Ideal-Based Algorithms for the Symbolic Verification of Well-Structured Systems

Philippe Schnoebelen

LSV, CNRS & ENS Paris Saclay, Cachan, France
phs@lsv.fr

Abstract

We explain how the downward-closed subsets of a well-quasi-ordering $(X, \leq)$ can be represented via the ideals of X and how this leads to simple and efficient algorithms for the verification of well-structured systems.

1998 ACM Subject Classification F.1.1 Models of Computation, F.3.1 Specifying and Verifying and Reasoning about Programs

Keywords and phrases Well-structured systems and verification, Order theory

Digital Object Identifier 10.4230/LIPIcs.MFCS.2017.85

Category Invited Talk

1 Summary of the talk

Well-structured systems, also known under the acronym "WSTSs", are a family of infinite-state models for which generic verification algorithms exist [1, 2, 13, 18, 23]. With WSTSs, the main ingredient for decidability is the existence of an ordering on configurations that enjoys two properties:

- it is a well-quasi-ordering (a WQO): every infinite sequence $c_0, c_1, c_2, \ldots$ of configurations contains an increasing pair $c_i \leq c_j$ with $i < j$;
- transitions are monotonic: if the system can perform a step $c \to c'$ then from any configuration $d \geq c$, a "similar" step is possible, i.e., there is some $d \to d'$ with $d' \geq c'$.

The most well-known instances of WSTSs are some families of counter machines or vector addition systems [8, 12]. For simplicity, we shall assume that the WQO set of configurations for these systems is $\mathit{Conf} = (\mathbb{N}^d, \leq_\times)$ for some dimension $d \in \mathbb{N}$, where the *component-wise ordering* $\leq_\times$ is given by $\boldsymbol{u} = (u_1, \ldots, u_d) \leq_\times \boldsymbol{v} = (v_1, \ldots, v_d) \overset{\text{def}}{\Leftrightarrow} u_1 \leq v_1 \wedge \cdots \wedge u_d \leq v_d$.

Another well-known instance are the lossy channel systems [4, 7], where for simplicity we assume that the set of configurations is $(\Sigma^*, \leq_*)$ for some finite alphabet $\Sigma = \{\mathtt{a}, \mathtt{b}, \ldots\}$ of messages, and where $\leq_*$ is the *subword ordering* [1] given by

$$u \leq v \overset{\text{def}}{\Leftrightarrow} \exists \mathtt{a}_1, \ldots, \mathtt{a}_\ell \in \Sigma : \exists v_0, \ldots, v_\ell \in \Sigma^* : u = \mathtt{a}_1\mathtt{a}_2 \cdots \mathtt{a}_\ell \wedge v = v_0\mathtt{a}_1v_1\mathtt{a}_2 \cdots \mathtt{a}_\ell v_\ell \,.$$

Algorithms for the verification of safety properties of WSTSs usually involve reasoning and computing with *upward-closed* and/or *downward-closed* sets of configurations. A set $U \subseteq \mathit{Conf}$ is upward-closed $\overset{\text{def}}{\Leftrightarrow} c \in U \wedge c \leq c' \implies c' \in U$, and there is a similar definition

[1] That $(\Sigma^*, \leq_*)$ is a WQO is known as Higman's Lemma.

42nd International Symposium on Mathematical Foundations of Computer Science (MFCS 2017).
Editors: Kim G. Larsen, Hans L. Bodlaender, and Jean-Francois Raskin; Article No. 85; pp. 85:1–85:4
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

for downward-closed subsets. These sets are usually infinite (like *Conf* itself) and symbolic representations or data structures are needed in algorithms handling them.

For upward-closed subsets, a well-known representation relies on the existence of *minimal bases*, i.e., the fact that the set of minimal elements of any subset is finite and unique (modulo equivalence). This representation is generic: it works for any WQO. Furthermore, it enjoys several nice algorithmic properties, e.g., testing inclusion between upward-closed subsets reduces to a quadratic number of comparisons between individual configurations, and the union of upward-closed sets is very easy to compute. In the case of $(\mathbb{N}^d, \leq_\times)$ or $(\Sigma^*, \leq_*)$, algorithms for computing intersections reduce to easy computations of least upper bounds between elements.

For downward-closed subsets, one cannot rely on a mirror notion of maximal elements and this makes symbolic computations harder to envision. The question of finding a generic approach for computing with downward-closed sets was first raised in [14].

