Security Requirements for Uniformly Parameterised Cooperations

Peter Ochsenschläger and **Roland Rieke**
peter-ochsenschlaeger@t-online.de, roland.rieke@sit.fraunhofer.de
Fraunhofer Institute for Secure Information Technology SIT, Darmstadt, Germany

PDP, February 2012

Fraunhofer SIT

Overview of Approach

1. Uniform parameterisations of cooperations are defined.
   - Each pair of partners cooperates in the same manner.
   - The mechanism (schedule) to determine how one partner may be involved in several cooperations, is the same for each partner.
2. By generalising this, self-similarity is formalised.
3. Uniformly parameterised behaviour properties are defined.

safety	something bad does not happen
liveness	something good eventually happens
+ fairness	something good eventually is possible

4. The parameterised problem of verifying such a property is reduced by self-similarity and simplicity of the abstraction to a finite state problem.
5. An example for a finite state verification of uniformly parameterised behaviour properties is given (in the paper).

Motivation - Security and Reliability of Scalable Systems

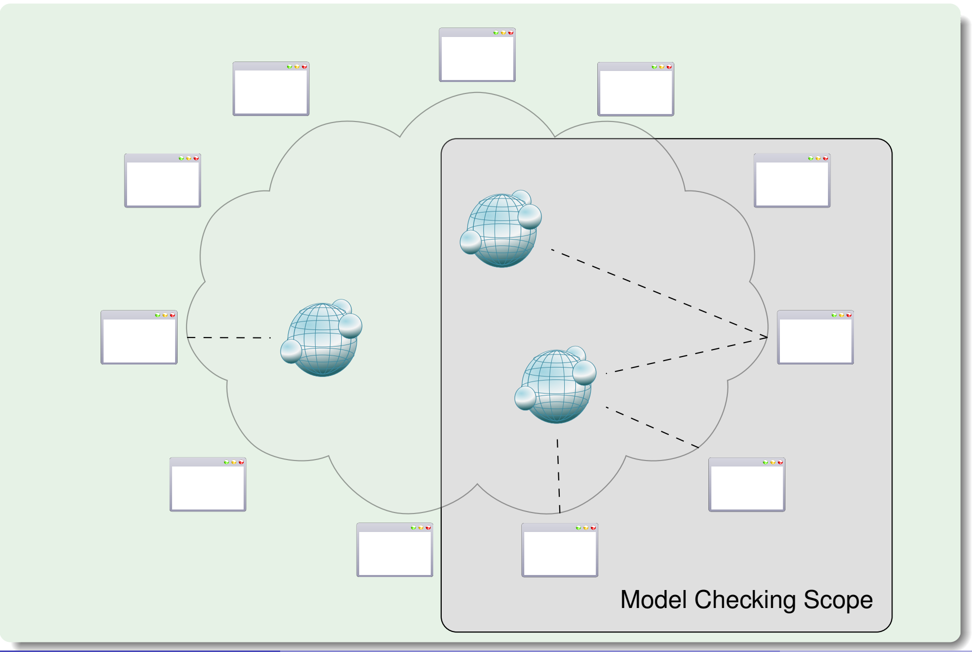

Parameterised Systems

- Parameterised systems describe families of systems that are finite-state in nature but scalable.
- Instances of the family can be obtained by fixing the parameters (e.g. number of replicated components).
- For safety/business critical systems, assuring the correctness - conformance to the intended purpose - is imperative.
- These systems must guarantee a variety of safety, liveness and security properties.
- Traditional model checking techniques to analyse behaviour of such systems are limited to systems with very few components.

We assume that uniformly parameterised structures are likely to appear in any highly scalable system (of systems), such as cloud computing platforms or cyber-physical systems.

