A correctness analysis for the algorithm presented in "Symbolic Computation of Boundary Unsafe States in Complex Resource Allocation Systems Using Partitioning Techniques" by Z. Fei, K. Akesson and S. Reveliotis, IEEE CASE 2015 (submitted)

Zhennan Fei, Knut Åkesson and Spyros Reveliotis

To prove the correctness of the considered algorithm, we need to show that (i) the algorithm terminates in a finite number of steps, (ii) the returned set χ_{FB} contains all the feasible boundary unsafe states, and furthermore, (iii) χ_{FB} does not contain any feasible safe state. We start by addressing item (i).

Theorem 0.1. *The considered algorithm terminates in a finite number of steps.*

Proof: To establish Theorem 0.1, we make the following remarks: From the formal statement of the considered algorithm, it is clear that it will terminate if and only if (*iff*) the condition of Line 29 is met at some iteration. But since the underlying state space is finite, and every non-terminating iteration increases χ_U by at least one state, to establish the eventual satisfaction of the condition of Line 29, it suffices to establish that the set $\chi_{U_{new}}$ generated at each iteration is non-overlapping with the set χ_U produced by the previous iterations. This is attained by stating and proving the following lemma.

Lemma 0.1. *In the iterations of the considered algorithm encoded by Lines 13-29,* $\chi_U \wedge \chi_{U_{new}} = 0$ *until the execution of Line 28.*

Proof: Lemma 0.1 can be proved by noticing that since (i) each characteristic function Δ_{σ_i} encodes the transitions of a single event (type) σ_i, and (ii) the FSA corresponding to the RAS-modeling EFA Φ is deterministic, the STD corresponding to each Δ_{σ_i} is a set of "in-trees" (i.e., each state appearing in this STD has possibly more than one incoming transitions but only one emanating transition). Because of this structure, each state appearing in some Δ_{σ_i} can be reached during the backtracing of the transitions of Δ_{σ_i} from a single path only, and therefore, only once.

Next, consider a deadlock-free unsafe state u, with emanating transitions corresponding to events $\sigma_k,\ k \in K \subseteq \{1, \ldots, \mu\}$. As long as there is a set Δ_{σ_k} for which state u does not belong in $\chi_{LU}^{\sigma_k}$ (i.e., u has not been reached yet through the backtracing steps in Δ_{σ_k}), u will be still in $\chi_{NU}^{\sigma_k}$, and therefore, it cannot be recognized as unsafe yet. Assuming that at some iteration u has been reached, through backtracing, in all $\Delta_{\sigma_k},\ k \in K$, then, u will enter $\chi_{U_{new}}$ at that iteration, through the execution of Lines 21-25, and eventually it will also enter χ_U, through the execution of Line 28. But due to the "in-tree" structure of the transition sets Δ_{σ_i}, u will not be encountered again in the subsequent backtracing of these sets. This last remark settles the validity of, both, Lemma 0.1 and Theorem 0.1. □

Next, we proceed to establish the soundness of the considered algorithm, i.e., items (ii) and (iii) in the requirements list that was provided in the opening paragraph of this document. We shall develop the sought results by establishing a series of lemmas.

Lemma 0.2. *The characteristic function χ_{FD} that is obtained from the symbolic operations performed in Lines 1-10 of the considered algorithm identifies correctly the feasible deadlock states w.r.t. the process-advancing events $\sigma_1, \ldots, \sigma_\mu$ from the transition sets $\Delta_{\sigma_1}, \ldots, \Delta_{\sigma_\mu}$.*

The validity of the above statement should be evident from the description of this part of the algorithm that is provided in the main text, and its formal proof is omitted for the sake of brevity. On the other hand, the next lemmas establish that the considered algorithm observes state feasibility.

Lemma 0.3. *For every transition (s, s') of the EFA $\Delta_{\mathbf{E}}$, feasibility of the target state s' implies also the feasibility of the source state s.*

Proof: We prove the contrapositive of the above statement, i.e., every transition (s, s') of the EFA $\Delta_{\mathbf{E}}$ with an infeasible source state s has also an infeasible target state s'. Infeasibility of state s implies that there exists some resource R_i with

$$vR_i + \sum_{j=1}^{n} \sum_{k=1}^{l(j)-1} \mathcal{A}_{jk}[i] * v_{jk} = d \neq C_i,$$

for the values of the variables vR_i and $v_{jk},\ j = 1, \ldots, n,\ k = 1, \ldots, l(j) - 1$ that define state s. But it can be easily checked that every forward-advancing transition from state s preserves the invariant

$$vR_i + \sum_{j=1}^{n} \sum_{k=1}^{l(j)-1} \mathcal{A}_{jk}[i] * v_{jk} = d,$$

and therefore, it cannot restore feasibility w.r.t. to the implied allocation of resource R_i. □

Lemma 0.4. *All states entering the sets U and χ_{FB} during the execution of Algorithm 2 are feasible.*

Proof: This lemma is an immediate implication of Lemmas 0.2 and 0.3, and of the fact that all the elements of the sets U and χ_{FB} are obtained by starting from some feasible deadlock state in χ_{FD} and backtracing upon some transitions in $\Delta_{\mathbf{E}}$. □

Having established the feasibility of the states that are generated by the considered algorithm, the next two lemmas address the additional properties of these states that define the algorithm correctness.

