Model Checking Synchronized Products of Infinite Transition Systems

Stefan Wöhrle and Wolfgang Thomas
Lehrstuhl für Informatik 7
RWTH Aachen, Germany
{woehrle,thomas}@informatik.rwth-aachen.de

Abstract

Formal verification using the model-checking paradigm has to deal with two aspects. The systems models are structured, often as products of components, and the specification logic has to be expressive enough to allow the formalization of reachability properties. The present paper is a study on what can be achieved for infinite transition systems under these premises. As models we consider products of infinite transition systems with different synchronization constraints. We introduce finitely synchronized transition systems, i.e. product systems which contain only finitely many synchronized transitions, and show that the decidability of FO(R), first-order logic extended by reachability predicates, of the product system can be reduced to the decidability of FO(R) of the components in a Feferman-Vaught like style. This result is optimal in the following sense. (1) If we allow semifinite synchronization, i.e. just in one component infinitely many transitions are synchronized, the FO(R)-theory of the product system is in general undecidable. (2) We cannot extend the expressive power of the logic under consideration. Already a weak extension of first-order logic with transitive closure, where we restrict the transitive closure operators to arity one and nesting depth two, is undecidable for an asynchronous (and hence finitely synchronized) product, namely for the infinite grid.

1. Introduction

In the theory of algorithmic verification, the classical framework for modeling systems is given by finite transition systems (often in the form of Kripke structures). Much effort is presently spent on extending this framework to cover infinite transition systems, and to deal adequately with the internal structure of the models under consideration, such as their composition from several components. The present paper is a study on the scope of model-checking over models which are composed from infinite components as products with various constraints on the synchronization of their transitions.

The synchronized product of transition systems is a crucial operation [1, 2] to model concurrency and interaction between processes. Synchronized products appear for example in the calculus of communicating systems CCS [19] or the communicating sequential processes CSP [13].

A natural approach for model-checking over products is to infer the truth-values of statements in the product structure from information about the components (factors). In model theory, such an inference method is provided by the celebrated Fefeman-Vaught Theorem [9]. For direct products $\mathcal{A} = \Pi_{i\in I}\mathcal{A}_i$ of structures $\mathcal{A}_i$ (respectively reduced products $\mathcal{A} = \Pi_D\mathcal{A}_i$ where D is a filter over I [6]), it allows to determine whether a sentence φ is true in $\mathcal{A}$ from the evaluation of (effectively computable) sentences $\psi_1, \ldots, \psi_k$ in the components $\mathcal{A}_i$ and a condition σ on the sets $I_j = \{i \in I \mid \mathcal{A}_i \models \psi_j\}$. Thus the Feferman-Vaught Theorem allows to compute the FO-theory of the product from the FO-theories of the component structures if the condition σ can be decided (which trivially holds for finite I). In this way a solution of the FO model-checking problem can be transferred from the components to the product model.

Many variants and extensions of the Feferman-Vaught technique have been developed. Already in [9] the composition by sums rather than products was considered (taking as universe the disjoint union of the component universes). Shelah and Gurevich extended the framework to monadic second-order logic, applying it to ordered sums of orderings and sums of trees (see [25, 11]). Generalized versions of sum operations were studied by Shelah in [26] and Rabinovich in [23]. A unified treatment and survey of Feferman-Vaught like results is given by Makowsky in [18].

For purposes of algorithmic verification, the direct and reduced products considered in the literature do not match the needs of modeling composition of computational systems. In the case of transition graphs, a direct product of two components has an a-labeled transition from (p, q) to (p', q') if there are such transitions from p to p' and from q

to q'. This is the case of complete synchronization. Other versions of synchronization have to be considered, including the extreme case of an asynchronous product, where a transition in one component does not affect the other components.

Regarding the logical framework, a basic requirement is that reachability properties have to be expressible. Among the numerous ways to incorporate these properties in the logic we consider four basic options of increasing expressiveness:

- Reachability logic FO(R), which is obtained from FO-logic by adjoining the transitive closure Reach_Γ for a subset of edge relations $\{E_a \mid a \in \Gamma\}$.
- FO(Reg) as a generalization of FO(R) in which path labels have to match a given regular expression.
- Transitive closure logic over binary relations, which allows to proceed from any definable relation (and not just from some edge relations) to its transitive closure.
- Monadic second-order logic MSO, which results from FO-logic by adjoining variables and quantifiers for sets (and in which transitive closure over binary relations can be expressed).

The purpose of this paper is to work out cases where the Feferman-Vaught technique is applicable over products of infinite transition systems with partial synchronization constraints, and where the logic includes reachability properties.

We introduce product structures which have possibly infinite components but are finitely synchronized. This situation applies to any system where the local computations may involve infinite state-spaces but where synchronization is possible only via finitely many transitions. We show a Feferman-Vaught result for FO(R)-logic over such product structures: The decidability of the FO(R)-theory of transition systems is preserved under forming finitely synchronized products. The solution is uniform in the sense that it transforms given model-checking algorithms for the components into a model-checking procedure for the product.

This extends a result of Rabinovich [24] on propositional modal logic extended by the modality EF over asynchronous products.

We show that our result is optimal in two ways.

Firstly, the result does not extend to a case where we allow a slight liberalization of the constraint on finite synchronization: We consider "semi-finite synchronization", in which all components except one can synchronize via finitely many transitions. In the presence of a single component with infinitely many synchronizing transitions we may obtain a structure with undecidable FO(R) model-checking problem, whereas the problem is decidable for the components individually.

Secondly, we investigate whether the FO(R)-logic can be extended in the above-mentioned preservation result. For a strong extension like MSO-logic it is clear that decidability of the component theories does not carry over to the theory of the product system. As is well-known, we may work with the asynchronous product of the successor structure of the natural numbers, which is the infinite $(\omega \times \omega)$-grid. (Note that the asynchronous product is finitely synchronized with an empty set of synchronizing transitions.) The grid has an undecidable monadic theory, whereas the component structures have decidable monadic theories.

We clarify the situation for weaker extensions of FO(R)-logic, namely FO(Reg) and transitive closure logic. We show that asynchronous products do not preserve the decidability of the FO(Reg)-theory. For transitive closure logic this undecidability result can already be obtained for a very simple example of an asynchronous product, namely the infinite grid as considered above. Furthermore, we show that this undecidability phenomenon only appears when the TC-operator is nested. For the fragment of transitive closure logic with unnested TC-operators interpreted over the infinite grid, we obtain a reduction to Presburger arithmetic and hence the decidability of the corresponding theory.

These undecidability results complement a theorem of Rabinovich [24] where the corresponding fact is shown for propositional modal logic extended by the modality EG over finite grids.

In our results the component structures are assumed to have a decidable theory in one of the logics considered above. Let us summarize some of the relevant classes and their closure properties with respect to synchronization.