Uniformly Parameterised Cooperations

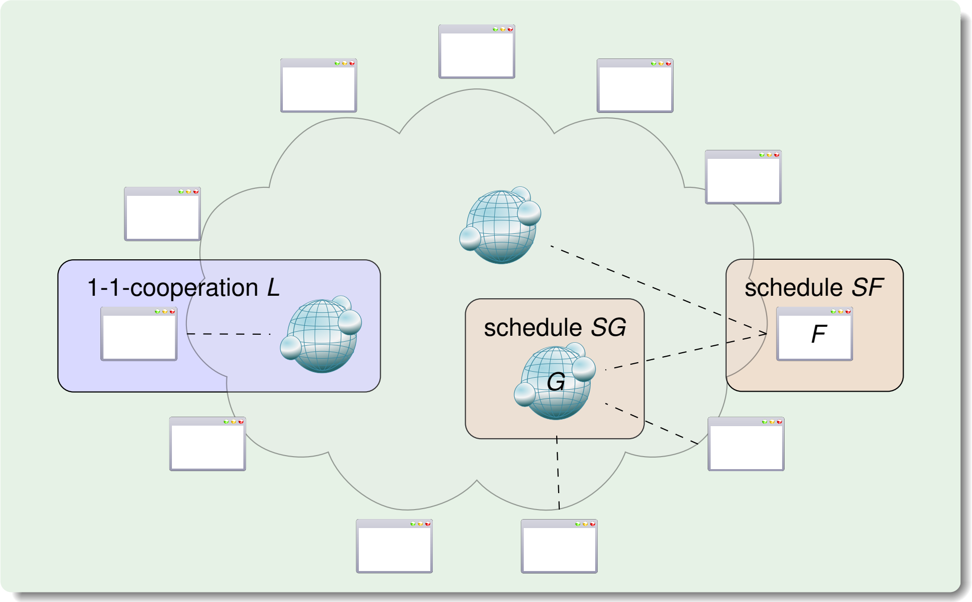

Prefix closed language *L* formally defines a two-sided cooperation

Example: Automaton for an iterated handshake of *F* and *G*

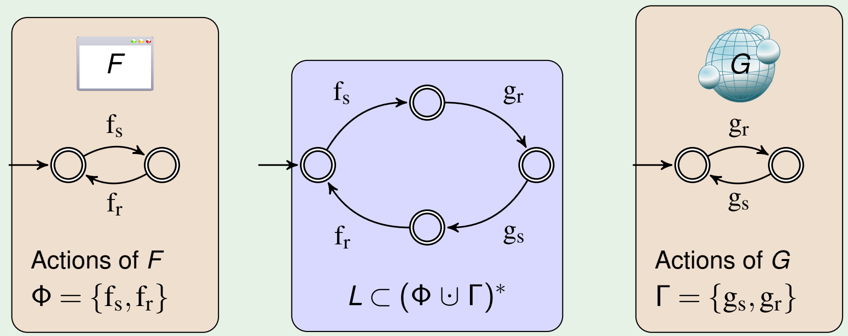

$\mathscr{L}_{IK} \subset \Sigma^*_{IK}$ describes a *parameterised cooperation*

Example: Automaton for the 1-2-cooperation $\mathscr{L}_{\{1\}\{1,2\}}$

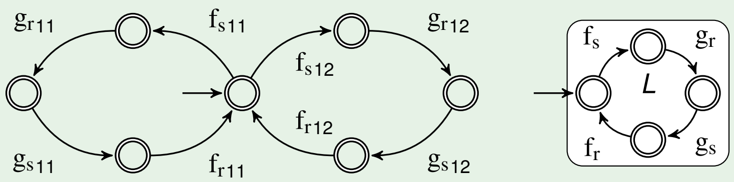

Each pair of partners cooperate restricted by *L* and each partner has to finish the handshake it just is involved in, before entering a new one.

For parameter sets *I*, *K* and $(i,k) \in I \times K$ let Σ_{ik} denote pairwise disjoint copies of Σ.

The elements of Σ_{ik} are denoted by a_{ik} and $\Sigma_{IK} := \dot{\bigcup}_{(i,k)\in I\times K} \Sigma_{ik}$.