Lemma 0.5. *The set U that is computed by the considered algorithm contains all the feasible unsafe states in Δ_E.*

Proof: By Lemma 0.2 and Line 11 of the considered algorithm, U contains all the feasible deadlocks. Next we will show that U also contains all the feasible deadlock-free unsafe states of the considered RAS.

Let us consider any such feasible deadlock-free unsafe state $\hat{u}$. The finite and acyclic nature of the paths that define the execution logic of the various process types in the considered RAS class, implies that the subspace that is reached from state $\hat{u}$ following only transitions in $\Delta_{\sigma_1} \vee \ldots \vee \Delta_{\sigma_\mu}$ has a finite, acyclic structure. This remark, when combined with the presumed unsafety of state $\hat{u}$, further imply that every path in $\Delta_{\sigma_1} \vee \ldots \vee \Delta_{\sigma_\mu}$ that emanates from state $\hat{u}$ is an acyclic path that terminates at some feasible deadlock state. Let ζ denote the longest length of these paths, where the length of a path is defined by the number of the involved transitions. Next, we will show, by induction on ζ, that state $\hat{u}$ will enter the state set U that is maintained by the considered algorithm before the termination of the

iteration in Lines 13-29. Also, in the following, we denote by u_k the respective states resulting from state $\hat{u}$ by executing its process-advancing events σ_k, where $k \in K \subseteq \{1, \ldots, \mu\}$.

First, we consider the base case of $\zeta = 1$. Then, each state u_k that is reached from $\hat{u}$ is contained in χ_{FD} as a deadlock state; therefore, each transition $(\hat{u}, u_k)$ will be contained in the corresponding $\Delta_U^{\sigma_k}$. Hence, after the operations performed in Lines 16-18, $\hat{u}$ will be contained in $\chi_{SU}^{\sigma_k}$ and $\chi_{LU}^{\sigma_k}$, for each $k \in K$. Also, we notice that $\hat{u}$ cannot be in any $\chi_{NU}^{\sigma_i}$, for all $i = 1, \ldots, \mu$. Therefore, state $\hat{u}$ will be correctly identified as an unsafe state in Lines 21-25, and eventually it will be included in the set χ_U.

Next, let us suppose that all the feasible unsafe states with a maximal path of length $\zeta - 1$ from the feasible deadlock states of χ_{FD} are correctly identified and included in set U by the considered algorithm. Since the target state of each process-advancing transition $(\hat{u}, u_k)$ that emanates from state $\hat{u}$ has a maximal path leading to χ_{FD} of length less than or equal to $\zeta - 1$, by the working hypothesis, each u_k eventually will be identified by the algorithm. Let us consider, in particular, the iteration where the last of these states, say u_l where $l \in K$, enters U. In the next iteration, the transition $(\hat{u}, u_l)$ will be in $\Delta_U^{\sigma_l}$ and $\hat{u}$ will be in $\chi_{SU}^{\sigma_l}$ (and thus, it is not in $\chi_{NU}^{\sigma_l}$). Note that also state $\hat{u}$ is not in $\chi_{NU}^{\sigma_k}$, for any $k \in K \setminus \{l\}$, since $\hat{u}$ has been added in $\chi_{LU}^{\sigma_k}$ at earlier iterations. Hence, $\hat{u}$ will be included in $\chi_{U_{new}}$ and eventually into χ_U. □

Lemma 0.6. *The set U that is computed by the considered algorithm contains no feasible safe state of Δ_E.*

Proof: We prove this lemma by induction on the number of iterations performed by the algorithm. The base case of zero iterations is covered by Lemma 0.2. Next, suppose that the statement of Lemma 0.6 is true for the set U constructed during the first n iterations. Then, as discussed in the algorithm description that is provided in the main text, at iteration $n + 1$, every state entering $\chi_{U_{new}}$ has all its emanating transitions leading into previously identified unsafe states, and therefore, it is correctly classified as a new unsafe state. □

Now we are ready to state and prove the main result regarding the soundness of the considered algorithm.

Theorem 0.2. *The set χ_{FB} returned by the considered algorithm possesses the following properties: (i) It contains only feasible states. (ii) It contains all the feasible boundary unsafe states in the underlying RAS state-space. (iii) It contains no feasible non-boundary unsafe state. (iv) It contains no safe state.*

Proof: Property (i) was established in Lemma 0.4.

Lemma 0.5, when combined with the logic of Lines 19 and 27 in the considered algorithm, imply that the states in χ_{FBA}, which are the target states of $\Delta_{U_{pre}}^{\sigma_i}, \forall i = 1, \ldots, \mu$, constitute all the boundary unsafe states that are reached by safe states through the processing-advancing events $\sigma_1, \ldots, \sigma_\mu$. Similarly, Lemma 0.5 and Lines 30-31 imply that the states in χ_{FBL} are all the boundary unsafe states that are reached by safe states through the loading events. Hence, Property (ii) holds.

Property (iii) can be established by noticing that (a) all states in χ_{FBL} are boundary unsafe since they are reached by some feasible safe states, and (b), by construction, the source states of $\Delta_{U_{pre}}^{\sigma_i}, \forall i = 1, \ldots, \mu$, are safe states while the target states are unsafe.

Finally, Property (iv) results from Lemma 0.6 and the fact that all transitions in Δ_{LU} and $\Delta_{U_{pre}}^{\sigma_i}, \forall i = 1, \ldots, \mu$, have target states in U. □