A fundamental result is that pushdown graphs have a decidable monadic second-order theory [21]. Since then several extensions like prefix recognizable graphs [4] or Caucal graphs [5] have been considered, see [28] for an overview. These classes form an increasing sequence in this order, and all of them enjoy a decidable MSO-theory. None of these classes is closed under asynchronous products.

Two other classes of infinite graphs we like to mention are the graphs of ground term rewriting systems [7] for which the FO(R)-theory is decidable, and ground tree rewriting systems [16] for which a temporal logic with reachability and recurrence operators is decidable. Both classes are closed under asynchronous products.

Classes which are closed under synchronized products are rational graphs [20], graphs of Thue specifications [22], or graphs of linear bounded machines [15]. However for all these classes already the FO-theory is undecidable and hence they are not suitable for model-checking purposes.

The paper is organized as follows. In Section 2 we give the definition of a synchronized product of a family

of graphs or transition systems, recall the definition of transitive closure logic, and define FO(R) and FO(Reg).

In Section 3 we show the composition theorem for finitely synchronized products and reachability logic and prove that this result cannot be extended to FO(Reg) or semifinite synchronization in general.

In Section 4 we investigate transitive closure logic over the infinite grid. We show that if we allow transitive closure operators of arity one without parameters but of nesting depth two the theory of the grid is undecidable. On the other hand we show that if no nesting of transitive closure operators is allowed, the respective theory is decidable even in presence of parameters in the scope of the transitive closure operators.

2. Preliminaries

Let $(V_i)_{1\leq i\leq n}$ be a family of sets. We denote by $\Pi_{1\leq i\leq n}V_i$ the Cartesian product of these sets. Tuples $v_1,\dots,v_n \in \Pi_{1\leq i\leq n}V_i$ are usually denoted by $\bar{v}$, and the ith component of $\bar{v}$ as v_i .

Let Σ be a finite set of labels. A *transition system* is a Σ-labeled directed graph $G = (V^G, (E^G_a)_{a\in\Sigma})$ where V^G is the set of vertices of G and $E^G_a \subseteq V^G \times V^G$ denotes the set of a-labeled edges in G.

2.1. Synchronized Products

For $1 \leq i \leq n$ let $G_i := (V_i, (E^i_a)_{a\in\Sigma_i})$ be a Σ_i-labeled graph. We assume that Σ_i is partitioned into a set Σ^l_i of *local* labels (or actions) and a set Σ^s_i of *synchronizing* labels, and to avoid notational complication we require the sets of local labels to be pairwise disjoint. Let $E^i_\varepsilon := \{(v,v) \mid v \in V_i\}$ and $\tilde{\Sigma}^s_i := \Sigma^s_i \cup \{\varepsilon\}$. A *synchronization constraint* is a set $C \subseteq \Pi_{1\leq i\leq n}\tilde{\Sigma}^s_i$.

The *synchronized product* of $(G_i)_{1\leq i\leq n}$ defined by C is the graph G with vertex set $V := \Pi_{1\leq i\leq n}V_i$, asynchronous transitions with labels $a \in \bigcup_{1\leq i\leq n}\Sigma^l_i$ defined by $E^G_a\bar{v}\bar{w}$ if $E^i_a v_i w_i$ and $v_j = w_j$ for $j \neq i$, and synchronized transitions with labels $\bar{c} \in C$ defined by $E^G_{\bar{c}}\bar{v}\bar{w}$ if $E^i_{c_i}v_iw_i$ for every $1 \leq i \leq n$.

Note that we slightly deviate from the definition in [1] since we require the sets of local labels and synchronizing labels to be disjoint, and implicitly assume an asynchronous behavior of local transitions.

If $E^G_{\bar{c}}$ is finite for every $\bar{c} \in C$ we call G a *finitely synchronized product*. Note that G cannot be finitely synchronized if for some $\bar{c} \in C$ with $c_i = \varepsilon$ the set V_i is infinite.

2.2. First-Order Logic and Extensions

We assume that the reader is familiar with first-order logic FO over graphs. We denote formulas by $\varphi(x_1,\dots,x_n)$ to express that the free variables of φ are among $x_1,\dots,x_n$. If G is a graph and $v_1,\dots,v_n$ are the vertices assigned to the variables $x_1,\dots,x_n$, we denote by $(G, v_1,\dots,v_n) \models \varphi(x_1,\dots,x_n)$ or shortly by $G \models \varphi[v_1,\dots,v_n]$ that the formula φ is satisfied in G under the respective variable assignment.

Transitive closure logic FO(TC) is defined by extending FO with formulas of the type

$$\psi := [\mathrm{TC}_{\bar{x},\bar{y}}\,\varphi(\bar{x},\bar{y},\bar{z})]\,\bar{s},\bar{t}$$

where $\varphi(\bar{x},\bar{y},\bar{z})$ is a FO(TC)-formula, $\bar{x},\bar{y}$ are disjoint tuples of free variables of the same length $k > 0$, $\bar{s},\bar{t}$ are tuples of variables of length k and $\mathrm{free}(\psi) := (\mathrm{free}(\varphi) \setminus \{\bar{x},\bar{y}\}) \cup \{\bar{s},\bar{t}\}$. Note that in the notation $[\mathrm{TC}_{\bar{x},\bar{y}}\,\varphi(\bar{x},\bar{y},\bar{z})]\,\bar{x},\bar{y}$ the variables inside the square brackets are bound while the variables at the end of the formula occur free.

Let G be a graph, let $\bar{c}$, $\bar{d}$, and $\bar{e}$ be the interpretations of the variables $\bar{z}$, $\bar{s}$, and $\bar{t}$ in φ. Let E be the relation on k-tuples defined by $E(\bar{c}) := \{(\bar{a},\bar{b}) \mid (G,\bar{a},\bar{b},\bar{c}) \models \varphi(\bar{x},\bar{y},\bar{z})\}$, and $E'(\bar{c})$ be its transitive closure, i.e. $(\bar{a},\bar{b}) \in E'(\bar{c})$ iff there exists a sequence $\bar{f}_0,\bar{f}_1,\dots,\bar{f}_l$ such that $\bar{f}_0 = \bar{a}$, $(\bar{f}_i,\bar{f}_{i+1}) \in E(\bar{c})$ for $1 \leq i < l$, and $\bar{f}_l = \bar{b}$. The semantics of the FO(TC)-formula above is defined by

$$(G,\bar{c},\bar{d},\bar{e}) \models [\mathrm{TC}_{\bar{x},\bar{y}}\,\varphi(\bar{x},\bar{y},\bar{z})]\,\bar{s},\bar{t} \Leftrightarrow (\bar{d},\bar{e}) \in E'(\bar{c}).$$

We call the variables $\bar{z}$ parameters for the transitive closure operator. By $\mathrm{FO(TC)}_{(k)}$ be denote the fragment of FO(TC) where the transitive closure operation is only allowed to define relations over tuples of length $\leq k$, i.e. the length of the tuples $\bar{x},\bar{y}$ in the definition above is bounded by k. For example, in $\mathrm{FO(TC)}_{(1)}$ we can only define binary relations using a transitive closure operator. For finite models the arity hierarchy $(\mathrm{FO(TC)}_{(k)})_{k\geq 0}$ is strict [10].

