

# 29th International Conference on Concurrency Theory

CONCUR 2018, September 4–7, 2018, Beijing, China

Edited by

Sven Schewe

Lijun Zhang



*Editors*

|  |  |
|--|--|
| Sven Schewe  | Lijun Zhang  |
| Department of Computer Science   | State Key Laboratory of Computer Science                 |
| University of Liverpool  | Institute of Software Chinese Academy of Sciences        |
| Liverpool, UK  | Beijing, China   |
| <a href="mailto:sven.schewe@liverpool.ac.uk">sven.schewe@liverpool.ac.uk</a> | <a href="mailto:zhanglj@ios.ac.cn">zhanglj@ios.ac.cn</a> |

*ACM Classification 2012*  
Theory of Computation

**ISBN 978-3-95977-087-3**

*Published online and open access by*  
Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH, Dagstuhl Publishing, Saarbrücken/Wadern, Germany. Online available at <http://www.dagstuhl.de/dagpub/978-3-95977-087-3>.

*Publication date*  
August, 2018

*Bibliographic information published by the Deutsche Nationalbibliothek*  
The Deutsche Nationalbibliothek lists this publication in the Deutsche Nationalbibliografie; detailed bibliographic data are available in the Internet at <http://dnb.d-nb.de>.

*License*  
This work is licensed under a Creative Commons Attribution 3.0 Unported license (CC-BY 3.0): <http://creativecommons.org/licenses/by/3.0/legalcode>.



In brief, this license authorizes each and everybody to share (to copy, distribute and transmit) the work under the following conditions, without impairing or restricting the authors' moral rights:

- Attribution: The work must be attributed to its authors.

The copyright is retained by the corresponding authors.

Digital Object Identifier: 10.4230/LIPIcs.CONCUR.2018.0

**ISBN 978-3-95977-087-3**

**ISSN 1868-8969**

**<http://www.dagstuhl.de/lipics>**

## LIPICs – Leibniz International Proceedings in Informatics

LIPICs is a series of high-quality conference proceedings across all fields in informatics. LIPICs volumes are published according to the principle of Open Access, i.e., they are available online and free of charge.

### *Editorial Board*

- Luca Aceto (*Chair*, Gran Sasso Science Institute and Reykjavik University)
- Susanne Albers (TU München)
- Christel Baier (TU Dresden)
- Javier Esparza (TU München)
- Michael Mitzenmacher (Harvard University)
- Madhavan Mukund (Chennai Mathematical Institute)
- Anca Muscholl (University Bordeaux)
- Catuscia Palamidessi (INRIA)
- Raimund Seidel (Saarland University and Schloss Dagstuhl – Leibniz-Zentrum für Informatik)
- Thomas Schwentick (TU Dortmund)
- Reinhard Wilhelm (Saarland University)

**ISSN 1868-8969**

**<http://www.dagstuhl.de/lipics>**



## ■ Contents

|  |          |
|--|----------|
| Preface                                  |          |
| <i>Sven Schewe and Lijun Zhang</i> ..... | 0:ix–0:x |

### Invited Contributions

|  |          |
|--|----------|
| The Siren Song of Temporal Synthesis   |          |
| <i>Moshe Y. Vardi</i> .....  | 1:1–1:1  |
| Bisimulations for Probabilistic and Quantum Processes                                      |          |
| <i>Yuxin Deng</i> .....  | 2:1–2:14 |
| Is Speed-Independent Mutual Exclusion Implementable?                                       |          |
| <i>Rob van Glabbeek</i> .....  | 3:1–3:1  |
| Verifying Arithmetic Assembly Programs in Cryptographic Primitives                         |          |
| <i>Andy Polyakov, Ming-Hsien Tsai, Bow-Yaw Wang, and Bo-Yin Yang</i> .....                 | 4:1–4:16 |
| Coalgebraic Theory of Büchi and Parity Automata: Fixed-Point Specifications, Categorically |          |
| <i>Ichiro Hasuo</i> .....  | 5:1–5:2  |

### Regular Papers

|   |            |
|---|------------|
| Universal Safety for Timed Petri Nets is PSPACE-complete  |            |
| <i>Parosh Aziz Abdulla, Mohamed Faouzi Atig, Radu Ciobanu, Richard Mayr, and Patrick Totzke</i> .....                       | 6:1–6:15   |
| It Is Easy to Be Wise After the Event: Communicating Finite-State Machines Capture First-Order Logic with “Happened Before” |            |
| <i>Benedikt Bollig, Marie Fortin, and Paul Gastin</i> .....   | 7:1–7:17   |
| Learning-Based Mean-Payoff Optimization in an Unknown MDP under Omega-Regular Constraints                                   |            |
| <i>Jan Křetínský, Guillermo A. Pérez, and Jean-François Raskin</i> .....  | 8:1–8:18   |
| Deciding Probabilistic Bisimilarity Distance One for Probabilistic Automata   |            |
| <i>Qiyi Tang and Franck van Breugel</i> .....   | 9:1–9:17   |
| Non-deterministic Weighted Automata on Random Words   |            |
| <i>Jakub Michaliszyn and Jan Otop</i> .....   | 10:1–10:16 |
| Ergodic Mean-Payoff Games for the Analysis of Attacks in Crypto-Currencies  |            |
| <i>Krishnendu Chatterjee, Amir Kafshdar Goharshady, Rasmus Ibsen-Jensen, and Yaron Velner</i> .....                         | 11:1–11:17 |
| Bounded Context Switching for Valence Systems   |            |
| <i>Roland Meyer, Sebastian Muskalla, and Georg Zetsche</i> .....  | 12:1–12:18 |
| Alternating Nonzero Automata  |            |
| <i>Paulin Fournier and Hugo Gimbert</i> .....   | 13:1–13:16 |

29th International Conference on Concurrency Theory (CONCUR 2018).

Editors: Sven Schewe and Lijun Zhang



Leibniz International Proceedings in Informatics  
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

|   |            |
|---|------------|
| Affine Extensions of Integer Vector Addition Systems with States<br><i>Michael Blondin, Christoph Haase, and Filip Mazowiecki</i> .....                               | 14:1–14:17 |
| Verifying Quantitative Temporal Properties of Procedural Programs<br><i>Mohamed Faouzi Atig, Ahmed Bouajjani, K. Narayan Kumar, and Prakash Saivasan</i> .....        | 15:1–15:17 |
| Narrowing down the Hardness Barrier of Synthesizing Elementary Net Systems<br><i>Ronny Tredup and Christian Rosenke</i> .....   | 16:1–16:15 |
| Up-To Techniques for Behavioural Metrics via Fibrations<br><i>Filippo Bonchi, Barbara König, and Daniela Petrişan</i> .....   | 17:1–17:17 |
| Completeness for Identity-free Kleene Lattices<br><i>Amina Doumane and Damien Pous</i> .....  | 18:1–18:17 |
| Reachability in Parameterized Systems: All Flavors of Threshold Automata<br><i>Jure Kukovec, Igor Konnov, and Josef Widder</i> .....                                  | 19:1–19:17 |
| Selective Monitoring<br><i>Radu Grigore and Stefan Kiefer</i> .....   | 20:1–20:16 |
| Synchronizing the Asynchronous<br><i>Bernhard Kragl, Shaz Qadeer, and Thomas A. Henzinger</i> .....   | 21:1–21:17 |
| A Semantics for Hybrid Iteration<br><i>Sergey Goncharov, Julian Jakob, and Renato Neves</i> .....   | 22:1–22:17 |
| GPU Schedulers: How Fair Is Fair Enough?<br><i>Tyler Sorensen, Hugues Evrard, and Alastair F. Donaldson</i> .....   | 23:1–23:17 |
| Linear Equations with Ordered Data<br><i>Piotr Hofman and Sławomir Lasota</i> .....   | 24:1–24:17 |
| A Coalgebraic Take on Regular and $\omega$ -Regular Behaviour for Systems with Internal Moves<br><i>Tomasz Brengos</i> .....  | 25:1–25:18 |
| Relating Syntactic and Semantic Perturbations of Hybrid Automata<br><i>Nima Roohi, Pavithra Prabhakar, and Mahesh Viswanathan</i> .....                               | 26:1–26:16 |
| Updating Probabilistic Knowledge on Condition/Event Nets using Bayesian Networks<br><i>Benjamin Cabrera, Tobias Heindel, Reiko Heckel, and Barbara König</i> .....    | 27:1–27:17 |
| Reachability in Timed Automata with Diagonal Constraints<br><i>Paul Gastin, Sayan Mukherjee, and B. Srivathsan</i> .....  | 28:1–28:17 |
| Parameterized complexity of games with monotonically ordered $\omega$ -regular objectives<br><i>Véronique Bruyère, Quentin Hautem, and Jean-François Raskin</i> ..... | 29:1–29:16 |
| A Universal Session Type for Untyped Asynchronous Communication<br><i>Stephanie Balzer, Frank Pfenning, and Bernardo Toninho</i> .....                                | 30:1–30:18 |
| Verification of Immediate Observation Population Protocols<br><i>Javier Esparza, Pierre Ganty, Rupak Majumdar, and Chana Weil-Kennedy</i> .....                       | 31:1–31:16 |

