 2) ))
63(define A0::bool (and grd0 st0 rwRest0))
64
65;; Defining the second action
66(define grd1::bool
67(and (/= x0 -1) (/= x0 x1) ))
68(define st1::bool (= x1p x0 ))
69(define A1::bool (and grd1 st1 rwRest1))
70
71;; Defining the third action
72(define grd2::bool
73(and (/= x1 -1) (/= x2 x1) ))
74(define st2::bool (= x2p x1 ))
75(define A2::bool (and grd2 st2 rwRest2))
76
77;; Defining the fourth action
78(define grd3::bool
79(and (/= x2 -1) (/= x2 x3) ))
80(define st3::bool (= x3p x2 ))
81(define A3::bool (and grd3 st3 rwRest3))
82
83;; Defining fault actions
84(define faultgrd::bool
85(or
86(and (/= x0 -1) (/= x1 -1))
87(and (/= x0 -1) (/= x2 -1))
88(and (/= x0 -1) (/= x3 -1))
89(and (/= x1 -1) (/= x2 -1))
90(and (/= x1 -1) (/= x3 -1))
91(and (/= x2 -1) (/= x3 -1) )))
92
93
94(define faultA0::bool
95(and (/= x3 -1) (/= x0 -1) (= x0p -1) (= x1 x1p) (= x2 x2p) (= x3 x3p)))
96
97(define faultA1::bool
98(and (/= x0 -1) (/= x1 -1) (= x1p -1) (= x0 x0p) (= x2 x2p) (= x3 x3p)))
99
100(define faultA2::bool
101(and (/= x1 -1) (/= x2 -1) (= x2p -1) (= x1 x1p) (= x0 x0p) (= x3 x3p)))

```

```

102
103 (define faultA3::bool
104 (and (/= x2 -1) (/= x3 -1) (= x3p -1) (= x1 x1p) (= x2 x2p) (= x0 x0p)))
105
106 ;; Calculating the f-span
107 ;; check whether fault actions violate the closure of the invariant
108 ;; uncomment the following actions one by one
109
110 ;; (assert (and Inv faultA0 (not Invp) ))
111 ;; (assert (and Inv faultA1 (not Invp) ))
112 ;; (assert (and Inv faultA2 (not Invp) ))
113 ;; (assert (and Inv faultA3 (not Invp) ))
114
115 ;; Does any fault action (starting in the f-span) directly violate safety?
116 ;; uncomment the following actions one by one
117
118 ;; (assert (and fs faultA0 safety ) ) ;; this action does not!
119 ;; (assert (and fs faultA1 safety ) ) ;; this action does not!
120 ;; (assert (and fs faultA2 safety ) ) ;; this action does not!
121 ;; (assert (and fs faultA3 safety ) ) ;; this action does not!
122
123 ;; Therefore, the set of offending state is empty.
124 ;; No need to execute Steps 1 and 2 of the synthesis node
125
126 ;; Does each action directly violate safety from f-span outside the invariant?
127 ;; yes when they propagate a corrupted value
128
129 ;; (assert (and fs (not Inv) A0 safety) )
130 ;; (assert (and fs (not Inv) A1 safety) )
131 ;; (assert (and fs (not Inv) A2 safety) )
132 ;; (assert (and fs (not Inv) A3 safety) )
133
134 ;; since the set of offending states is empty,
135 ;; no deadlock states are created in the invariant.
136 ;; the invariant remains unchanged.
137
138 (check)

```