By $\mathrm{FO(TC)}^l_{(k)}$ we denote the fragment of $\mathrm{FO(TC)}_{(k)}$ where the nesting depth of transitive closure operations is bounded by l.

In transitive closure logic we can express that from a vertex x a vertex y is reachable via a path with labels from some set $\Sigma' \subseteq \Sigma$ by

$$\mathrm{Reach}_{\Sigma'}(x,y) := \Big[\,\mathrm{TC}_{x,y}\,\big(x = y \vee \bigvee_{a\in\Sigma'} E_a xy\big)\Big]\,x,y.$$

We call the restriction of FO(TC) where the only transitive closure formulas allowed are of the form $\mathrm{Reach}_{\Sigma'}(x,y)$ for $\Sigma' \subseteq \Sigma$ *reachability logic* and denote it by FO(R).

The expressive power of the reachability predicates in FO(R) is limited, e.g. we cannot express that there is a path between vertex v and w in the graph whose labels form a word in a given regular language.

We denote by FO(Reg) first-oder logic extended by reachability predicates $\mathrm{Reach}_r(x,y)$ for regular expressions r over Σ, where $G \models \mathrm{Reach}_r[v,w]$ if there is a path in

G from v to w labeled by a word contained in the language described by r.

3. Synchronization and FO(R)

In this section we show that synchronization preserves the decidability of the FO(R)-theory if (and only if) the product is finitely synchronized. For this case we prove a composition theorem that reduces the evaluation of a formula in the product graph to the evaluation of several formulas in the component graphs and a Boolean combination of these truth values. This result does not extent to the case of FO(Reg).

Furthermore we show that *semifinite* synchronization of two components, where in just one of the components infinitely many edges are allowed to be synchronized, does in general not preserve the decidability of the FO(R)-theory.

Theorem 1 *Let G be a finitely synchronized product of a family $(G_i)_{1\leq i\leq n}$ of graphs with decidable FO(R)-theories. Then the FO(R)-theory of G is also decidable, and for an FO(R)-formula φ we can effectively construct sets of formulas Ψ_i and a Boolean formula α such that $G \models \varphi$ iff α is true under an Boolean interpretation defined by the truth values of the formulas in Ψ_i.*

Proof: Let $(G_i)_{1\leq i\leq n}$ be a family of graphs whose signatures $\Sigma_i := \Sigma_i^l \cup \Sigma_i^s$ are partitioned into local and synchronizing labels. Let $C \subseteq \Pi_{1\leq i\leq n}\tilde{\Sigma}_i^s$ be a synchronization constraint such that the product G of $(G_i)_{1\leq i\leq n}$ is finitely synchronized with respect to C.

We show by induction that for every FO(R)-formula over $\bigcup_{1\leq i\leq n}\Sigma_i^l \cup C$ there are finite sets Ψ_i of Σ_i-formulas and a Boolean formula α over predicates $p_i(\psi_j^i)$ $(1 \leq i \leq n,\ 1 \leq j \leq |\Psi_i|)$ such that

$$(G, \bar{v}_1, \dots \bar{v}_m) \models \varphi(x_1, \dots, x_m) \Leftrightarrow I \models \alpha \qquad (1)$$

where I is the Boolean interpretation defined by

$$I(p_i(\psi_j)) = \begin{cases} \text{true} & \text{if } (G_i, v_1^i, \dots, v_m^i) \models \psi_j^i \\ \text{false} & \text{otherwise.} \end{cases}$$

We start with the atomic formulas. For $x = y$ let $\psi_i := (x = y)$, for $E_a xy$ with $a \in \Sigma_i^l$ let $\psi_i := E_a xy$ and $\psi_j := (x = y)$ for $i \neq j$, and for $E_{\bar{c}}xy$ with $\bar{c} \in C$ let $\psi_i := E_{c_i}xy$. For every formula above let $\alpha := \bigwedge_{1\leq i\leq n} p_i(\psi_i)$. Obviously (1) holds in all cases, so the remaining "atomic" formulas we have to take care of are of the form $\mathrm{Reach}_\Gamma(x, y)$ for $\Gamma \subseteq \bigcup_{1\leq i\leq n}\Sigma_i^l \cup C$. Since G is finitely synchronized the sets $\{(\bar{v}, \bar{w}) \mid E_{\bar{c}}^G \bar{v}\bar{w}\}$ are finite for every $\bar{c} \in C$, say of size $l(\bar{c})$. Let $\bar{c}_1, \dots, \bar{c}_k$ be an enumeration of C, and $S := \{\bar{c}_1^1, \dots, \bar{c}_1^{l(\bar{c}_1)}, \dots, \bar{c}_k^1, \dots, \bar{c}_k^{l(\bar{c}_k)}\}$ where $\bar{c}_j^i := \bar{c}_j$ for $1 \leq i \leq l(\bar{c}_j)$ and $1 \leq j \leq k$. For $R \subseteq S$ of size r let $\mathrm{Perm}(R)$ denote the set of all permutations $\pi : \{1, \dots, r\} \to R$.

For $c \in \Sigma_i^s$ we define the formula

$$\chi_c^i(x, y) := \exists z_1 \big(E_c x z_1 \wedge \mathrm{Reach}_{\Gamma\cap\Sigma_i^l}(z_1, y)\big)$$

which expresses that x is the source of a synchronizing transition c and y can be reached using such a transition from x and only local ones from Γ afterwards. Now for every $R \subseteq S$ of size r and every $\pi \in \mathrm{Perm}(R)$ we define a formula

$$\theta_{R,\pi}^i := \exists z_1, \dots, z_r \Big(\mathrm{Reach}_{\Gamma\cap\Sigma_i^l}(x, z_1) \wedge \bigwedge_{1\leq j\leq r} \chi_{\pi(j)_i}^i(z_j, z_{j+1}) \wedge z_r = y \Big)$$

which expresses that there is a path from x to y which contains $\pi(1)_i, \dots, \pi(r)_i$ as subsequence of synchronized transitions. Let

$$\alpha := \bigvee_{R\subseteq S} \bigvee_{\pi\in\mathrm{Perm}(R)} \bigwedge_{1\leq i\leq n} p_i(\theta_{R,\pi}^i).$$

To prove the correctness of the construction above note that every path from x to y in G which contains a synchronizing edge twice can be shortened to one which traverses every edge only once, and that G is an asynchronous product with respect to local transitions. From these properties (1) follows immediately.