|   |            |
|---|------------|
| The Satisfiability Problem for Unbounded Fragments of Probabilistic CTL<br><i>Jan Křetínský and Alexej Rotar</i> .....  | 32:1–32:16 |
| Automatic Analysis of Expected Termination Time for Population Protocols<br><i>Michael Blondin, Javier Esparza, and Antonín Kučera</i> .....  | 33:1–33:16 |
| On Runtime Enforcement via Suppressions<br><i>Luca Aceto, Ian Cassar, Adrian Francalanza, and Anna Ingólfssdóttir</i> .....   | 34:1–34:17 |
| Regular Separability of Well-Structured Transition Systems<br><i>Wojciech Czerwiński, Sławomir Lasota, Roland Meyer, Sebastian Muskalla,<br/>K. Narayan Kumar, and Prakash Saivasan</i> ..... | 35:1–35:18 |
| Separable GPL: Decidable Model Checking with More Non-Determinism<br><i>Andrey Gorlin and C. R. Ramakrishnan</i> .....  | 36:1–36:16 |
| (Metric) Bisimulation Games and Real-Valued Modal Logics for Coalgebras<br><i>Barbara König and Christina Mika-Michalski</i> .....  | 37:1–37:17 |
| The Complexity of Rational Synthesis for Concurrent Games<br><i>Rodica Condurache, Youssouf Oualhadj, and Nicolas Troquard</i> .....  | 38:1–38:15 |
| Logics Meet 1-Clock Alternating Timed Automata<br><i>Shankara Narayanan Krishna, Khushraj Madnani, and Paritosh K. Pandya</i> .....   | 39:1–39:17 |
| Progress-Preserving Refinements of CTA<br><i>Massimo Bartoletti, Laura Bocchi, and Maurizio Murgia</i> .....  | 40:1–40:19 |
| Automated Detection of Serializability Violations Under Weak Consistency<br><i>Kartik Nagar and Suresh Jagannathan</i> .....  | 41:1–41:18 |
| Effective Divergence Analysis for Linear Recurrence Sequences<br><i>Shaull Almagor, Brynmor Chapman, Mehran Hosseini, Joël Ouaknine, and<br/>James Worrell</i> .....                          | 42:1–42:15 |





## ■ Preface

This volume contains the proceedings of the 29th Conference on Concurrency Theory, which was held in Beijing, China, on September 4–7, 2018. CONCUR 2018 was organised by the Institute of Software, Chinese Academy of Sciences.

CONCUR is a forum for the development and dissemination of leading research in concurrency theory and its applications. Its aim is to bring together researchers, developers, and students to exchange and discuss latest theoretical developments and learn about challenging practical problems. CONCUR is the reference annual event for researchers in the field.

The principal topics include basic models of concurrency such as abstract machines, domain-theoretic models, game-theoretic models, process algebras, graph transformation systems, Petri nets, hybrid systems, mobile and collaborative systems, probabilistic systems, real-time systems, biology-inspired systems, and synchronous systems; logics for concurrency such as modal logics, probabilistic and stochastic logics, temporal logics, and resource logics; verification and analysis techniques for concurrent systems such as abstract interpretation, atomicity checking, model checking, race detection, pre-order and equivalence checking, run-time verification, state-space exploration, static analysis, synthesis, testing, theorem proving, type systems, and security analysis; distributed algorithms and data structures: design, analysis, complexity, correctness, fault tolerance, reliability, availability, consistency, self-organisation, self-stabilisation, protocols. The theoretical foundations of more applied topics like architectures, execution environments, and software development for concurrent systems such as geo-replicated systems, communication networks, multiprocessor and multi-core architectures, shared and transactional memory, resource management and awareness, compilers and tools for concurrent programming, programming models such as component-based, object- and service-oriented can also be found at CONCUR.

This edition of the conference attracted 101 full paper submissions, and we thank the authors for their interest in CONCUR 2018. After careful reviewing and discussions, the Program Committee selected 37 papers for presentation at the conference. Each submission was reviewed by at least three reviewers who wrote detailed evaluations and gave insightful comments. We warmly thank the members of the Program Committee and the additional reviewers for their excellent work, including the constructive discussions. The full list of reviewers is available as part of these proceedings.

The conference programme was greatly enriched by the invited talks by Moshe Vardi, Yuxin Deng, Rob van Glabbeek, and Bow-Yaw Wang, as well as the tutorial delivered by Ichiro Hasuo. We thank the speakers for having accepted our invitation and their excellent presentations.

This year, the conference was jointly organised with the 16th International Conference on Formal Modelling and Analysis of Timed Systems (FORMATS), the 15th International Conference on Quantitative Evaluation of SysTems (QEST), and the fourth Symposium on Dependable Software Engineering (SETTA) in an overarching event, CONFESTA, organised by the Institute of Software, Chinese Academy of Sciences.

CONFESTA included four more satellite events: the combined 25th International Workshop on Expressiveness in Concurrency and 15th Workshop on Structural Operational Semantics (EXPRESS/SOS), the 3rd International workshop on Timing Performance engineering for Safety critical systems (TIPS'18), the 7th IFIP WG 1.8 Workshop on Trends in Concurrency Theory (TRENDS), and the 8th Young Researchers Workshop on Concurrency

29th International Conference on Concurrency Theory (CONCUR 2018).

Editors: Sven Schewe and Lijun Zhang



Leibniz International Proceedings in Informatics  
LIPIC Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Theory (YR-CONCUR). as well as a number of tutorials. CONFESTA was preceded by two further associated events, a Summer School on Formal Methods and a CAP Project Workshop.

The CONCUR proceedings are available for open access via LIPIcs, and we thank the staff from Schloss Dagstuhl, in particular Michael Wagner, for helping us with the preparation. Last, but not least, we thank the authors and the participants for making this year's CONCUR a successful and inspiring event.

Sven Schewe (University of Liverpool)

Lijun Zhang (Institute of Software, Chinese Academy of Sciences)

## ■ Committees

### Programme Committee

Parosh Abdulla

Uppsala University (Sweden)

Christel Baier

TU Dresden (Germany)

Roderick Bloem

Graz University of Technology (Austria)

Ahmed Bouajjani

IRIF, University Paris Diderot (France)

Taolue Chen

Birkbeck, University of London (UK)

Yu-Fang Chen

Academia Sinica (Taiwan)

Alessandro Cimatti

Fondazione Bruno Kessler (Italy)

Pedro R. D'Argenio

Universidad Nacional de Córdoba -  
CONICET (Argentina)

Josée Desharnais

Université Laval (Canada)

Wan Fokkink

Vrije Universiteit Amsterdam (The  
Netherlands)

Erich Grädel

RWTH Aachen University (Germany)

Ichiro Hasuo

National Institute of Informatics (Japan)

Fei He

Tsinghua University (China)

Anna Ingólfssdóttir

Reykjavík University (Iceland)

Stefan Kiefer

University of Oxford (UK)

Shankara Narayanan Krishna

IIT Bombay (India)

Antonín Kučera

Masaryk University (Czech Republic)

Salvatore La Torre

Università degli Studi di Salerno (Italy)

Jérôme Leroux

CNRS (France)

Parthasarathy Madhusudan

University of Illinois at Urbana-Champaign  
(USA)

Rupak Majumdar

MPI-SWS (Germany)

Radu Mardare

Aalborg University (Denmark)

Roland Meyer

TU Braunschweig (Germany)

Angelo Montanari

University of Udine (Italy)

Sriram Sankaranarayanan

University of Colorado, Boulder (USA)

Alexandra Silva

University College London (UK)

Ana Sokolova

University of Salzburg (Austria)

Mariëlle Stoelinga

University of Twente (The Netherlands)

Franck van Breugel

York University (Canada)

Verena Wolf

Saarland University (Germany)

### Co-Chairs

Sven Schewe

University of Liverpool (UK)

Lijun Zhang

Institute of Software, Chinese Academy of  
Sciences (China)

29th International Conference on Concurrency Theory (CONCUR 2018).

