Foundations of the UML

Lecture 10: Regular MSCs

Joost-Pieter Katoen

Lehrstuhl für Informatik 2
Software Modeling and Verification Group

http://moves.rwth-aachen.de/i2/370

23. November 2009

RWTH AACHEN UNIVERSITY

Realisabiliy and safe realisability

Definition (Realisability)

1. MSC M is realisable whenever $\{M\} = L(\mathcal{A})$ for some CFM $\mathcal{A}$.
2. A finite set $\{M_1, \ldots, M_n\}$ of MSCs is realisable whenever $\{M_1, \ldots, M_n\} = L(\mathcal{A})$ for some CFM $\mathcal{A}$.
3. MSG G is realisable whenever $L(G) = L(\mathcal{A})$ for some CFM $\mathcal{A}$.

Definition (Safe realisability)

Same as above except that the CFM should be deadlock-free

RWTH AACHEN UNIVERSITY

Summary of results

Approach so far:

The (safe) realisation of a (finite) set of MSCs by a weak CFM is the one where the automaton $\mathcal{A}_p$ of process p generates the projections of these MSCs on p.

Results so far:

1. Conditions for (safe) realisability for languages obtained by finite sets of MSCs.
2. Checking safe realisability for such languages is in P.
3. Checking realisability for such languages is co-NP complete.

Some remaining questions

- Can results be obtained for larger classes of MSGs?
- What happens if we allow synchronisation messages?
  - recall that weak CFMs in fact do not involve synchronisation messages
- How do we obtain a CFM realising an MSG algorithmically?
  - in particular, for non-local choice MSGs
- Are there simple conditions on MSGs that guarantee realisability?
  - e.g., easily identifiable subsets of (safe) realisable MSGs

Regular MSCs

Let $\mathbb{M}$ be the set of MSCs over $\mathcal{P}$ and $\mathcal{C}$.

Definition (Regular)

1. $\mathcal{M} = \{ M_1, \ldots, M_n \}$ with $n \in \mathbb{N} \cup \{ \infty \}$ is called regular if $Lin(\mathcal{M}) = \bigcup_{i=1}^{n} Lin(M_i)$ is a regular word language over Act^*.
2. MSG G is regular if $Lin(G)$ is a regular word language over Act^*.
3. CFM $\mathcal{A}$ is regular if $Lin(\mathcal{A})$ is a regular word language over Act^*.

Note that $\mathbb{M}$ itself is not regular.

Obviously we have:

Any $\forall$-bounded CFM is regular.

Examples

On the black board.

A decidability result

Theorem [Henriksen et. al, 2005]

The decision problem "does a regular language $L \subseteq Act^*$ represent a set of well-formed words"? —that is, does L represent a set of MSCs?— is decidable.

Proof

Since L is regular, there exists a minimal DFA $\mathcal{A} = (S, Act, s_0, \delta, F)$ that accepts L. Consider the productive states in this DFA, i.e., all states from which some state in F can be reached. We label any productive state s with a channel-capacity function $K_s : Ch \to \mathbb{N}$ such that 4 constraints (cf. next slide) are fulfilled. Then: L is a regular set of well-formed words iff each productive state in the DFA $\mathcal{A}$ can be labeled with K_s satisfying these constraints. In fact, if a state-labeling violates any of these constraints, it is due to a word that is not well-formed.

Constraints on state-labelling

Constraints on channel-capacity function K_s for state s:

1. If $s \in F \cup \{s_0\}$, then $K_s((p, q)) = 0$ for any channel (p, q)
2. If $s, s' \in S$ are productive and $\delta(s, !(p, q, a)) = s'$, then $K_{s'}((p, q)) := K_s((p, q)) + 1$, and $K_{s'} = K_s$ for all other channels
3. If $s, s' \in S$ are productive and $\delta(s, ?(p, q, a)) = s'$, then $K_s((q, p)) > 0$, $K_{s'}((q, p)) := K_s((q, p)) - 1$, and $K_{s'} = K_s$ for all other channels
4. (The "diamond" property).
   If $\delta(s, \alpha) = s_1$ and $\delta(s_1, \beta) = s_2$ with $\alpha \in Act_p$ and $\beta \in Act_q$, $p \neq q$, then if:

   $$\text{not } (\alpha \neq !(p, q, a) \text{ and } \beta \neq ?(q, p, a)), \text{ or } K_s((p, q)) > 0$$

   then $\delta(s, \beta) = s'_1$ and $\delta(s'_1, \alpha) = s_2$ for some $s'_1 \in S$.