The case of Boolean connectives may be solved in the standard way. Let $\varphi_1(\bar{x})$ and $\varphi_2(\bar{y})$ be FO(R)-formulas and α_1, $(\Psi_i^1)_{1\leq i\leq n}$ as well as α_2, $(\Psi_i^2)_{1\leq i\leq n}$ be given by the induction hypothesis. Then, for $\neg\varphi_1(\bar{x})$ we can choose the same $(\Psi_i^1)_{1\leq i\leq n}$ and the Boolean formula to be $\neg\alpha_1$, and for $\varphi_1(\bar{x}) \vee \varphi_2(\bar{y})$ we choose $\Psi_i := \Psi_i^1 \cup \Psi_i^2$ and $\alpha = \alpha_1 \vee \alpha_2$.

To finish the proof let $\varphi(x_1, \dots, x_n) := \exists x_{n+1}\varphi_1(x_1, \dots, x_{n+1})$. Let Ψ_i^1 and α_1 be the formulas computed for $\varphi_1(x_1, \dots, x_{n+1})$. Let $\mathcal{I}$ be the set of all satisfying assignments for α_1. For every $I \in \mathcal{I}$ let $I_i := \{j \mid I(p_i^j) = \text{true}\}$. Then sets Ψ_i for $1 \leq i \leq n$ are constructed by adding for every $I \in \mathcal{I}$ the formula

$$\psi_i^I(x_1, \dots, x_n) := \exists x_{n+1} \Big(\bigwedge_{j\in I_i} \psi_i^j \wedge \bigwedge_{j\notin I_i} \neg\psi_i^j \Big).$$

Then we can define $\alpha := \bigvee_{I\in\mathcal{I}} \bigwedge_{1\leq i\leq n} p(\psi_i^I)$. □

For a complexity analysis of this algorithm we can estimate the number of formulas that have to be evaluated for every component graph to obtain an analysis independent of the complexity of the evaluation algorithms for the components.

The atomic subformulas $\mathrm{Reach}_\Gamma(x,y)$ require to evaluate $2^{(s^2)}$ formulas in each component where s is the size of the set of synchronizing transitions in the product graph. Note that this set is the only information about the product graph we need to evaluate a formula, and apart from this the data complexity is the same as data complexity of the algorithms for the components. Compared to the size of the representations of the infinite components we expect the number of synchronizations s to be small.

However, the main contribution to the overall complexity of the algorithm is the exponential increase of the sets Ψ_i which comes from the quantification step and results in an overall non-elementary formula complexity.

It is easy to see that Theorem 1 also covers FO(Reg)-formulas with regular expressions built from Γ_i^* for $\Gamma_i \subseteq \Sigma$ using $\cdot$ and $+$. However, if we allow reachability predicates with regular expressions of the form $(\Gamma_1 \cdot \Gamma_2)^*$ the decidability of the corresponding theory will be lost.

Theorem 2 *Asynchronous products do not preserve the decidability of the FO(Reg)-theory.*

Proof: We reduce the reachability problem for 2-PDAs (pushdown automata with two stacks), a problem which is well known to be undecidable [14].

A 2-PDA is a tuple $\mathcal{A} = (Q, \Sigma, \Gamma, q_0, \Delta, f)$ where Q is a finite set of states, q_0 is the initial state, f is the final state, and $\Delta \subseteq Q \times \Sigma \times \Gamma^2 \times \Gamma^2 \times Q$. We split $\mathcal{A}$ into two component pushdown automata

$$\mathcal{A}_1 = (Q, \Sigma \times \Delta, \Gamma, q_0, \Delta_1, f)$$
$$\mathcal{A}_2 = (Q, \bar{\Sigma} \times \bar{\Delta}, \Gamma, q_0, \Delta_2, \bar{f})$$

where for every $\delta = (q, a, \gamma_1, \gamma_2, \gamma_3, \gamma_4, p) \in \Delta$ the following transitions are added

$$(q, (a, \delta), \gamma_1, \gamma_3, p) \text{ to } \Delta_1$$
$$(q, (\bar{a}, \bar{\delta}), \gamma_2, \gamma_4, p) \text{ to } \Delta_2.$$

The graphs generated by $\mathcal{A}_1$ and $\mathcal{A}_2$ have a decidable FO(Reg)-theory (since even the MSO-theory is decidable). Let $\mathcal{B}$ their asynchronous product. By construction we obtain that $\mathcal{A}$ reaches f iff

$$\mathcal{B} \models \exists z_1 z_2 \, \mathrm{Reach}_r((q_0, q_0), z_1) \wedge \bigvee_{\substack{\delta \in \Delta \\ a \in \Sigma}} E_{(a,\delta)} z_1 z_2 \wedge E_{(\bar{a},\bar{\delta})} z_2 (f, f),$$

where $r = \left(\bigvee_{\substack{\delta \in \Delta \\ a \in \Sigma}} (a, \delta)(\bar{a}, \bar{\delta}) \right)^*$ expresses that a transition of $\mathcal{A}_1$ is followed by the corresponding transition of $\mathcal{A}_2$. □

We now turn to the proof that semifinite synchronization in general does not preserve the decidability of the FO(R)-theory. We reduce the halting problem of deterministic Turing machines to the model checking problem for FO(R) for synchronized products of finite graphs and infinite graphs which are generated by ground tree rewriting systems (GTRS). The GTRS graphs we will construct are of finite out-degree and hence have a decidable FO(R)-theory [16, 17].

The GTRS graph will encode computations of the Turing machine M, but not all of them are valid. We will use the synchronization with a finite graph to eliminate computations which are not valid.

Our construction of the GTRS graph encoding computations of M follows ideas of [17]. Before we start the proof we give a short definition of the Turing machine model we use and of ground tree rewriting systems. For a more detailed description we refer to [14] and [17].

A *deterministic Turing machine* is a tuple $M = (Q, \Gamma, q_0, q_f, \delta)$ where Q is a finite set of states, Γ is an alphabet containing a designated blank symbol $\sqcup$, q_0 is the initial state, q_f is the halting state, and $\delta : Q \times \Gamma \to Q \times \Gamma \times \{L, R\}$ is the transition function. A *configuration* of M is a sequence $a_1, \ldots a_k, q, b_l, b_{l-1} \ldots b_1$ where $a_i, b_i \in \Gamma$, $q \in Q$ and b_l denotes the symbol currently read by the head of the machine. We consider two configurations to be equivalent if they differ only in heading or trailing blank symbols, and do not distinguish between equivalent configurations.