Editors: Sven Schewe and Lijun Zhang



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

**Steering Committee**

Jos Baeten  
Centrum Wiskunde & Informatica (CWI)  
(The Netherlands)

Pedro R. D'Argenio  
Universidad Nacional de Córdoba  
(Argentina)

Javier Esparza  
Technische Universität München (Germany)

Joost-Pieter Katoen  
RWTH Aachen (Germany)

Kim G. Larsen  
Aalborg University (Denmark)

Ugo Montanari  
Università di Pisa (Italy)

Catuscia Palamidessi  
INRIA and LIX, École Polytechnique  
(France)

**Local Organisers**

Teng Fei  
Institute of Software, Chinese Academy of  
Sciences (China)

David N. Jansen  
Institute of Software, Chinese Academy of  
Sciences (China)

Yongjian Li  
Institute of Software, Chinese Academy of  
Sciences (China)

Yi Lv  
Institute of Software, Chinese Academy of  
Sciences (China)

Andrea Turrini  
Institute of Software, Chinese Academy of  
Sciences (China)

Shuling Wang  
Institute of Software, Chinese Academy of  
Sciences (China)

Peng Wu  
Institute of Software, Chinese Academy of  
Sciences (China)

Bai Xue  
Institute of Software, Chinese Academy of  
Sciences (China)

Rongjie Yan  
Institute of Software, Chinese Academy of  
Sciences (China)

Li Zhang  
Institute of Software, Chinese Academy of  
Sciences (China)

Xueyang Zhu  
Institute of Software, Chinese Academy of  
Sciences (China)

**Local Organisation Chair**

Zhilin Wu  
Institute of Software, Chinese Academy of  
Sciences (China)

**Publicity Co-Chairs**

Ernst Moritz Hahn  
University of Liverpool (UK)

Meng Sun  
Peking University (China)

## ■ List of External Reviewers

|                        |                        |
|------------------------|------------------------|
| Luca Aceto             | David de Frutos Escrig |
| Dan Alistarh           | Giorgio Delzanno       |
| Baskar Anguraj         | Stéphane Demri         |
| Stavros Aronis         | Catalin Dima           |
| S. Arun-Kumar          | Brijesh Dongol         |
| Mohamed Faouzi Atig    | Cezara Dragoi          |
| Giorgio Bacci          | Clemens Dubslaff       |
| Giovanni Bacci         | Jérémy Dubut           |
| Michael Backenköhler   | Constantin Enea        |
| Eric Badouel           | Gidon Ernst            |
| Nikhil Balaji          | Marco Faella           |
| Borja Balle            | Uli Fahrenberg         |
| Francesco Belardinelli | Nathanaël Fijalkow     |
| Dietmar Berwanger      | Brendan Fong           |
| František Blahoudek    | Ignacio Fábregas       |
| Laura Bocchi           | Pierre Ganty           |
| Marco Bozzano          | Paul Gastin            |
| Laura Bozzelli         | Simon Gay              |
| Tomas Brazdil          | Sergey Goncharov       |
| Simon Castellan        | Alexander Graf-Brill   |
| Ilaria Castellani      | Alberto Griggio        |
| Pablo Castro           | Gerrit Grossmann       |
| Didier Caucal          | Stefan Göller          |
| Mariano Ceccato        | Vojtěch Havlena        |
| Rohit Chadha           | Frédéric Herbreteau    |
| Liqian Chen            | Lukas Holik            |
| Xin Chen               | Hung-Wei Hsu           |
| Peter Chini            | Omar Inverso           |
| Corina Cirstea         | Ahmed Irfan            |
| Emanuele D’Osualdo     | Rinat Iusupov          |
| Vrunda Dave            | Petr Jancar            |

29th International Conference on Concurrency Theory (CONCUR 2018).

Editors: Sven Schewe and Lijun Zhang



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

**0:xiv External Reviewers**

|                           |                     |
|---------------------------|---------------------|
| Nils Jansen               | Gustavo Petri       |
| Claude Jard               | Thomas Place        |
| Peter Gjøøl Jensen        | Gabriele Puppis     |
| Thomas Kahl               | David Pym           |
| Benjamin Lucien Kaminski  | Jorge A. Pérez      |
| Anja Karl                 | Hadi Ravanbakhsh    |
| Joachim Klein             | Vojtech Rehak       |
| Bettina Koenighofer       | Antoine Rollet      |
| Clemens Kupke             | Jurriaan Rot        |
| Marcel Kyas               | Marco Roveri        |
| Charalampos Kyriakopoulos | Enno Ruijters       |
| Rom Langerak              | Prakash Saivasan    |
| Kung-Kiu Lau              | Pietro Sala         |
| Marijana Lazic            | Tetsuya Sato        |
| Ondřej Lengál             | Sylvain Schmitz     |
| Christoph Lenzen          | Lutz Schröder       |
| Hsin-Hung Lin             | Roberto Segala      |
| Alexander Lück            | Ilya Sergey         |
| Khushraj Madnani          | Mahsa Shirmohammadi |
| Konstantinos Mamouras     | David Sprunger      |
| Richard Mayr              | Daniel Stan         |
| Filip Mazowiecki          | Caleb Stanford      |
| Alberto Molinari          | Ivan Stojic         |
| J. Garrett Morris         | Eijiro Sumii        |
| Mohammad Mousavi          | Grégoire Sutre      |
| Sergio Mover              | Toru Takisaka       |
| Sebastian Muskalla        | Qiyi Tang           |
| Elisabeth Neumann         | Peter Thiemann      |
| Jan Obdrzalek             | Chun Tian           |
| Oded Padon                | Simone Tini         |
| Vincent Penelle           | Stefano Tonetta     |
| Adriano Peron             | Tigran Tonoyan      |
| Kirstin Peters            | Andrea Turrini      |

Henning Urbat

Jaco van de Pol

Rob van Glabbeek

Dominik Velan

Walter Vogler

Masaki Waga

Hengfeng Wei

Tim Willemse

Sebastian Wolff

Nicolás Wolovick

James Worrell

Bo Wu

Zhilin Wu

Sascha Wunderlich

Akihisa Yamada

Shaofa Yang

Fabio Zanasi

Georg Zetsche

29th International Conference on Concurrency Theory (CONCUR 2018).

Editors: Sven Schewe and Lijun Zhang



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



## ■ List of Authors

Parosh Aziz Abdulla  
Uppsala University  
Sweden  
parosh@it.uu.se

Luca Aceto  
Reykjavik University  
Iceland  
luca@ru.is

Shaul Almagor  
Oxford University, UK  
United Kingdom  
shaull.almagor@mail.huji.ac.il

Mohamed Faouzi Atig  
Uppsala University  
Sweden  
mohamed\_faouzi.atig@it.uu.se

Stephanie Balzer  
Carnegie Mellon University  
United States  
balzers@cs.cmu.edu

Massimo Bartoletti  
Università degli Studi di Cagliari  
Italy  
bart@unica.it

Michael Blondin  
Technical University of Munich  
Germany  
blondin@in.tum.de

Laura Bocchi  
University of Kent  
United Kingdom  
L.Bocchi@kent.ac.uk

Benedikt Bollig  
LSV, ENS Cachan, CNRS  
France  
bollig@lsv.ens-cachan.fr

Filippo Bonchi  
University of Pisa  
Italy  
filippo.bonchi@ens-lyon.fr

Ahmed Bouajjani  
IRIF, University Paris Diderot  
France  
abou@irif.fr

Tomasz Brengos  
Warsaw University of Technology  
Poland  
t.brengos@mini.pw.edu.pl

Véronique Bruyère  
University of Mons  
Belgium  
veronique.bruyere@umons.ac.be

Benjamin Cabrera  
University of Duisburg-Essen  
Germany  
benjamin.cabrera@uni-due.de