In the case of $(\mathbb{N}^d, \leq_\times)$, a symbolic technique was popularized by Karp and Miller with their classic algorithm for coverability in VAS [19]. They define $\mathbb{N}_\omega = \mathbb{N} \cup \{\omega\}$ —where the set of natural numbers is completed with a new infinite element ω that is larger than any finite number— and consider d-tuples over $\mathbb{N}_\omega$. It turns out that this is exactly what we need to represent downward-closed subsets of $\mathbb{N}^d$. For $\sigma = (s_1, \ldots, s_d) \in \mathbb{N}_\omega^d$, we let $\downarrow\sigma = \{c \in \mathbb{N}^d \mid c \leq_\times \sigma\}$ denote the downward-closed subset of $\mathbb{N}^d$ generated by σ and call it an *ideal* of $(\mathbb{N}^d, \leq_\times)$. Then downward-closed subsets of $\mathbb{N}^d$ can be denoted in a unique way by finite unions of incomparable ideals. Computing unions and intersections with such representations, and deciding inclusion between them, use simple algorithms that are uncannily similar to what happened with the finite-basis representation for upward-closed subsets.

If we now consider $(\Sigma^*, \leq_*)$, a very elegant representation for downward-closed subsets was proposed by Abdulla et al. in [3]. They show that any downward-closed $D \subseteq \Sigma^*$ can be represented by a *simple regular expression* (a SRE), obtained as a union of concatenations of *atoms* of the form Γ^* for a subalphabet $\Gamma \subseteq \Sigma$, or of the form $\mathtt{a} + \epsilon$ for some letter $\mathtt{a} \in \Sigma$. Furthermore, these SREs support simple and efficient algorithms for unions, intersections, comparisons, and more.

It turns out that concatenations of atoms denote exactly the ideals of $(\Sigma^*, \leq_*)$. Formally, an ideal of a WQO $(X, \leq)$ is a nonempty downward-closed directed subset $D \subseteq X$. Being directed means that for all $x, y \in D$ there is some $z \in D$ with $x \leq z \wedge y \leq z$. Given any WQO $(X, \leq)$, the downward-closed subsets of X can be written as unions of finitely many pairwise incomparable ideals, and this decomposition is unique. This property explains the nice algorithmic properties we observed with $\mathbb{N}_\omega^d$ and the SREs over Σ, and it generalizes to any WQO where we can provide effective characterizations for the ideals.

In the second part of the talk, we show how such effective characterizations exist for most of the WQOs one encounters in practice. This is done by considering the most common ways of constructing new WQOs from previous ones (sequence extension, powerset, but also substructures and quotients) and characterizing the ideals of the new WQOs in terms of the ideals of the earlier ones.

We illustrate these constructions with lesser known WSTSs like priority channel systems and higher-order channel systems [17], or data nets [20] and timed-arc Petri nets [5].

Acknowledgments. This talk is based on joint work with J. Goubault-Larrecq, S. Halfon, P. Karandikar, K. Narayan Kumar, S. Schmitz, and it has further profited from many discussions with A. Finkel, J. Leroux and G. Sutre. Most of the presented definitions and

results can be found in recent works like [6, 9, 10, 11, 16, 21, 22]. A full version of these notes is in preparation [15].

References

1 P. A. Abdulla. Well (and better) quasi-ordered transition systems. *Bull. Symbolic Logic*, 16(4):457–515, 2010. doi:10.2178/bsl/1294171129.

2 P. A. Abdulla, K. Čerāns, B. Jonsson, and Yih-Kuen Tsay. Algorithmic analysis of programs with well quasi-ordered domains. *Information and Computation*, 160(1/2):109–127, 2000. doi:10.1006/inco.1999.2843.

3 P. A. Abdulla, A. Collomb-Annichini, A. Bouajjani, and B. Jonsson. Using forward reachability analysis for verification of lossy channel systems. *Formal Methods in System Design*, 25(1):39–65, 2004. doi:10.1023/B:FORM.0000033962.51898.1a.

4 P. A. Abdulla and B. Jonsson. Verifying programs with unreliable channels. *Information and Computation*, 127(2):91–101, 1996. doi:10.1006/inco.1996.0053.

5 B. Bérard, F. Cassez, S. Haddad, D. Lime, and O. H. Roux. The expressive power of time Petri nets. *Theoretical Computer Science*, 474, 2012. doi:10.1016/j.tcs.2012.12.005.