A *ground tree rewriting system* is a tuple $\mathcal{R} = (A, \Sigma, R, t_0)$ where A is a ranked alphabet, Σ is a set of labels for the rules, R is a finite set of rules, and t_0 is a finite tree over A. We denote the set of all finite trees over A by T_A. A *rewriting rule* r is of the form $t \xrightarrow{b} t'$ with $t, t' \in T_A$ and $b \in \Sigma$. A rule r is applicable to a tree s if there is a subtree s_1 of s equal to t, and the result of an application of r to s is a tree s' obtained from s by replacing s_1 with t'. $\mathcal{R}$ generates a Σ-labeled graph whose vertices are the trees that can be obtained from t_0 by applying rewriting rules from R, with a b-labeled edge between s and s' if s' results from s by an application of a rule of the form $t \xrightarrow{b} t' \in R$.

Theorem 3 *Semifinite synchronization does not preserve the decidability of the FO(R)-theory.*

Proof: Let $M = (Q, \Gamma, q_0, q_f, \delta)$ be a deterministic Turing machine. We assume that $q_0 \neq q_f$, $Q \cap \Gamma = \emptyset$, $X \notin Q \cup \Gamma$ and encode a configuration $a_1, \ldots, a_k, q, b_l, b_{l-1}, \ldots b_1$ of M by a tree

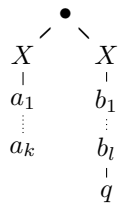

Every transition of the Turing machine will be simulated by the rewriting system in two steps, by first rewriting the right branch of the configuration tree, and then rewriting the left branch. The labels of the rewriting rules will indicate which letter from Γ has to be added $(+)$ or removed $(-)$ from the left branch of the configuration tree, and $\top$ respectively $\bot$ indicate whether the halting state has been reached or not.

More precisely we define a GTRS $\mathcal{R} = (A, \Sigma, R, t_0)$ where $A_2 = \{\bullet\}$, $A_1 = \Gamma \cup \{X\}$, $A_0 = A_1 \cup Q$, $\Sigma = \{+,-\} \times (\Gamma \cup \bar{\Gamma}) \times \{\bot, \top\}$ and

$$t_0 := \begin{array}{ccc} & \bullet & \\ X & & \begin{array}{c} X \\ \vdots \\ q_0 \end{array} \end{array}.$$

The set R is defined by adding for $\delta(q,b) = (p,c,L)$ and every $a \in \Gamma$ the rules

$$\begin{array}{c} b \\ \vdots \\ q \end{array} \xrightarrow{(-,a,*)} \begin{array}{c} c \\ \vdots \\ a \\ \vdots \\ p \end{array} \text{ and } \begin{array}{c} X \\ \vdots \\ q \end{array} \xrightarrow{(-,a,*)} \begin{array}{c} c \\ \vdots \\ a \\ \vdots \\ p \end{array} \text{ if } b = \text{␣},$$

and for $\delta(q,b) = (p,c,R)$ with $p \neq q_f$ and every $a \in \Gamma$ the rules

$$\begin{array}{c} a \\ \vdots \\ b \\ \vdots \\ q \end{array} \xrightarrow{(+,c,*)} \begin{array}{c} a \\ \vdots \\ p \end{array} \text{ and } \begin{array}{c} X \\ \vdots \\ q \end{array} \xrightarrow{(+,c,*)} \begin{array}{c} X \\ \vdots \\ p \end{array} \text{ if } b = \text{␣}$$

where $* = \top$ if $p = q_f$ and $* = \bot$ otherwise. Note that these rules can only be applied to the right branch of a configuration tree. For the left branch we add for every $a, c \in \Gamma$ and $* \in \{\bot, \top\}$ the rules

$$a \xrightarrow{(-,\bar{a},*)} \varepsilon \text{ and } X \xrightarrow{(-,\bar{a},*)} X \text{ if } a = \text{␣},$$

as well as

$$a \xrightarrow{(+,\bar{c},*)} \begin{array}{c} a \\ \vdots \\ c \end{array}$$

and

$$X \xrightarrow{(+,\bar{c},*)} \begin{array}{c} X \\ \vdots \\ c \end{array} \text{ if } c \neq \text{␣} \text{ and } X \xrightarrow{(+,\bar{c},*)} X \text{ if } c = \text{␣}.$$

By construction, a path through the graph G generated by $\mathcal{R}$ corresponds to a valid computation of M started on the empty tape iff every transition with label $(+,a,*)$ respectively $(-,a,*)$ is followed by its counterpart labeled $(+,\bar{a},*)$ respectively $(-,\bar{a},*)$. Let H be the star graph with $|\Sigma|+1$ many vertices where the center vertex v has for every $(\$,a,*) \in \{+,-\} \times \Gamma \times \{\bot,\top\}$ a single outgoing edge with this label to a vertex w and the single corresponding incoming edge from w labeled $(\$,\bar{a},*)$. If we define the synchronization constraint $C := \{(\sigma,\sigma) \mid \sigma \in \Sigma\}$, the synchronized product of G and H will contain exactly the valid computations of M. To decide whether M halts on the empty tape we thus have to check the truth of the formula

$$\exists xy \Big[\forall z \bigwedge_{\sigma \in \Sigma} \neg E_\sigma zx \\ \wedge \exists z \Big(\mathrm{Reach}_\Sigma(x,z) \wedge \bigwedge_{\sigma \in \{+,-\} \times \bar{\Gamma} \times \{\top\}} E_\sigma zy \Big) \Big]$$

in the semifinitely synchronized product of G and H. □

4. Transitive Closure Logic over the Infinite Grid

The infinite grid is the structure $\mathcal{G} = (\omega^2, S_1, S_2)$ with two successor relations S_1 and S_2. It can be viewed as the asynchronous and hence finitely synchronized product of two copies of the natural numbers with successor relation, $\mathcal{N}_1 = (\omega, S_1)$ and $\mathcal{N}_2 = (\omega, S_2)$, defined by the empty synchronization constraint.

We show in this section how to interpret the first-order theory of addition and multiplication of the natural numbers in $\mathrm{FO(TC)}^2_{(1)}$ (without parameters) over the infinite grid. $\mathrm{FO(TC)}^2_{(1)}$ allows only transitive closure operators of arity one and a nesting depth of two.

It is well known that the FO-theory of addition and multiplication of $\mathcal{N}$ is undecidable. However, since $\mathrm{FO(TC)}_{(1)}$ can be interpreted in MSO, $\mathrm{FO(TC)}_{(1)}$ is decidable over $\mathcal{N}$. From these results we can conclude that the $\mathrm{FO(TC)}^2_{(1)}$-theory is not preserved by finitely synchronized products and thus obtain that we cannot extend FO(R) much without losing decidability for finitely synchronized products.