Ian Cassar  
University of Malta & Reykjavik University  
Malta & Iceland  
ian.cassar.10@um.edu.mt

Brynmor Chapman  
MIT CSAIL & EECS, Cambridge, MA  
United States  
brynmor@mit.edu

Krishnendu Chatterjee  
Institute of Science and Technology (IST)  
Austria  
krish.chat@gmail.com

Radu Ciobanu  
The University of Edinburgh  
United Kingdom  
R.Ciobanu@sms.ed.ac.uk

Rodica Condurache  
Universite Paris Est, Creteil, and  
Universite Libre de Bruxelles  
France & Belgium  
rodica.bozianu@gmail.com

Wojciech Czerwinski  
University of Warsaw  
Poland  
wczerwin@mimuw.edu.pl

Yuxin Deng  
East China Normal University  
China  
yxdeng@sei.ecnu.edu.cn

Alastair Donaldson  
Imperial College London  
United Kingdom  
alastair.donaldson@imperial.ac.uk

Amina Doumane  
CNRS - ENS Lyon  
France  
Amina.Doumane@ens-lyon.fr

Javier Esparza  
Technical University of Munich  
Germany  
esparza@in.tum.de

Hugues Evrard  
Imperial College London  
United Kingdom  
h.evrard@imperial.ac.uk

Marie Fortin  
LSV, ENS Paris-Saclay,  
CNRS, Université Paris-Saclay  
France  
marie.fortin@lsv.fr

Paulin Fournier  
LS2N, Université de Nantes  
France  
paulin.fournier@gmail.com

Adrian Francalanza  
University of Malta  
Malta  
adrian.francalanza@um.edu.mt

Pierre Ganty  
IMDEA Software Institute  
Spain  
pierre.ganty@imdea.org

Paul Gastin  
LSV, ENS Paris-Saclay,  
CNRS, Université Paris-Saclay  
France  
gastin@lsv.fr

Hugo Gimbert  
CNRS, LABRI, Bordeaux  
France  
hugo.gimbert@labri.fr

Amir Kafshdar Goharshady  
IST Austria  
Austria  
goharshady@ist.ac.at

Sergey Goncharov  
FAU Erlangen-Nürnberg  
Germany  
Sergey.Goncharov@fau.de

Andrey Gorlin  
Stony Brook University  
United States  
agorlin@cs.stonybrook.edu

Radu Grigore  
University of Kent  
United Kingdom  
radugrigore@gmail.com

Christoph Haase  
University of Oxford  
United Kingdom  
Christoph.Haase@cs.ox.ac.uk

Ichiro Hasuo  
National Institute of Informatics  
Japan  
i.hasuo@acm.org

Quentin Hautem  
UMONS  
Belgium  
quentin.hautem@umons.ac.be

Reiko Heckel  
University of Leicester  
United Kingdom  
reiko@mcs.le.ac.uk

Tobias Heindel  
DIKU, University of Copenhagen  
Denmark  
tobias.heindel@googlemail.com

Thomas A. Henzinger  
IST Austria  
Austria  
tah@ist.ac.at

Piotr Hofman  
University of Warsaw  
Poland  
piotrek.hofman@gmail.com

Mehran Hosseini  
University of Oxford  
United Kingdom  
mehran.hosseini@cs.ox.ac.uk

Rasmus Ibsen-Jensen  
IST Austria  
Austria  
ribsen@ist.ac.at

Anna Ingólfssdóttir  
Reykjavik University  
Iceland  
annai@ru.is

Suresh Jagannathan  
Purdue University  
United States  
suresh@cs.purdue.edu

Julian Jakob  
FAU Erlangen-Nürnberg  
Germany  
Julian.Jakob@fau.de

Kartik Nagar  
Purdue University  
United States  
nagark@purdue.edu

Stefan Kiefer  
University of Oxford  
United Kingdom  
stefan.kiefer@cs.ox.ac.uk

Barbara König  
Universität Duisburg-Essen  
Germany  
barbara\_koenig@uni-due.de

Igor Konnov  
INRIA Nancy (LORIA)  
France  
igor.konnov@inria.fr

Bernhard Kragl  
IST Austria  
Austria  
bkragl@ist.ac.at

Jan Křetínský  
Technical University of Munich  
Germany  
jan.kretinsky@gmail.com

Antonin Kučera  
Masaryk University  
Czechia  
tony@fi.muni.cz

Jure Kukovec  
Vienna University of Technology  
Austria  
jkukovec@forsyte.at

K Narayan Kumar  
Chennai Mathematical Institute  
India  
kumar@cmi.ac.in

Sławomir Lasota  
University of Warsaw  
Poland  
sl@mimuw.edu.pl

Khushraj Madnani  
IIT Bombay  
India  
khushraj@cse.iitb.ac.in

Rupak Majumdar  
Max Planck Institute for Software Systems  
Germany  
rupak@mpi-sws.org

Richard Mayr  
The University of Edinburgh  
United Kingdom  
rmayr@staffmail.ed.ac.uk

Filip Mazowiecki  
LaBRI, Université de Bordeaux  
France  
filip.mazowiecki@u-bordeaux.fr

Roland Meyer  
TU Braunschweig  
Germany  
roland.meyer@tu-bs.de

Jakub Michaliszyn  
University of Wrocław  
Poland  
jakub.michaliszyn@gmail.com

Christina Mika-Michalski  
University Duisburg-Essen  
Germany  
christine.mika@uni-due.de

Sayan Mukherjee  
Chennai Mathematical Institute  
India  
sayanm@cmi.ac.in

Maurizio Murgia  
University of Kent  
United Kingdom  
M.Murgia@kent.ac.uk

Sebastian Muskalla  
TU Braunschweig  
Germany  
s.muskalla@tu-bs.de

Renato Neves  
INESC TEC (HASLab) and  
University of Minho  
Portugal  
nevrenato@di.uminho.pt

Jan Otop  
University of Wroclaw  
Poland  
jotop@cs.uni.wroc.pl

Joel Ouaknine  
Max Planck Institute for Software Systems  
Germany  
joel@mpi-sws.org

Youssef Oualhadj  
Université Paris Est Créteil  
France  
youssef.oualhadj@lacl.fr

Paritosh Pandya  
TIFR  
India  
pandya@tifr.res.in

Guillermo Perez  
Université libre de Bruxelles  
Belgium  
gperezme@ulb.ac.be

Daniela Petrisan  
Université Paris Diderot - Paris 7  
France  
daniela.petrisan@gmail.com

Frank Pfenning  
Carnegie Mellon University  
United States  
fp@cs.cmu.edu

Andy Polyakov  
The OpenSSL project  
Sweden  
appro@openssl.org

Damien Pous  
CNRS - ENS Lyon  
France  
Damien.Pous@ens-lyon.fr

Pavithra Prabhakar  
Kansas State University  
United States  
pprabhakar@ksu.edu

Shaz Qadeer  
Microsoft  
United States  
qadeer@microsoft.com

C. R. Ramakrishnan  
Stony Brook University  
United States  
cram@cs.stonybrook.edu

Jean-Francois Raskin  
Université Libre de Bruxelles  
Belgium  
jraskin@ulb.ac.be

Nima Roohi  
University of Pennsylvania  
United States  
roohi2@cis.upenn.edu

Christian Rosenke  
University of Rostock  
Germany  
christian.rosenke@uni-rostock.de

Alexej Rotar  
Technical University of Munich  
Germany  
alexejrotar@gmail.com

Krishna S  
IIT Bombay  
India  
krishnas@cse.iitb.ac.in

Prakash Saivasan  
TU Braunschweig  
Germany  
p.saivasan@tu-bs.de

Tyler Sorensen  
Imperial College London  
United Kingdom  
t.sorensen15@imperial.ac.uk