State space explosion problem

Example: Automaton for the 2-1-cooperation $\mathscr{L}_{\{1,2\}\{1\}}$

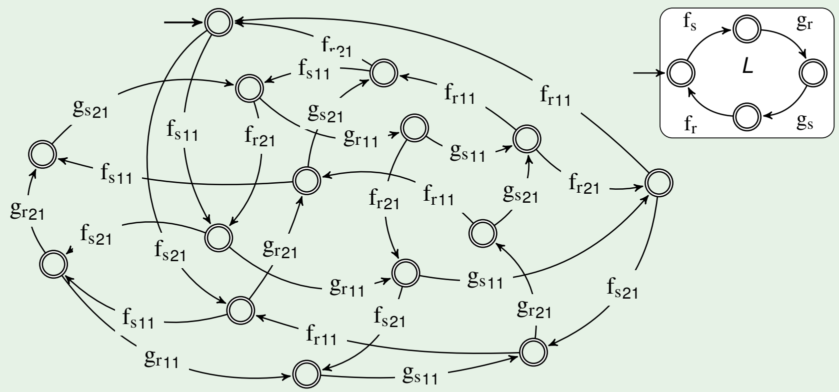

A 3-3-cooperation with the same simple behaviour of partners already requires an automaton with 916 states and 3168 state transitions.

Each pair of partners cooperate restricted by L

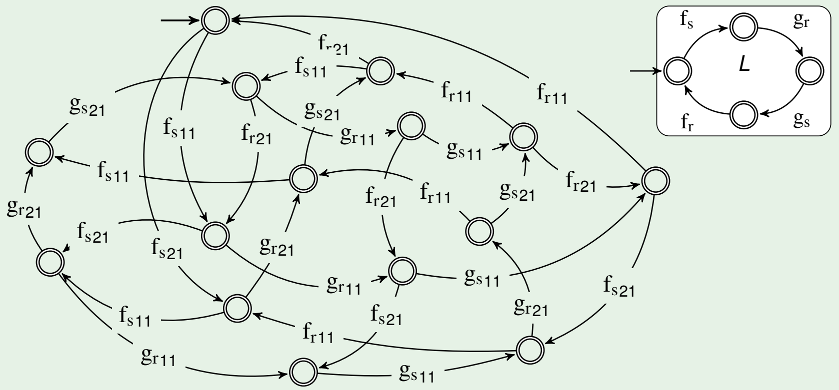

For uniformly parameterised systems $\mathscr{L}_{IK}$ we generally want to have

$$\mathscr{L}_{IK} \subset \bigcap_{(i,k)\in I\times K} ((\pi^{IK}_{ik})^{-1}(L)), \text{ with } \pi^{IK}_{ik}(a_{rs}) = \begin{cases} a \mid & a_{rs} \in \Sigma_{ik} \\ \varepsilon \mid & a_{rs} \in \Sigma_{IK} \setminus \Sigma_{ik} \end{cases}$$

Automaton for the 2-1-cooperation $\mathscr{L}_{\{1,2\}\{1\}}$

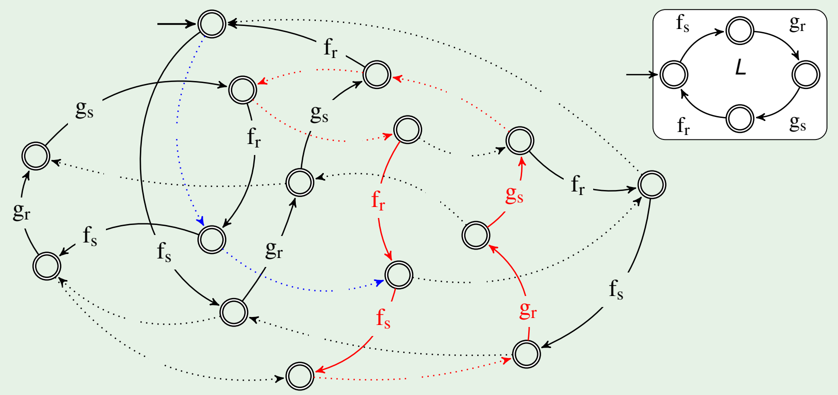

$$\pi^{IK}_{ik}(a_{rs}) = \begin{cases} a \mid & a_{rs} \in \Sigma_{ik} \\ \varepsilon \mid & a_{rs} \in \Sigma_{IK} \setminus \Sigma_{ik} \end{cases}, \text{ with } i = 2, k = 1$$