6 M. Blondin, A. Finkel, and P. McKenzie. Well behaved transition systems. arXiv:1608.02636 [cs.LO], August 2016. To appear in Logical Meth. Comp. Sci. URL: http://arxiv.org/abs/1608.02636.

7 G. Cécé, A. Finkel, and S. Purushothaman Iyer. Unreliable channels are easier to verify than perfect channels. *Information and Computation*, 124(1):20–31, 1996. doi:10.1006/inco.1996.0003.

8 C. Dufourd, P. Jančar, and Ph. Schnoebelen. Boundedness of Reset P/T nets. In *CONCUR '99, LNCS 1644*, pages 301–310. Springer, 1999. doi:10.1007/3-540-48523-6_27.

9 A. Finkel. The ideal theory for WSTS. In *RP 2016, LNCS 9899*, pages 1–22. Springer, 2016. doi:10.1007/978-3-319-45994-3_1.

10 A. Finkel and J. Goubault-Larrecq. Forward analysis for WSTS, part I: Completions. In *STACS 2009, LIPIcs 3*, pages 433–444. Leibniz-Zentrum für Informatik, 2009. doi:10.4230/LIPIcs.STACS.2009.1844.

11 A. Finkel and J. Goubault-Larrecq. Forward analysis for WSTS, part II: Complete WSTS. *Logical Methods in Comp. Science*, 8(4), 2012. doi:10.2168/LMCS-8(3:28)2012.

12 A. Finkel, P. McKenzie, and C. Picaronny. A well-structured framework for analysing Petri nets extensions. *Information and Computation*, 195(1–2):1–29, 2004. doi:10.1016/j.ic.2004.01.005.

13 A. Finkel and Ph. Schnoebelen. Well-structured transition systems everywhere! *Theoretical Computer Science*, 256(1–2):63–92, 2001. doi:10.1016/S0304-3975(00)00102-X.

14 G. Geeraerts, J.-F. Raskin, and L. Van Begin. Expand, enlarge and check: New algorithms for the coverability problem of WSTS. *Journal of Computer and System Sciences*, 72(1):180–203, 2006. doi:10.1016/j.jcss.2005.09.001.

15 J. Goubault-Larrecq, S. Halfon, P. Karandikar, K. Narayan Kumar, and Ph. Schnoebelen. The ideal approach to computing closed subsets in well-quasi-orderings. In preparation, 2017.

16 J. Goubault-Larrecq and S. Schmitz. Deciding piecewise testable separability for regular tree languages. In *ICALP 2016, LIPIcs 55*, pages 97:1–97:15. Leibniz-Zentrum für Informatik, 2016. doi:10.4230/LIPIcs.ICALP.2016.97.

17 Ch. Haase, S. Schmitz, and Ph. Schnoebelen. The power of priority channel systems. *Logical Methods in Comp. Science*, 10(4:4), 2014. doi:10.2168/LMCS-10(4:4)2014.

18 T. A. Henzinger, R. Majumdar, and J.-F. Raskin. A classification of symbolic transition systems. *ACM Trans. Computational Logic*, 6(1):1–32, 2005. doi:10.1145/1042038.1042039.

19 R. M. Karp and R. E. Miller. Parallel program schemata. *Journal of Computer and System Sciences*, 3(2):147–195, 1969. doi:10.1016/S0022-0000(69)80011-5.

20 R. Lazić, T. Newcomb, J. Ouaknine, A. W. Roscoe, and J. Worrell. Nets with tokens which carry data. *Fundamenta Informaticae*, 88(3):251–274, 2008.

21 R. Lazić and S. Schmitz. The complexity of coverability in ν-Petri nets. In *LICS 2016*, pages 467–476. ACM Press, 2016. doi:10.1145/2933575.2933593.

22 J. Leroux and S. Schmitz. Ideal decompositions for vector addition systems. In *STACS 2016, LIPIcs 47*, pages 1:1–1:13. Leibniz-Zentrum für Informatik, 2016. doi:10.4230/LIPIcs.STACS.2016.1.

23 S. Schmitz and Ph. Schnoebelen. The power of well-structured systems. In *CONCUR 2013, LNCS 8052*, pages 5–24. Springer, 2013. doi:10.1007/978-3-642-40184-8_2.