B Srivathsan  
Chennai Mathematical Institute  
India  
sri@cmi.ac.in

Qiyi Tang  
York University, Toronto  
Canada  
qiyitang@eecs.yorku.ca

Bernardo Toninho  
Universidade NOVA de Lisboa  
Portugal  
btoninho@gmail.com

Patrick Totzke  
University of Edinburgh  
United Kingdom  
p.totzke@ed.ac.uk

Ronny Tredup  
University of Rostock  
Germany  
ronny.tredup2@uni-rostock.de

Nicolas Troquard  
Free University of Bozen  
Italy  
nicolas.troquard@unibz.it

Ming-Hsien Tsai  
Academia Sinica  
Taiwan  
mhstai208@gmail.com

Franck van Breugel  
York University, Toronto  
Canada  
franck@eecs.yorku.ca

Rob van Glabbeek  
CSIRO  
Australia  
rvg@cs.stanford.edu

Moshe Y. Vardi  
Rice University  
USA  
vardi@cs.rice.edu

Yaron Velner  
Tel Aviv University  
Israel  
yaron172@yahoo.com

Mahesh Viswanathan  
University of Illinois at Urbana-Champaign  
United States  
vmahesh@illinois.edu

Bow-Yaw Wang  
Academia Sinica  
Taiwan  
bywang@iis.sinica.edu.tw

Josef Widder  
Vienna University of Technology  
Austria  
widder@forsyte.at

James Worrell  
University of Oxford  
United Kingdom  
jbw@cs.ox.ac.uk

Bo-Yin Yang  
Academia Sinica  
Taiwan  
by@crypto.tw

Georg Zetsche  
IRIF, CNRS & Université Paris-Diderot  
France  
zetsche@irif.fr



# On Runtime Enforcement via Suppressions

**Luca Aceto**

Gran Sasso Science Institute, L'Aquila, Italy; and  
Reykjavik University, Reykjavik, Iceland  
luca.aceto@gssi.it

**Ian Cassar**

Reykjavik University, Reykjavik Iceland; and  
University of Malta, Msida, Malta  
ianc@ru.is

**Adrian Francalanza**

University of Malta, Msida, Malta  
adrian.francalanza@um.edu.mt

**Anna Ingólfssdóttir**

Reykjavik University, Reykjavik, Iceland  
annai@ru.is

---

## Abstract

Runtime enforcement is a dynamic analysis technique that uses monitors to enforce the behaviour specified by some correctness property on an executing system. The enforceability of a logic captures the extent to which the properties expressible via the logic can be enforced at runtime. We study the enforceability of Hennessy-Milner Logic with Recursion ( $\mu$ HML) with respect to suppression enforcement. We develop an operational framework for enforcement which we then use to formalise when a monitor enforces a  $\mu$ HML property. We also show that the safety syntactic fragment of the logic, sHML, is enforceable by providing an automated synthesis function that generates correct suppression monitors from sHML formulas.

**2012 ACM Subject Classification** Theory of computation  $\rightarrow$  Logic and verification, Software and its engineering  $\rightarrow$  Software verification, Software and its engineering  $\rightarrow$  Dynamic analysis

**Keywords and phrases** Enforceability, Suppression Enforcement, Monitor Synthesis, Logic

**Digital Object Identifier** 10.4230/LIPIcs.CONCUR.2018.34

**Related Version** <https://arxiv.org/abs/1807.01004>

**Acknowledgements** The research work disclosed in this publication is partially supported by the projects “Developing Theoretical Foundations for Runtime Enforcement” (184776-051) and “TheoFoMon: Theoretical Foundations for Monitorability” (163406-051) of the Icelandic Research Fund, and by the Endeavour Scholarship Scheme (Malta), part-financed by the European Social Fund (ESF) – Operational Programme II – Cohesion Policy 2014-2020.

## 1 Introduction

Runtime monitoring [22, 24] is a dynamic analysis technique that is becoming increasingly popular in the turbid world of software development. It uses code units called *monitors* to aggregate system information, compare system execution against correctness specifications, or steer the execution of the observed system. The technique has been used effectively to offload certain verification tasks to a post-deployment phase, thus complementing other



© Luca Aceto, Ian Cassar, Adrian Francalanza, and Anna Ingólfssdóttir;  
licensed under Creative Commons License CC-BY

29th International Conference on Concurrency Theory (CONCUR 2018).

Editors: Sven Schewe and Lijun Zhang; Article No. 34; pp. 34:1–34:17

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

(static) analysis techniques in multi-pronged verification strategies – see e.g., [6, 12, 27, 18, 28]. *Runtime enforcement* (RE) [33, 34, 21] is a specialized monitoring technique, used to ensure that the behaviour of a system-under-scrutiny (SuS) is *always* in agreement with some correctness specification. It employs a specific kind of monitor (referred to as a *transducer* [9, 42, 4] or an *edit-automaton* [33, 34]) to anticipate incorrect behaviour and counter it. Such a monitor thus acts as a proxy between the SuS and the surrounding environment interacting with it, encapsulating the system to form a composite (monitored) system: at runtime, the monitor *transforms* any incorrect executions exhibited by the SuS into correct ones by either *suppressing*, *inserting* or *replacing* events on behalf of the system.

We extend a recent line of research [25, 24, 2, 1] and study RE approaches that adopt a *separation of concerns* between the correctness specification, describing *what* properties the SuS should satisfy, and the monitor, describing *how* to enforce these properties on the SuS. Our work considers system properties expressed in terms of the process logic  $\mu$ HML [30, 32], and explores what properties can be operationally enforced by monitors that can suppress system behaviour. A central element for the realisation of such an approach is the *synthesis* function: it automates the translation from the *declarative*  $\mu$ HML specifications to *algorithmic* descriptions formulated as executable monitors. Since analysis tools ought to form part of the trusted computing base, enforcement monitoring should be, in and of itself, correct. However, it is unclear what is to be expected of the synthesised monitor to adequately enforce a  $\mu$ HML formula. Nor is it clear for which type of specifications should this approach be expected to work effectively – it has been well established that a number of properties are *not* monitorable [15, 40, 16, 25, 2] and it is therefore reasonable to expect similar limits in the case of enforceability [19]. We therefore study the relationship between  $\mu$ HML specifications and suppression monitors for enforcement, which allows us to address the above-mentioned concerns and make the following contributions:

**Modelling:** We develop a general framework for enforcement instrumentation that is parametrisable by any system behaviour that is expressed via labelled transitions, and can express suppression, insertion and replacement enforcement, Figure 2.

**Correctness:** We give formal definitions for asserting when a monitor correctly enforces a formula defined over labelled transition systems, Definitions 3 and 8. These definitions are parametrisable with respect to an instrumentation relation, an instance of which is our enforcement framework of Figure 2.

**Expressiveness:** We provide enforceability results, Theorems 14 and 18 (but also Proposition 24), by identifying a subset of  $\mu$ HML formulas that can be (correctly) enforced by suppression monitors.

As a by-product of this study, we also develop a formally-proven correct synthesis function, Definition 12, that then can be used for tool construction, along the lines of [8, 7].

The setup selected for our study serves a number of purposes. For starters, the chosen logic,  $\mu$ HML, is a branching-time logic that allows us to investigate enforceability for properties describing computation graphs. Second, the use of a highly expressive logic allows us to achieve a good degree of generality for our results, and so, by working in relation to logics like  $\mu$ HML (a reformulation of the  $\mu$ -calculus), our work would also apply to other widely used logics (such as LTL and CTL [17]) that are embedded within this logic. Third, since the logic is verification-technique agnostic, it fits better with the realities of software verification in the present world, where a *variety* of techniques (e.g., model-checking and testing) straddling both pre- and post-deployment phases are used. In such cases, knowing which properties can be verified statically and which ones can be monitored for and enforced at runtime is crucial for devising effective multi-pronged verification strategies. Equipped



**Syntax**

$$\begin{array}{llll}
\varphi, \psi \in \mu\text{HML} ::= \text{tt} & (\text{truth}) & | \text{ff} & (\text{falsehood}) & | \bigvee_{i \in I} \varphi_i & (\text{disjunction}) \\
& | \bigwedge_{i \in I} \varphi_i & (\text{conjunction}) & | \langle \{p, c\} \rangle \varphi & (\text{possibility}) & | [\{p, c\}] \varphi & (\text{necessity}) \\
& | \min X. \varphi & (\text{least fp.}) & | \max X. \varphi & (\text{greatest fp.}) & | X & (\text{fp. variable})
\end{array}$$