Restriction by local schedule SG

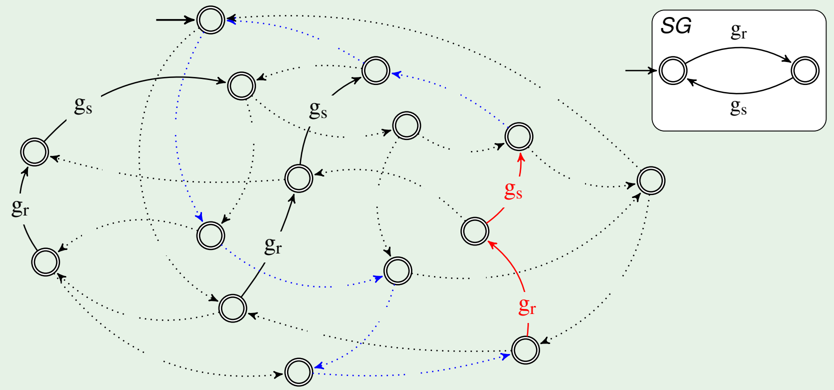

Local schedules determine how each "*version of a partner*" can participate in "*different cooperations*".

$$\mathscr{L}_{IK} \subset \bigcap_{k\in K} (\gamma^{IK}_{k})^{-1}(SG), \text{ with } \gamma^{IK}_{k}(a_{rs}) = \begin{cases} a \mid & a_{rs} \in \Gamma_{I\{k\}} \\ \varepsilon \mid & a_{rs} \in \Sigma_{IK} \setminus \Gamma_{I\{k\}} \end{cases}$$

Uniformly Parameterised Cooperation

- Each pair of partners cooperates in the same manner.
- The mechanism (schedule) to determine how one partner may be involved in several cooperations, is the same for each partner.

Definition (uniformly parameterised cooperation $\mathscr{L}_{IK}$)

Let I, K be finite parameter sets and $\pi_\Phi(L) \subset SF$, $\pi_\Gamma(L) \subset SG$, then

$$\mathscr{L}_{IK} := \bigcap_{(i,k)\in I\times K} (\pi^{IK}_{ik})^{-1}(L) \cap \bigcap_{i\in I} (\varphi^{IK}_{i})^{-1}(SF) \cap \bigcap_{k\in K} (\gamma^{IK}_{k})^{-1}(SG)$$

$\pi_\Phi : \Sigma^* \to \Phi^*$ and $\pi_\Gamma : \Sigma^* \to \Gamma^*$ are defined by

$$\pi_\Phi(a) = \begin{cases} a \mid & a \in \Phi \\ \varepsilon \mid & a \in \Gamma \end{cases} \text{ and } \pi_\Gamma(a) = \begin{cases} a \mid & a \in \Gamma \\ \varepsilon \mid & a \in \Phi \end{cases}.$$

Remark: $\mathscr{L}_{\{1\}\{1\}}$ isomorphic to L.

Self-similarity

Abstracting point of view

- Only actions of some selected partners $\Sigma_{I'K'}$ are considered.
- The complex system $\mathscr{L}_{IK}$ of all partners behaves like the smaller subsystem $\Sigma_{I'K'}$ of selected partners.

Definition (Self-similarity)

A uniformly parameterised cooperation $\mathscr{L}_{IK}$ is *self-similar* iff

$$\Pi^{IK}_{I'K'}(\mathscr{L}_{IK}) = \mathscr{L}_{I'K'} \text{ for each } I' \times K' \subset I \times K,$$

$$\text{where } \Pi^{IK}_{I'K'}(a_{rs}) = \begin{cases} a_{rs} \mid & a_{rs} \in \Sigma_{I'K'} \\ \varepsilon \mid & a_{rs} \in \Sigma_{IK} \setminus \Sigma_{I'K'}. \end{cases}$$

Counterexample $\Pi^{IK}_{I'K'}(\mathscr{L}_{IK}) \not\subset \mathscr{L}_{I'K'}$