To interpret the theory of addition and multiplication in $\mathrm{FO(TC)}^2_{(1)}$ over the infinite grid we first connect the transitive closure theories of $\mathcal{N}$ and $\mathcal{G}$.

Lemma 4 *Let $k \geq 1$.*

(a) For every $\mathrm{FO(TC)}^n_{(k)}$-sentence φ there is a $\mathrm{FO(TC)}^n_{(2k)}$-sentence $\tilde{\varphi}$ such that $\mathcal{G} \models \varphi \Leftrightarrow \mathcal{N} \models \tilde{\varphi}$.

(b) For every $\mathrm{FO(TC)}^n_{(2k)}$-sentence φ there is a $\mathrm{FO(TC)}^n_{(k)}$-sentence $\hat{\varphi}$ such that $\mathcal{N} \models \varphi \Leftrightarrow \mathcal{G} \models \hat{\varphi}$.

Proof: For (a) there is almost nothing to show. It suffices to split every variable x (interpreted as vertex of the grid) into coordinate variables x_1 and x_2 (interpreted as natural numbers) and to replace the atomic formulas S_1xy by Sx_1y_1 and S_2xy by Sx_2y_2.

For (b) we identify every $x \in \omega$ with $(x,0) \in \omega^2$. To reduce the number of variables needed in a TC operator we

represent a pair of variables x_1, x_2 by a single variable $x = (x_1, x_2)$ to be interpreted as a vertex of the grid.

To finish the proof it suffices to show that the following operations are FO(TC)$_{(1)}$ definable:

(i) π_i with $\pi_1((x_1, x_2)) := (x_1, 0)$ and $\pi_2((x_1, x_2)) := (0, x_2)$,

(ii) swap$_i$ with swap$_1((x, 0)) := (0, x)$ and swap$_2((0, x)) := (x, 0)$,

(iii) comb with comb$((x, 0), (0, y)) := (x, y)$

Then a FO(TC)$_{(2k)}$ formula $[\mathrm{TC}_{\bar{x},\bar{y}}\, \varphi(\bar{x}, \bar{y}, \bar{z})]\bar{x}, \bar{y}$ is equivalent to the FO(TC)$_{(k)}$ formula

$$\exists \bar{u}\bar{v}\Big(\bigwedge_{1 \le i \le k} \big(u_i = \text{comb}(x_{2i-1}, \text{swap}_2(x_{2i})) \wedge v_i = \text{comb}(y_{2i-1}, \text{swap}_2(y_{2i}))\big) \wedge [\mathrm{TC}_{\bar{u},\bar{v}}\, \tilde{\varphi}(\bar{u}, \bar{v}, \bar{z})]\bar{u}, \bar{v}\Big)$$

where

$$\tilde{\varphi} := \exists \bar{x}\bar{y}\Big(\bigwedge_{1 \le i \le k} \big(x_{2i-1} = \pi_1(u_i) \wedge x_{2i} = \text{swap}_2(\pi_2(u_i)) \wedge y_{2i-1} = \pi_1(v_i) \wedge y_{2i} = \text{swap}_2(\pi_2(v_i))\big) \wedge \varphi(\bar{x}, \bar{y})\Big)$$

and in φ every occurrence of the symbol S is replaced by S_1.

Let us now define the operations above:

$$\begin{aligned}
\pi_1(x) = y &\leftrightarrow y \le_2 x \wedge \forall z(z \le_2 x \to z = y)\\
\text{swap}_1(x) = y &\leftrightarrow \forall z(z \le_2 x \to z = x)\\
&\quad \wedge [\mathrm{TC}_{x,y}\, \exists z(S_1 xz \wedge S_2 yz)]x, y\\
&\quad \wedge \forall z(z \le_1 y \to z = y)\\
\text{comb}(x, y) = z &\leftrightarrow \forall u(u \le_2 x \to u = x)\\
&\quad \wedge \forall u(u \le_1 y \to u = y)\\
&\quad \wedge x \le_1 z \wedge y \le_2 z
\end{aligned}$$

Observe that if the formula φ has no TC operators with parameters, the neither $\tilde{\varphi}$ nor $\hat{\varphi}$ has (in $\tilde{\varphi}$ only TC-formulas without parameters are introduced), and that the nesting depth is not increased. □

Let us now turn to the undecidability proof.

Theorem 5 *The FO(TC)$^2_{(1)}$-theory of the infinite grid is undecidable.*

Proof: We define addition and multiplication in FO(TC)$_{(1)}$ over $\mathcal{G}$ without the use of parameters. By Lemma 4 it is enough to define these operations in FO(TC)$_{(2)}$ over $\mathcal{N}$. The definition of addition is straightforward.

$$a + b = c \leftrightarrow \mathcal{N} \models [\mathrm{TC}_{x_1x_2,y_1y_2}\, Sx_1y_1 \wedge Sx_2y_2]0\, a, b\, c$$

To define multiplication note that $x \cdot y = \frac{(x+y)^2 - x^2 - y^2}{2}$, hence it suffices to define the square function. To define x^2 note that $x^2 = \sum_{i=0}^{x-1} 2i + 1$. The formula

$$\psi(x, y) = [\mathrm{TC}_{x_1x_2,y_1y_2}\, y_2 = x_2 + (x_2 - x_1) + 2 \wedge y_1 = x_2]0\, 1, xy$$

defines all pairs of square numbers

$$\Big(\sum_{i=1}^{k-2} 2i + 1, \sum_{i=1}^{k-1} 2i + 1\Big) \text{ for } k \ge 3.$$

Hence $\mathcal{N} \models \psi[a, b]$ iff $b - a = 2k - 1$ for some $k \ge 2$. Let

$$\chi(x, y) = \exists z_1\Big(\psi(z_1, y) \wedge \frac{y - z_1 + 1}{2} = x\Big).$$

Then $\mathcal{N} \models \chi[a, b]$ iff $b = a^2$. □

A similar technique was used in [3] to define multiplication in $(\omega, +, 0)$ using a transitive closure operator of arity one.

The nesting of transitive closure operators in the previous proof is necessary. If we disallow nesting, even in the presence of parameters in the transitive closure formulas, the theory of the infinite grid is decidable.