**Semantics**

$$\begin{array}{lll}
\llbracket \text{tt}, \rho \rrbracket \stackrel{\text{def}}{=} \text{Sys} & \llbracket \text{ff}, \rho \rrbracket \stackrel{\text{def}}{=} \emptyset & \llbracket X, \rho \rrbracket \stackrel{\text{def}}{=} \rho(X) \\
\llbracket \bigwedge_{i \in I} \varphi_i, \rho \rrbracket \stackrel{\text{def}}{=} \bigcap_{i \in I} \llbracket \varphi_i, \rho \rrbracket & \llbracket \max X. \varphi, \rho \rrbracket \stackrel{\text{def}}{=} \bigcup \{ S \mid S \subseteq \llbracket \varphi, \rho[X \mapsto S] \rrbracket \} \\
\llbracket \bigvee_{i \in I} \varphi_i, \rho \rrbracket \stackrel{\text{def}}{=} \bigcup_{i \in I} \llbracket \varphi_i, \rho \rrbracket & \llbracket \min X. \varphi, \rho \rrbracket \stackrel{\text{def}}{=} \bigcap \{ S \mid \llbracket \varphi, \rho[X \mapsto S] \rrbracket \subseteq S \} \\
\llbracket [\{p, c\}] \varphi, \rho \rrbracket \stackrel{\text{def}}{=} \{ s \mid (\forall \alpha, r. s \xrightarrow{\alpha} r \text{ and } (\exists \sigma \cdot \text{mtch}(p, \alpha) = \sigma \text{ and } c\sigma \Downarrow \text{true})) \text{ implies } q \in \llbracket \varphi\sigma, \rho \rrbracket \} \\
\llbracket \langle \{p, c\} \rangle \varphi, \rho \rrbracket \stackrel{\text{def}}{=} \{ s \mid \exists \alpha, r, \sigma \cdot (s \xrightarrow{\alpha} r \text{ and } \text{mtch}(p, \alpha) = \sigma \text{ and } c\sigma \Downarrow \text{true} \text{ and } q \in \llbracket \varphi\sigma, \rho \rrbracket) \}
\end{array}$$

■ **Figure 1**  $\mu\text{HML}$  Syntax and Semantics.

with such knowledge, one could also employ standard techniques [36, 5, 31] to decompose a non-enforceable property into a collection of smaller properties, a subset of which can then be enforced at runtime.

**Structure of the paper.** Section 2 revisits labelled transition systems and our touchstone logic,  $\mu\text{HML}$ . The operational model for enforcement monitors and instrumentation is given in Section 3. In Section 4 we formalise the interdependent notions of correct enforcement and enforceability. These act as a foundation for the development of a synthesis function in Section 5, that produces *correct-by-construction* monitors. In Section 6 we consider alternative definitions for enforceability for logics with a specific additional interpretation, and show that our proposed synthesis function is still correct with respect to the new definition. Section 7 concludes and discusses related work.

## 2 Preliminaries

**The Model.** We assume systems described as *labelled transition systems* (LTSs), triples  $\langle \text{SYS}, \text{ACT} \cup \{\tau\}, \rightarrow \rangle$  consisting of a set of *system states*,  $s, r, q \in \text{SYS}$ , a set of *observable actions*,  $\alpha, \beta \in \text{ACT}$ , and a distinguished silent action  $\tau \notin \text{ACT}$  (where  $\mu \in \text{ACT} \cup \{\tau\}$ ), and a *transition relation*,  $\rightarrow \subseteq (\text{SYS} \times \text{ACT} \cup \{\tau\} \times \text{SYS})$ . We write  $s \xrightarrow{\mu} r$  in lieu of  $(s, \mu, r) \in \rightarrow$ , and use  $s \xRightarrow{\mu} s'$  to denote weak transitions representing  $s(\xrightarrow{\tau})^* \cdot \xrightarrow{\mu} \cdot (\xrightarrow{\tau})^* s'$ . We refer to  $s'$  as a  $\mu$ -derivative of  $s$ . Traces,  $t, u \in \text{ACT}^*$  range over (finite) sequences of observable actions, and we write  $s \xRightarrow{t} r$  to denote a sequence of weak transitions  $s \xrightarrow{\alpha_1} \dots \xrightarrow{\alpha_n} r$  for  $t = \alpha_1, \dots, \alpha_n$ . We also assume the classic notion of *strong bisimilarity* [39, 43] for our model,  $s \sim r$ , using it as our touchstone system equivalence. The syntax of the regular fragment of CCS [39] is occasionally used to concisely describe LTSs in our examples.

**The Logic.** We consider a slightly generalised version of  $\mu\text{HML}$  [32, 3] that uses *symbolic actions* of the form  $\{p, c\}$ . *Patterns*,  $p$ , abstract over actions using *data variables*  $d, e, f \in \text{VAR}$ ; in a pattern, they may either occur free,  $d$ , or as binders,  $(d)$  where a *closed pattern* is one without any free variables. We assume a (partial) *matching function* for closed patterns  $\text{mtch}(p, \alpha)$  that returns a substitution  $\sigma$  (when successful) mapping variables in  $p$  to the corresponding values in  $\alpha$ , i.e., if we instantiate every bound variable  $d$  in  $p$  with

$\sigma(d)$  we obtain  $\alpha$ . The *filtering condition*,  $c$ , contains variables found in  $p$  and evaluates wrt. the substitutions returned by successful matches. Put differently, a *closed* symbolic action  $\{p, c\}$  is one where  $p$  is closed and  $\mathbf{fv}(c) \subseteq \mathbf{bv}(p)$ ; it denotes the *set* of actions  $\llbracket \{p, c\} \rrbracket \stackrel{\text{def}}{=} \{ \alpha \mid \exists \sigma \cdot \text{mtch}(p, \alpha) = \sigma \text{ and } c\sigma \Downarrow \text{true} \}$  and allows more adequate reasoning about LTSs with infinite actions (e.g., actions carrying data from infinite domains).

The logic syntax is given in Figure 1 and assumes a countable set of logical variables  $X, Y \in \text{LVAR}$ . Apart from standard logical constructs such as conjunctions and disjunctions ( $\bigwedge_{i \in I} \varphi_i$  describes a *compound* conjunction,  $\varphi_1 \wedge \dots \wedge \varphi_n$ , where  $I = \{1, \dots, n\}$  is a finite set of indices, and similarly for disjunctions), and the characteristic greatest and least fixpoints ( $\max X.\varphi$  and  $\min X.\varphi$  bind free occurrences of  $X$  in  $\varphi$ ), the logic uses necessity and possibility modal operators with symbolic actions,  $\llbracket \{p, c\} \rrbracket \varphi$  and  $\langle \{p, c\} \rangle \varphi$ , where  $\mathbf{bv}(p)$  bind free data variables in  $c$  and  $\varphi$ . Formulas in  $\mu\text{HML}$  are interpreted over the system powerset domain where  $S \in \mathcal{P}(\text{SYS})$ . The semantic definition of Figure 1,  $\llbracket \varphi, \rho \rrbracket$ , is given for *both* open and closed formulas. It employs a valuation from logical variables to sets of states,  $\rho \in (\text{LVAR} \rightarrow \mathcal{P}(\text{SYS}))$ , which permits an inductive definition on the structure of the formulas;  $\rho' = \rho[X \mapsto S]$  denotes a valuation where  $\rho'(X) = S$  and  $\rho'(Y) = \rho(Y)$  for all other  $Y \neq X$ . The only non-standard cases are those for the modal formulas, due to the use of symbolic actions. Note that we recover the standard logic for symbolic actions  $\{p, c\}$  whose pattern  $p$  does not contain variables ( $p = \alpha$  for some  $\alpha$ ) and whose condition holds trivially ( $c = \text{true}$ ); in such cases we write  $[\alpha]\varphi$  and  $\langle \alpha \rangle \varphi$  for short. We generally assume *closed* formulas, i.e., without free logical and data variables, and write  $\llbracket \varphi \rrbracket$  in lieu of  $\llbracket \varphi, \rho \rrbracket$  since the interpretation of a closed  $\varphi$  is independent of  $\rho$ . A system  $s$  *satisfies* formula  $\varphi$  whenever  $s \in \llbracket \varphi \rrbracket$  whereas a formula  $\varphi$  is *satisfiable*,  $\varphi \in \text{SAT}$ , whenever there exists a system  $r$  such that  $r \in \llbracket \varphi \rrbracket$ .