Example

On the black board.

Boundedness and regularity

Definition (B-bounded words)

Let $B \in \mathbb{N}$ and $B > 0$. A word $w \in Act^*$ is called B-bounded if for any prefix u of w and any channel $(p, q) \in Ch$:

$$0 \leqslant \sum_{a \in \mathcal{C}} |u|_{!(p,q,a)} - \sum_{a \in \mathcal{C}} |u|_{?(q,p,a)} \leqslant B$$

Corollary:

For any regular, well-formed language L, there exists $B \in \mathbb{N}$ and $B > 0$ such that any $w \in L$ is B-bounded.

Proof

The bound is the largest value attained by the channel-capacity functions assigned to productive states in the proof of the previous theorem.

Regularity and realisability

Theorem: [Henriksen et al., 2005, Baudru & Morin, 2007]

For any set L of well-formed words, the following statements are equivalent:

1. L is regular.
2. L is realisable by a $\forall$-bounded CFM.
3. L is realisable by a deterministic $\forall$-bounded CFM.
4. L is safely realisable by a $\forall$-bounded CFM.

Note:

The maximal size of the CFM realising L is such that for each process p, $|S_p|$, the number of states of local automaton $\mathcal{A}_p$, is double exponential in the bound B, and n^2 where $n = |\mathcal{P}|$, and exponential in $m \log m$ where m is the size of a minimal DFA representing L.

Regularity for MSGs is undecidable

Theorem [Henriksen et. al, 2005]

The decision problem “is MSG G regular“? is undecidable.

Proof

By a reduction from the (undecidable) problem to determine whether the trace-closure of a regular language L over alphabet Σ with respect to an independence relation $I \subseteq \Sigma \times \Sigma$.
(Proof omitted in this lecture.)

Towards structural conditions for regular MSGs

- MSG G is regular if $Lin(G)$ is a regular language
- Regularity yields deterministic, or safe, but bounded CFMs
- But, "is MSG G regular"? is unfortunately undecidable
- Is it possible to impose structural conditions on MSGs that guarantee regularity?
- Yes we can. For instance, by constraining:
  1. the communication structure of the MSCs in loops of G, or
  2. the structure of rational expressions describing the MSCs in G

Communication graph

Definition (Communication graph)

The communication graph of the MSC $M = (\mathcal{P}, E, \mathcal{C}, l, m, <)$ is the directed graph $(V, \rightarrow)$ with:

- $V = \mathcal{P} \setminus \{ p \in \mathcal{P} \mid E_p = \varnothing \}$, the set of active processes
- $(p, q) \in \rightarrow$ if and only if $l(e) = !(p, q, a)$ for some $e \in E$

Example

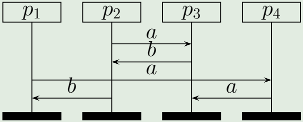

MSC

an example

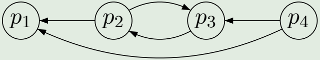

its communication graph

Communication closedness

Definition

Communication graph MSG G is communication-closed if for any loop $\pi = v_1 v_2 \ldots v_n$ (with $v_1 = v_n$) in G, the MSC $M(\pi)$ has a strongly connected communication graph.

Example

On the black board.

Communication-closed vs. regularity

Theorem:

Any communication-closed MSG G is regular.

Example

Example on the black board.

Note:

The converse does not hold (cf. next slide).