Theorem 6 *The FO(TC)$^1_{(1)}$-theory of the infinite grid is decidable.*

Proof: We reduce the FO(TC)$^1_{(1)}$-theory of the infinite grid $\mathcal{G}$ to Presburger arithmetic, the first-order theory of $\mathcal{N}_+ = (\omega, +, 0)$, in the following sense: For every FO(TC)$^1_{(1)}$-formula $\varphi(x_1, \ldots, x_n)$ there is a Presburger formula $\tilde{\varphi}(x_{11}, x_{12}, \ldots, x_{n1}, x_{n2})$ such that

$$\mathcal{G} \models \varphi[(k_1, l_1), \ldots, (k_n, l_n)] \Leftrightarrow \mathcal{N}_+ \models \tilde{\varphi}[k_1, l_1, \ldots, k_n, l_n]. \quad (2)$$

In order to construct $\tilde{\varphi}$ it suffices to consider the case

$$\varphi(x_1, \ldots, x_n) = [\mathrm{TC}_{x_1,x_2}\, \psi(x_1, \ldots, x_n)]x_1, x_2,$$

or for better readability

$$\varphi(x_1, \ldots, x_n) = [\mathrm{TC}_{x,y}\, \psi(x, y, x_3, \ldots, x_n)]x_1, x_2$$

where ψ is a first-order formula. The second notation emphasizes that $x_3, \ldots, x_n$ serve as parameters in the transitive closure formula.

In a first step we rewrite ψ in a normal form, applying Hanf's Theorem for first-order logic over graphs (see [12, 8, 27]).

For this purpose we recall some definitions. The r-sphere $r\text{-sph}(d)$ around a vertex $d \in \omega^2$ is the set of grid vertices which are of distance less or equal to r from d, where we allow to traverse the edges in either direction. Invoking the distributive normal form and Hanf's Theorem, there exists a suitable $r > 0$ such that $\psi(x_1, \ldots, x_n)$ is equivalent to a disjunction of formulas $\varphi_\tau(x_1, \ldots, x_n)$ where each φ_τ describes the isomorphism type τ of $\bigcup_{1\le i\le n} r\text{-sph}(c_i)$ for some tuple $c_1, \ldots, c_n$ of grid vertices. Let T be the set of all such types. Since T is finite it suffices to consider only finitely many tuples $c_1, \ldots, c_n$.

Remark. In the general case, over an arbitrary graph instead of the infinite grid, Hanf's Theorem involves a statement on the number (up to a certain threshold) of spheres outside $\bigcup_{1\le i\le n} r\text{-sph}(c_i)$. This statement is superfluous here due to the regular structure of the infinite grid. (For technical convenience we assume that $(0,0)$ is included in the set of parameters, so every isomorphism type realizable in $\mathcal{G}$ outside $\bigcup_{1\le i\le n} r\text{-sph}(c_i)$ occurs an infinite number of times.)

Due to the special structure of the grid, which we depict as a diagram with the bottom row and left column as margins, open upwards and to the right, every formula $\varphi_\tau(x_1, \ldots, x_n)$ can be expressed by conditions on the vertices $x_1, \ldots, x_n$ which fix their distances up to the radius r from the left margin as well as the bottom margin, and their relative distances up to $2r$.

It is convenient to express $\varphi_\tau(x_1, \ldots, x_n)$ in terms of the $2n$ components of the vertices, obtaining a formula $\tilde\varphi_\tau(x_{11}, x_{12}, \ldots, x_{n1}, x_{n2})$ The formula $\tilde\varphi_\tau$ is interpreted over ω and equivalent to φ in the sense of (2) above. It is a conjunction of statements

- $x_{ih} = k$ for $0 \le k \le r$ or $x_{ih} > r$
- $(x_{i1}, x_{i2}) = (x_{j1}, x_{j2}) + (k, l)$ for $-2r \le k, l \le 2r$
- $\mathrm{dist}((x_{i1}, x_{i2}), (x_{j1}, x_{j2})) > 2r$

where $1 \le i, j \le n$ and $h \in \{1, 2\}$.

We now have to evaluate formulas of the form

$$\Big[\mathrm{TC}_{(x_{11},x_{12}),(x_{21},x_{22})} \bigvee_{\tau\in T'} \tilde\varphi_\tau(x_{11}, x_{12}, \ldots, x_{n1}, x_{n2})\Big](s,t),(u,v) \quad (3)$$

for some $T' \subseteq T$.

In a first step we note that it is possible to add disjuncts to (3) such that vertices tied to occur in a $2r$-sphere around a parameter (x_{i1}, x_{i2}) for $i > 2$ only need to appear as start vertex or as end vertex of any path described by (3). Hence vertices tied to parameters can be handled without the use of TC, by an appropriate modification of the formula.

Let I be an initial segment of the grid encompassing the $2r$-spheres around parameters (x_{i1}, x_{i2}) for $i > 2$. Outside this initial segment, in a second step, it suffices to consider formulas (3) in which only type formulas $\tilde\varphi_\tau$ which contain

$$x_{11} = k_1 \wedge x_{12} > r \quad \text{or} \quad x_{11} > r \wedge x_{12} = k_2 \quad \text{or} \quad x_{11} > r \wedge x_{12} > r$$

for $k_1, k_2 \le r$ and

$$x_{21} = l_1 \wedge x_{22} > r \quad \text{or} \quad x_{21} > r \wedge x_{22} = l_2 \quad \text{or} \quad x_{21} > r \wedge x_{22} > r$$

for $l_1, l_2 \le r$ and

$$\mathrm{dist}((x_{11}, x_{12}), (x_{21}, x_{22})) > 2r \quad \text{or } (x_{11}, x_{12}) = (x_{21}, x_{22}) + (k, l)$$

for $-2r \le k, l \le 2r$ appear.

It is now possible to apply a finite saturation process to obtain a formula

$$\Big[\mathrm{TC}_{(x_{11},x_{12}),(x_{21},x_{22})} \bigvee_{1\le j\le m} \tilde\varphi_j(x_{11}, x_{12}, \ldots, x_{n1}, x_{n2})\Big](s,t),(u,v) \quad (4)$$

which is equivalent to (3) and where TC and $\bigvee$ commute, i.e.

$$\mathcal{G} \models \Big[\mathrm{TC}_{(x_{11},x_{12}),(x_{21},x_{22})} \bigvee_{1\le j\le m} \tilde\varphi_j\Big](s,t),(u,v) \Leftrightarrow \mathcal{G} \models \bigvee_{1\le j\le m} \Big[\mathrm{TC}_{(x_{11},x_{12}),(x_{21},x_{22})}\, \tilde\varphi_j\Big](s,t),(u,v).$$

The subformulas $\tilde\varphi_j$ in (4) have the same format as the subformulas $\tilde\varphi_\tau$ in (3) except that the center of the excluded $2r$-sphere around (x_{11}, x_{12}) may be shifted by a bounded distance from (x_{11}, x_{12}) or be missing, or $\tilde\varphi_\tau$ defines the complete relation outside I and the border stripes of width r. Thus it remains to consider two cases.

Case 1. If $\tilde\varphi_j$ contains a conjunct excluding some $2r$-sphere then the relation defined by $[\mathrm{TC}_{(x_{11},x_{12}),(x_{21},x_{22})}\, \tilde\varphi_j](s,t),(u,v)$ is cofinite (w.r.t. the grid excluding I and border stripes of width r, or a fixed line in one of the border stripes) and hence definable without the use of a transitive closure operator.