► **Example 1.** Consider two systems (a good system,  $s_g$ , and a bad one,  $s_b$ ) implementing a server that interacts on port  $i$ , repeatedly accepting *requests* that are *answered* by outputting on the same port, and terminating the service once a *close* request is accepted (on the same port). Whereas  $s_g$  outputs an answer ( $i!\text{ans}$ ) for every request ( $i?\text{req}$ ),  $s_b$  occasionally refuses to answer a given request (see the underlined branch). Both systems terminate with  $i?\text{cls}$ .

$$s_g = \text{rec } x.(i?\text{req}.i!\text{ans}.x + i?\text{cls}.\text{nil}) \quad s_b = \text{rec } x.(i?\text{req}.i!\text{ans}.x + \underline{i?\text{req}.x} + i?\text{cls}.\text{nil})$$

We can specify that two consecutive requests on port  $i$  indicate invalid behaviour via the  $\mu\text{HML}$  formula  $\varphi_0 \stackrel{\text{def}}{=} \max X.[i?\text{req}].[i!\text{ans}]X \wedge [i?\text{req}]\text{ff}$ ; it defines an invariant property ( $\max X.(\dots)$ ) requiring that whenever a system interacting on  $i$  inputs a request, it cannot input a subsequent request, i.e.,  $[i?\text{req}]\text{ff}$ , unless it outputs an answer beforehand, in which case the formula recurses, i.e.,  $[i!\text{ans}]X$ . Using symbolic actions, we can generalise  $\varphi_0$  by requiring the property to hold for *any* interaction happening on *any* port number *except*  $j$ .

$$\varphi_1 \stackrel{\text{def}}{=} \max X.[\{(d)?\text{req}, d \neq j\}][\{d!\text{ans}, \text{true}\}]X \wedge [\{d?\text{req}, \text{true}\}]\text{ff}$$

In  $\varphi_1$ ,  $(d)?\text{req}$  binds the free occurrences of  $d$  found in  $d \neq j$  and  $[\{d!\text{ans}, \text{true}\}]X \wedge [\{d?\text{req}, \text{true}\}]\text{ff}$ . Using Figure 1, one can check that  $s_g \in \llbracket \varphi_1 \rrbracket$ , whereas  $s_b \notin \llbracket \varphi_1 \rrbracket$  since  $s_b \xrightarrow{i?\text{req}} \cdot \xrightarrow{i?\text{req}} \dots$

### 3 An Operational Model for Enforcement

Our operational mechanism for enforcing properties over systems uses the (symbolic) transducers  $m, n \in \text{TRN}$  defined in Figure 2. The transition rules in Figure 2 assume closed terms, i.e., for every *symbolic-prefix transducer*,  $\{p, c, p'\}.m$ ,  $p$  is closed and  $(\mathbf{fv}(c) \cup \mathbf{fv}(p') \cup \mathbf{fv}(m)) \subseteq \mathbf{bv}(p)$ , and yield an LTS with labels of the form  $\gamma \blacktriangleright \mu$ , where  $\gamma \in (\text{ACT} \cup \{\bullet\})$ . Our syntax

**Syntax**

$$m, n \in \text{TRN} ::= \text{id} \quad | \quad \{p, c, p'\}.m \quad | \quad \sum_{i \in I} m_i \quad | \quad \text{rec } x.m \quad | \quad x$$

**Dynamics**

$$\begin{array}{c} \text{EID} \frac{}{\text{id} \xrightarrow{\mu \blacktriangleright \mu} \text{id}} \quad \text{ESEL} \frac{m_j \xrightarrow{\gamma \blacktriangleright \mu} n_j}{\sum_{i \in I} m_i \xrightarrow{\gamma \blacktriangleright \mu} n_j} \quad j \in I \quad \text{EREC} \frac{m\{\text{rec } x.m/x\} \xrightarrow{\gamma \blacktriangleright \mu} n}{\text{rec } x.m \xrightarrow{\gamma \blacktriangleright \mu} n} \\ \\ \text{ETRN} \frac{\text{mtch}(p, \gamma) = \sigma \quad c\sigma \Downarrow \text{true} \quad \mu = p'\sigma}{\{p, c, p'\}.m \xrightarrow{\gamma \blacktriangleright \mu} m\sigma} \end{array}$$

**Instrumentation**

$$\begin{array}{c} \text{ITRN} \frac{s \xrightarrow{\alpha} s' \quad m \xrightarrow{\alpha \blacktriangleright \mu} n}{m[s] \xrightarrow{\mu} n[s']} \quad \text{IASY} \frac{s \xrightarrow{\tau} s'}{m[s] \xrightarrow{\tau} m[s']} \quad \text{IINS} \frac{m \xrightarrow{\bullet \blacktriangleright \mu} n}{m[s] \xrightarrow{\mu} n[s']} \quad \text{ITER} \frac{s \xrightarrow{\alpha} s' \quad m \xrightarrow{\alpha} m \xrightarrow{\bullet} \cdot}{m[s] \xrightarrow{\alpha} \text{id}[s']} \end{array}$$

■ **Figure 2** A model for transducers ( $I$  is a finite index set and  $m \xrightarrow{\gamma \blacktriangleright \mu}$  means  $\nexists \mu, n \cdot m \xrightarrow{\gamma \blacktriangleright \mu} n$ ).

assumes a well-formedness constraint where for every  $\{p, c, p'\}.m$ ,  $\text{bv}(c) \cup \text{bv}(p') = \emptyset$ . Intuitively, a transition  $m \xrightarrow{\alpha \blacktriangleright \mu} n$  denotes the fact that the transducer in state  $m$  *transforms* the visible action  $\alpha$  (produced by the system) into the action  $\mu$  (which can possibly become silent) and transitions into state  $n$ . In this sense, the transducer action  $\alpha \blacktriangleright \tau$  represents the *suppression* of action  $\alpha$ , action  $\alpha \blacktriangleright \beta$  represents the *replacing* of  $\alpha$  by  $\beta$ , and  $\alpha \blacktriangleright \alpha$  denotes the *identity* transformation. The special case  $\bullet \blacktriangleright \alpha$  encodes the *insertion* of  $\alpha$ , where  $\bullet$  represents that the transition is not induced by any system action.

The key transition rule in Figure 2 is ETRN. It states that the symbolic-prefix transducer  $\{p, c, p'\}.m$  can transform an (extended) action  $\gamma$  into the concrete action  $\mu$ , as long as the action matches with pattern  $p$  with substitution  $\sigma$ ,  $\text{mtch}(p, \gamma) = \sigma$ , and the condition is satisfied by  $\sigma$ ,  $c\sigma \Downarrow \text{true}$  (the matching function is lifted to extended actions and patterns in the obvious way, where  $\text{mtch}(\bullet, \bullet) = \emptyset$ ). In such a case, the transformed action is  $\mu = p'\sigma$ , i.e., the action  $\mu$  resulting from the instantiation of the free data variables in pattern  $p'$  with the corresponding values mapped by  $\sigma$ , and the transducer state reached is  $m\sigma$ . By contrast, in rule EID, the transducer  $\text{id}$  acts as the identity and leaves actions unchanged. The remaining rules are fairly standard and unremarkable.

Figure 2 also describes an *instrumentation* relation which relates the behaviour of the SuS  $s$  with the transformations of a transducer monitor  $m$  that *agrees* with the (observable) actions  $\text{ACT}$  of  $s$ . The term  $m[s]$  thus denotes the resulting *monitored system* whose behaviour is defined in terms of  $\text{ACT} \cup \{\tau\}$  from the system's LTS. Concretely, rule ITRN states that when a system  $s$  transitions with an observable action  $\alpha$  to  $s'$  and the transducer  $m$  can *transform* this action into  $\mu$  and transition to  $n$ , the instrumented system  $m[s]$  transitions with action  $\mu$  to  $n[s']$ . However, when  $s$  transitions with a silent action, rule IASY allows it to do so independently of the transducer. Dually, rule IINS allows the transducer to *insert* an action  $\mu$  independently of  $s$ 's behaviour. Rule ITER is analogous to standard monitor instrumentation rules for premature termination of the transducer [22, 25, 23, 1], and accounts for underspecification of transformations. Thus, if a system  $s$  transitions with an observable action  $\alpha$  to  $s'$ , and the transducer  $m$  does not specify how to transform it ( $m \xrightarrow{\alpha}$ ), nor can it transition to a new transducer state by inserting an action ( $m \xrightarrow{\bullet}$ ), the system is still allowed to transition while the transducer's transformation activity is ceased, i.e., it acts like the identity  $\text{id}$  from that point onwards.