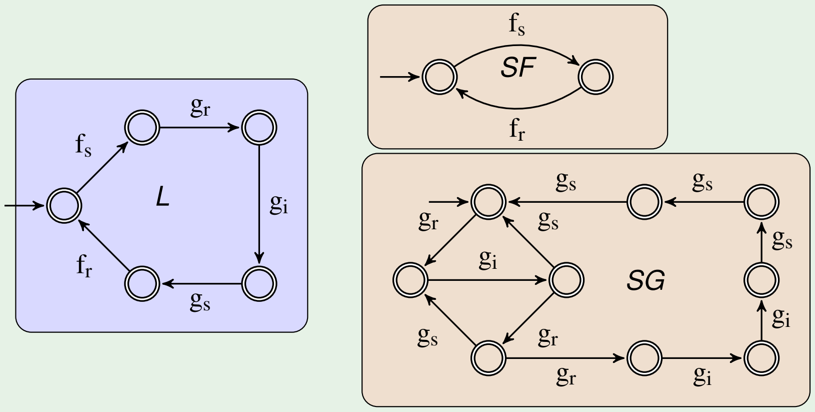

$$\cancel{\mathrm{f}_{s11}}\mathrm{f}_{s21}\mathrm{f}_{s31}\cancel{\mathrm{g}_{r11}}\cancel{\mathrm{g}_{i11}}\mathrm{g}_{r21}\mathrm{g}_{r31} \in \mathscr{L}_{\{1,2,3\}\{1\}}$$

$$\text{Hence } \mathrm{f}_{s21}\mathrm{f}_{s31}\mathrm{g}_{r21}\mathrm{g}_{r31} \in \Pi^{\{1,2,3\}\{1\}}_{\{2,3\}\{1\}}(\mathscr{L}_{\{1,2,3\}\{1\}})$$

$$\text{but } \mathrm{f}_{s21}\mathrm{f}_{s31}\mathrm{g}_{r21}\mathrm{g}_{r31} \notin \mathscr{L}_{\{2,3\}\{1\}}$$

System Properties

A property E of a system is a subset of Σ^ω. If $S \subset \Sigma^\omega$ represents the behaviour of a system, then S *linearly satisfies* E iff $S \subset E$.

Alpern/Schneider: Each property E is the intersection of a safety and a liveness property

Safety properties $E_s \subset \Sigma^\omega$ are of the form $E_s = \Sigma^\omega \setminus F\Sigma^\omega$ with $F \subset \Sigma^*$, where F is the set of "bad things".
Liveness properties $E_l \subset \Sigma^\omega$ are characterised by $\mathrm{pre}(E_l) = \Sigma^*$.

Reliability: Typical example of a liveness property

$$E_l = (\Sigma^* M)^\omega \text{ with } \emptyset \neq M \subset \Sigma^+. \quad (1)$$

"always eventually a finite action sequence $m \in M$ happens"

A word u is called a *prefix* of a word v if there is a word x such that $v = ux$.
The set of all prefixes of a word u is denoted by $\mathrm{pre}(u)$

Approximate Satisfaction

Linear satisfaction is too strong w.r.t. to liveness properties, because $S = \lim(\hat{B})$ can contain "unfair" infinite behaviours, which are not in E.

If in a 2-1-cooperation infinite action sequences exist where only the partners with index 1 cooperate, i.e., $\lim(\widehat{\mathscr{L}_{IK}}) \cap \Sigma^\omega_{\{1\}\{1\}} \neq \emptyset$,
then e.g. $E = \Sigma^*_{IK}\Sigma_{\{2\}\{1\}}\Sigma^\omega_{IK}$ is not linearly satisfied; $\lim(\widehat{\mathscr{L}_{IK}}) \not\subset E$.

A weaker satisfaction relation implicitly expresses a kind of fairness.

A system $S \subset \hat{\Sigma}^\omega$ *approximately satisfies* a property $E \subset \hat{\Sigma}^\omega$ iff each finite behaviour (finite prefix of an element of S) can be continued to an infinite behaviour, which belongs to E, i.e. $\mathrm{pre}(S) \subset \mathrm{pre}(S \cap E)$.

With respect to approximate satisfaction, liveness properties stipulate that "something good" eventually is possible.

For safety properties linear and approximate satisfaction are equivalent.

Property Preserving Abstractions

Theorem (1)

Simple homomorphisms define exactly the class of such abstractions, for which holds that each property is approximately satisfied by the abstract behaviour if and only if the “corresponding” property is approximately satisfied by the concrete behaviour of the system.

