Abstract (for the general public):

Frontiers of automatic analysis of concurrent systems

Principal Investigator: Sławomir Lasota

Motivation. Concurrent systems are ubiquitous, and the impact of their reliability is continuously increasing. In consequence, we observe an increasing call for improving the quality of concurrent systems. It is well known that concurrent system designs are prone to subtle bugs that are inherently difficult to find by humans, and examples of costly damages caused by such bugs abound. One way of achieving improvement of quality of such systems is formal verification, i.e., automatic analysis methods.

Goal. We are going to push further theoretical foundations of automatic analysis of models of concurrent systems, by investigating the frontiers of such analysis, and to enhance its practical applicability. We will specifically concentrate on the model of Petri nets, and investigate questions like *coverability* or *reachability* of a given configuration.

Description of the project. Petri nets consist of *places* representing resources or control states of processes, and *transitions* representing their dynamical behaviour. For illustration, consider a Petri net with two places $P = \{p_1, p_2\}$ and two transitions $T = \{t_1, t_2\}$, as depicted below, together with its configuration that stores three tokens on place p_1 and two tokens on p_2:

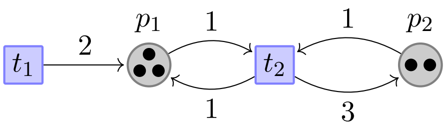

Every execution of transition t_1 puts additional two tokens on place p_1, while every execution of t_2 takes one token from p_1 and one from p_2, and then puts back one token on place p_1 and three tokens on place p_2. Here is an example *coverability* question: is there a run (i.e., a sequence of executions of transitions) starting in the given configuration which puts *at least* four tokens on p_1 and at least six tokens on p_2? Here is an example *reachability* question: is there a run that puts *exactly* four tokens on p_1 and exactly six tokens on p_2? In our example, coverability holds (in fact, both places p_1 and p_2 are *unbounded*, i.e., there is no finite bound on the number of tokens on these places) but reachability does not. Indeed, the parity of the number of tokens on both places is preserved by both transitions.

The coverability problem is a key for safety verification of concurrent systems. On the other hand, the reachability problem is of key theoretical importance, since a number of problems from formal languages, logic, concurrent systems, process calculi and other areas, are known to admit reductions to/from the reachability problem. In consequence, establishing its exact complexity is considered as the central challenge in formal verification of concurrent systems. We note a recent breakthrough: an improvement of the lower bound for the reachability problem, for the first time since 40 years, the result co-authored by the PI of this project. The current project aims at pushing further this new and fascinating line of research.

Expected results. The expected results are twofold. On one hand, we hope to successfully solve (at least special cases of) a number of extremely challenging open theoretical problems, like decidability status of Petri nets with stack, or exact complexity of the Petri nets reachability problem in fixed dimension. On the other hand, we would like to investigate subclasses of the general model, hoping for identifying tractable cases applicable in practical verification.