Communication-closed vs. regularity

Communication-closedness is not a necessary condition for regularity:

G:

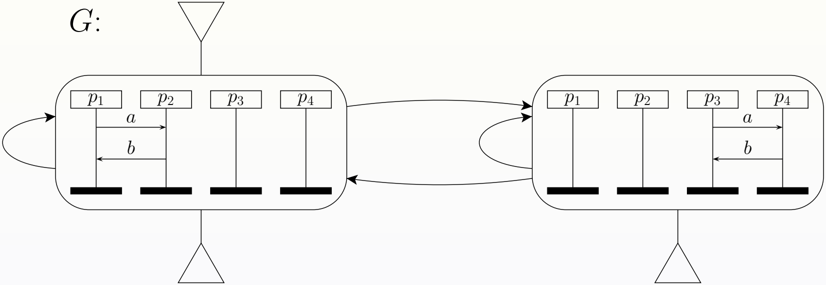

MSG G is not communication-closed, but $Lin(G)$ is regular.

Communication-closed vs. regularity

Definition (Asynchronous iteration)

For $\mathcal{M}_1, \mathcal{M}_2 \subseteq \mathbb{M}$ sets of MSCs, let:

$$\mathcal{M}_1 \bullet \mathcal{M}_2 \;=\; \{\, M_1 \bullet M_2 \mid M_1 \in \mathcal{M}_1, M_2 \in \mathcal{M}_2 \,\}$$

For $\mathcal{M} \subseteq \mathbb{M}$ let

$$\mathcal{M}^i \;=\; \begin{cases} \{M_\epsilon\} & \text{if } i{=}0\text{, where } M_\epsilon \text{ denotes the empty MSC} \\ \mathcal{M} \bullet \mathcal{M}^{i-1} & \text{if } i > 0 \end{cases}$$

The asynchronous iteration of $\mathcal{M}$ is now defined by:

$$\mathcal{M}^* = \bigcup_{i \geqslant 0} \mathcal{M}^i.$$

Communication-closed vs. regularity

Definition (Finitely generated)

Set of MSCs $\mathcal{M} \subseteq \mathbb{M}$ is finitely generated if there is a finite set of MSCs $\widehat{\mathcal{M}} \subseteq \mathbb{M}$ such that $\mathcal{M} \subseteq \widehat{\mathcal{M}}^*$.

Notes:

1. Each set of MSCs defined by MSG G is finitely generated.
2. Not every regular language of well-formed words is finitely generated.
3. Not every finitely generated set of MSCs is regular.
4. It is decidable to check whether a set of MSCs is finitely generated.

Communication-closed vs. regularity

Theorem: [Henriksen et. al, 2005]

Let $\mathcal{M}$ be a set of MSCs. Then:

$$\mathcal{M} \text{ is finitely generated and regular}$$

$$\text{iff}$$

$$\mathcal{M} = L(G) \text{ for communication-closed MSG } G.$$

Checking communication-closedness

Theorem: [Genest et. al, 2006]

The decision problem "is MSG G communication closed" is co-NP complete.

Proof

1. Membership in co-NP can be proven in a standard way: guess a subgraph of G, check in polynomial time whether this subgraph has a loop passing through all its vertices, and check whether its communication graph is not strongly connected.
2. It can be shown that the problem is co-NP hard by a reduction from the 3-SAT problem.

Local communication-closedness

Definition (Local communication-closedness)

MSG G is locally communication-closed if for each vertex (v, v') in G, the MSCs $\lambda(v_1)$, $\lambda(v_2)$ and $\lambda(v_1) \bullet \lambda(v_2)$ have weakly connected communication graphs.

Notes:

1. A directed graph is weakly connected if its induced undirected graph is strongly connected.
2. Checking whether MSG G is locally communication-closed can be done in linear time.

Theorem:

For any locally communication-closed MSG G, there exists a CFM $\mathcal{A}$ with $L(\mathcal{A}) = L(G)$ of size $n^{\mathcal{O}(|\mathcal{P}|)}$ where n is the number of vertices in G.

Summary of realisability [Lohrey, 2003]

Computability and complexity results for FIFO communication:

	finite MSGs*	communication-closed MSGs	general MSGs
realisabilty	co-NP complete	undecidable	undecidable
safe realisability	PTIME	EXPSPACE-complete	undecidable

* MSG G is finite if $L(G)$ is a finite set of MSCs.

Computability and complexity results for non-FIFO communication:

	finite MSGs*	communication-closed MSGs	general MSGs
realisabilty	co-NP complete	PSPACE-hard	undecidable
safe realisability	PTIME	EXPSPACE-complete	undecidable

RWTH AACHEN UNIVERSITY