Case 2. If $\tilde\varphi_j$ fixes relations of the form

$$(x_{21}, x_{22}) = (x_{11}, x_{12}) + (k_i, l_i) \quad (5)$$

for $i = 1, \ldots, N$ and $-2r \leq k_i, l_i \leq 2r$. the formula

$$[\mathrm{TC}_{(x_{11},x_{12}),(x_{21},x_{22})}\, \tilde{\varphi}_j](s,t),(u,v)$$

expresses that there is a path from (s,t) to (u,v) consisting of steps of the form (4). The set of vertices (u,v) reachable in this way from (s,t) can be represented as the union of paths in the finite initial segment I of the grid and finitely many sets of the form

$$\{(u,v) \mid (u,v) = (s',t') + y_1(k_1,l_1) + \ldots + y_N(k_N,l_N)\}.$$

Here $y_i \geq 0$, the (s',t') range over boundary vertices of I, and the (k_i,l_i) are from (5). It follows that the relation defined by (3) is definable in Presburger arithmetic. □

5. Conclusion

We have proved a result on compositional model-checking for a logic including reachability predicates, and we have shown tight limitations for possible extensions of this result.

Let us mention some questions left open in this paper (and partly subject of ongoing work).

1. The Feferman-Vaught type composition result (Theorem 1) should be generalized to infinite products.

2. In Theorem 1, one may extend FO(R) by an operator for "recurrent reachability" (existence of an infinite path which visits a designated set infinitely often), or one can consider stronger logics like CTL.

3. Interesting subcases of Theorem 1 should be found where the mentioned blow-up of complexity can be avoided.

4. The distinction between products which are asynchronous, finitely synchronized, or synchronized should be refined, by allowing other means of coordination between component structures, also incorporating the special case of synchronization of parameterized systems composed from identical components.

Acknowledgment

We thank C. Löding for pointing us to GTRS-graphs to prove Theorem 3 and the anonymous referees for many helpful comments and pointers to related literature.

References

[1] A. Arnold. *Finite Transition Systems*. Prentice Hall, 1994.

[2] A. Arnold. Nivat's processes and their synchronization. *Theoretical Computer Science*, 281:31–36, 2002.

[3] A. Avron. Transitive closure and the mechanization of mathematics. In F. Kamareddine, editor, *Thirty Five Years of Automating Mathematics*, pages 149–171. Kluwer Academic Publishers, 2003.

[4] D. Caucal. On infinite transition graphs having a decidable monadic theory. In *Proceedings of the 23rd International Colloquium on Automata, Languages and Programming*, volume 1099 of *Lecture Notes in Computer Science*, pages 194–205, 1996.

[5] D. Caucal. On infinite terms having a decidable theory. In *Proceedings of the 27th International Symnposium on Mathematical Foundations of Computer Science*, volume 2420 of *Lecture Notes in Computer Science*, pages 165–176. Springer, 2002.

[6] C. Chang and H. Keisler. *Model Theory*. North-Holland, 1973.

[7] T. Colcombet. On families of graphs having a decidable first order theory with reachability. In *Proceedings of the 29th International Conference on Automata, Languages, and Programming*, volume 2380 of *Lecture Notes in Computer Science*, pages 98–109, 2002.

[8] H. Ebbinghaus and J. Flum. *Finite Model Theory*. Springer, 1995.

[9] S. Feferman and R. Vaught. The first-order properties of products of algebraic systems. *Fundamenta Mathematicae*, 47:57–103, 1959.

[10] M. Grohe. Arity hierarchies. *Annals of Pure and Applied Logic*, 82:103–163, 1996.

[11] Y. Gurevich. Monadic second-order theories. In J. Barwise and S. Feferman, editors, *Model-Theoretic Logics*, pages 479–506. Springer, 1985.

[12] W. Hanf. Model-theortic methods in the study of elementary logic. In *Proceedings of the Symposium on the Theory of Models*, pages 132–145. North Holland, 1965.

[13] C. Hoare. Communicating sequential processes. *ACM Communications*, 21:666–677, 1978.

[14] J. Hopcroft and J. Ullman. *Introduction to Automata Theory, Languages, and Computation*. Addison-Wesley, 1979.

[15] T. Knapik and É. Payet. Synchronized product of linear bounded machines. In *Proceedings of the 12th International Symposium on Fundamentals of Computation Theory*, volume 1684 of *Lecture Notes in Computer Science*, pages 362–373. Springer, 1999.

[16] C. Löding. Model-checking infinite systems generated by ground tree rewriting. In *Proceedings of the 5th International Conference on Foundations of Software Science and Computation Structures*, volume 2303 of *Lecture Notes in Computer Science*, pages 280–294. Springer, 2002.

[17] C. Löding. *Infinite Graphs Generated by Tree Rewriting*. PhD thesis, RWTH Aachen, 2003.

[18] J. Makowsky. Algorithmic aspects of the Feferman-Vaught theorem. *Annals of Pure and Applied Logic*, 126:159–213, 2004.

[19] R. Milner. Lectures on a calculus for communicating systems. In *Seminar on Concurrency*, volume 197 of *Lecture Notes in Computer Science*, pages 197–220, 1985.

[20] C. Morvan. On rational graphs. In *Proceedings of the 3rd International Conference on Foundations of Software Science and Computation Structures*, volume 1784 of *Lecture Notes in Computer Science*, pages 252–266. Springer, 2000.

[21] D. Muller and P. Schupp. The theory of ends, pushdown automata, and second-order logic. *Theoretical Computer Science*, 37:51–75, 1985.

[22] É. Payet. Thue specifications, infinite graphs and synchronized product. *Fundamenta Informaticae*, 44:265–290, 2000.

[23] A. Rabinovich. Composition theorem for generalized sum. Submitted. Available at http://www.math.tau.ac.il/˜rabinoa/journal.html.

[24] A. Rabinovich. On the compositional method and its limitations. Technical Report EDI-INF-RR-0035, University of Edinburgh, 2001.

[25] S. Shelah. The monadic theory of orders. *Annals of Mathematics*, pages 379–419, 1975.

[26] S. Shelah. On the very weak $0 - 1$ law for random graphs with orders. *Journal of Logic an Computation*, 6:137–159, 1996.

[27] W. Thomas. Languages, automata, and logic. In G. Rozenberg and A. Salomaa, editors, *Handbook of Formal Languages*, volume 3, pages 389–455. Springer, 1997.

[28] W. Thomas. Constructing infinite graphs with a decidable MSO-theory. In *Proceedings of the 28th International Symposium on Mathematical Foundations of Computer Science*, volume 2747 of *Lecture Notes in Computer Science*, pages 113–124. Springer, 2003.