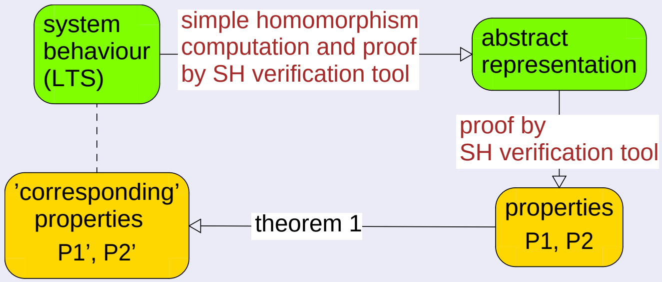

A system *approximately satisfies* a property if and only if each finite behaviour can be continued to an infinite behaviour, which satisfies the property.

Property Preserving Abstractions (Counterexample)

Abstract representations may hide resticted behaviour

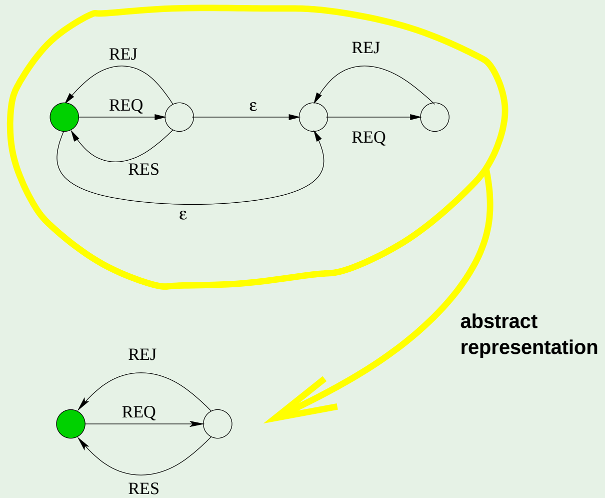

The property $\mathbb{G}(\mathbb{F}(RES))$ is approximately satisfied in the abstract representation but NOT in the concrete system.

Combining the introduced Concepts

Theorem

Let I, K, $\mathring{I}$ and $\mathring{K}$ be finite index sets with $|\mathring{I}| \leq |I|$ and $|\mathring{K}| \leq |K|$.
Let $\mathscr{L}_{IK}$ be a

- *uniformly parameterised,*
- *self-similar regular system of cooperations, and*
- *the abtracting view $\Pi^{IK}_{I'K'}$ on actions of selected partners I', K' be simple on $\mathscr{L}_{IK}$ for each $I' \subset I$, $K' \subset K$ with $|\mathring{I}| = |I'|$, $|\mathring{K}| = |K'|$.*

If an $\boxed{\text{abstract system}}$ $\lim(\widehat{\mathscr{L}_{\mathring{I}\mathring{K}}})$

approximately satisfies a property $\mathring{E} \subset \widehat{\Sigma_{\mathring{I}\mathring{K}}}^{\omega}$,

then the $\boxed{\text{concrete system}}$ $\lim(\widehat{\mathscr{L}_{IK}})$

approximately satisfies the “corresponding” family of properties $\mathscr{E}^{\mathring{E}}_{IK}$.

Limit of prefix closed languages $\lim(B) := \{w \in \Sigma^{\omega} | \mathrm{pre}(w) \subset B\}$.
In $\hat{B}$ the maximal words of B are continued by arbitrary many $\# \notin \Sigma$.

Conclusions

Previous work

- The parameterised problem of verifying a uniformly parameterised safety property can be reduced to finitely many finite state problems.

Main result of the presented work

- A formal framework for uniformly parameterised behaviour properties capturing the full spectrum of safety and liveness.
- Uniformly parameterisation of behaviour properties fits to the reliability issues of scalable systems, e.g., cloud computing.
- A combination of these properties can now be used to specify security requirements for such kinds of systems.

Further Work

- Concept to prove simplicity of $\Pi^{IK}_{I'K'}$ on $\mathscr{L}_{IK}$
- Construction principles for uniformly parameterised self-similar systems

Roland Rieke developed the work presented here in the context of the project MASSIF (ID 257475) being co-funded by the European Commission within FP7.