## 34:6 On Runtime Enforcement via Suppressions

► **Example 2.** Consider the insertion transducer  $m_i$  and the replacement transducer  $m_r$  below:

$$\begin{aligned} m_i &\stackrel{\text{def}}{=} \{\bullet, \text{true}, i?\text{req}\}.\{\bullet, \text{true}, i!\text{ans}\}.\text{id} \\ m_r &\stackrel{\text{def}}{=} \text{rec } x.(\{(d)?\text{req}, \text{true}, j?\text{req}\}.x + \{(d)!\text{ans}, \text{true}, j!\text{ans}\}.x + \{(d)?\text{cls}, \text{true}, j?\text{cls}\}.x) \end{aligned}$$

When instrumented with a system,  $m_i$  inserts the two successive actions  $i?\text{req}$  and  $i!\text{ans}$  before behaving as the identity. Concretely in the case of  $s_b$  we can only start the computation as:

$$m_i[s_b] \xrightarrow{i?\text{req}} \{\bullet, \text{true}, i!\text{ans}\}.\text{id}[s_b] \xrightarrow{i!\text{ans}} \text{id}[s_b] \xrightarrow{\alpha} \dots \quad (\text{where } s_b \xrightarrow{\alpha})$$

By contrast,  $m_r$  transforms input actions with either payload  $\text{req}$  or  $\text{cls}$  and output actions with payload  $\text{ans}$  on any port name, into the respective actions on port  $j$ . For instance:

$$m_r[s_b] \xrightarrow{j?\text{req}} m_r[i!\text{ans}.s_b] \xrightarrow{j!\text{ans}} m_r[s_b] \xrightarrow{j?\text{cls}} m_r[\text{nil}]$$

Consider now the two suppression transducers  $m_s$  and  $m_t$  for actions on ports other than  $j$ :

$$\begin{aligned} m_s &\stackrel{\text{def}}{=} \text{rec } x.(\{(d)?\text{req}, d \neq j, \tau\}.x + \{(d)!\text{ans}, \text{true}, d!\text{ans}\}.x) \\ m_t &\stackrel{\text{def}}{=} \text{rec } x.(\{(d)?\text{req}, d \neq j, d?\text{req}\}.\text{rec } y.(\{d!\text{ans}, \text{true}, d!\text{ans}\}.x + \{d?\text{req}, \text{true}, \tau\}.y)) \end{aligned}$$

Monitor  $m_s$  suppresses any requests on ports other than  $j$ , and continues to do so after any answers on such ports. When instrumented with  $s_b$ , we can observe the following behaviour:

$$m_s[s_b] \xrightarrow{\tau} m_s[i!\text{ans}.s_b] \xrightarrow{i!\text{ans}} m_s[s_b] \xrightarrow{\tau} m_s[i!\text{ans}.s_b] \xrightarrow{i!\text{ans}} m_s[s_b] \dots$$

Note that  $m_s$  does not specify a transformation behaviour for when the monitored system produces inputs with payload other than  $\text{req}$ . The instrumentation handles this underspecification by ceasing suppression activity; in the case of  $s_b$  we get  $m_s[s_b] \xrightarrow{i?\text{cls}} \text{id}[\text{nil}]$ . The transducer  $m_t$  performs slightly more elaborate transformations. For interactions on ports other than  $j$ , it suppresses consecutive input requests following any serviced request (i.e., an input on  $\text{req}$  followed by an output on  $\text{ans}$ ) sequence. For  $s_b$  we can observe the following:

$$\begin{aligned} m_t[s_b] &\xrightarrow{i?\text{req}} \text{rec } y.(\{i!\text{ans}, \text{true}, i!\text{ans}\}.m_t + \{i?\text{req}, \text{true}, \tau\}.y)[s_b] \\ &\xrightarrow{\tau} \text{rec } y.(\{i!\text{ans}, \text{true}, i!\text{ans}\}.m_t + \{i?\text{req}, \text{true}, \tau\}.y)[i!\text{ans}.s_b] \xrightarrow{i!\text{ans}} m_t[s_b] \end{aligned}$$

In the sequel, we find it convenient to refer to  $\underline{p}$  as the transformed pattern  $p$  where all the binding occurrences ( $d$ ) are converted to free occurrences  $d$ . As shorthand notation, we elide the second pattern  $p'$  in a transducer  $\{p, c, p'\}.m$  whenever  $p' = \underline{p}$  and simply write  $\{p, c\}.m$ ; note that if  $\text{bv}(p) = \emptyset$ , then  $\underline{p} = p$ . Similarly, we elide  $c$  whenever  $c = \text{true}$ . This allows us to express  $m_t$  from Example 2 as  $\text{rec } x.(\{(d)?\text{req}, d \neq j\}.\text{rec } y.(\{d!\text{ans}\}.x + \{d?\text{req}, \tau\}.y))$ .

### 4 Enforceability

The *enforceability* of a logic rests on the relationship between the semantic behaviour specified by the logic on the one hand, and the ability of the operational mechanism (the transducers and instrumentation of Section 3 in our case) to enforce the specified behaviour on the other.

► **Definition 3 (Enforceability).** A logic  $\mathcal{L}$  is enforceable iff every formula  $\varphi \in \mathcal{L}$  is enforceable. A formula  $\varphi$  is enforceable iff there exists a transducer  $m$  such that  $m$  enforces  $\varphi$ .

Definition 3 depends on what is considered to be an adequate definition for “ $m$  enforces  $\varphi$ ”. It is reasonable to expect that the latter definition should concern *any* system that the transducer  $m$ — hereafter referred to as the *enforcer* — is instrumented with. In particular, for *any* system  $s$ , the resulting composite system obtained from instrumenting the enforcer  $m$  with it should satisfy the property of interest,  $\varphi$ , whenever this property *is satisfiable*.

► **Definition 4 (Sound Enforcement).** Enforcer  $m$  *soundly enforces* a formula  $\varphi$ , denoted as  $\text{senf}(m, \varphi)$ , iff for *all*  $s \in \text{SYS}$ ,  $\varphi \in \text{Sat}$  implies  $m[s] \in \llbracket \varphi \rrbracket$  holds.

► **Example 5.** Recall  $\varphi_1$ ,  $s_g$  and  $s_b$  from Example 1 where  $s_g \in \llbracket \varphi_1 \rrbracket$  (hence  $\varphi_1 \in \text{SAT}$ ) and  $s_b \notin \llbracket \varphi_1 \rrbracket$ . For the enforcers  $m_i$ ,  $m_r$ ,  $m_s$  and  $m_t$  presented in Example 2, we have:

- $m_i[s_b] \notin \llbracket \varphi_1 \rrbracket$ , since  $m_i[s_b] \xrightarrow{i?\text{req}} \cdot \xrightarrow{i!\text{ans}} \text{id}[s_b] \xrightarrow{i?\text{req}} \text{id}[s_b] \xrightarrow{i?\text{req}} \text{id}[s_b]$ . This counter example implies that  $\neg \text{senf}(m_i, \varphi_1)$ .
- $m_r[s_g] \in \llbracket \varphi_1 \rrbracket$  and  $m_r[s_b] \in \llbracket \varphi_1 \rrbracket$ . Intuitively, this is because the ensuing instrumented systems only generate (replaced) actions that are not of concern to  $\varphi_1$ . Since this behaviour applies to any system  $m_r$  is composed with, we can conclude that  $\text{senf}(m_r, \varphi_1)$ .
- $m_s[s_g] \in \llbracket \varphi_1 \rrbracket$  and  $m_s[s_b] \in \llbracket \varphi_1 \rrbracket$  because the resulting instrumented systems never produce inputs with `req` on a port number other than  $j$ . We can thus conclude that  $\text{senf}(m_s, \varphi_1)$ .
- $m_t[s_g] \in \llbracket \varphi_1 \rrbracket$  and  $m_t[s_b] \in \llbracket \varphi_1 \rrbracket$ . Since the resulting instrumentation suppresses consecutive input requests (if any) after any number of serviced requests on any port other than  $j$ , we can conclude that  $\text{senf}(m_t, \varphi_1)$ .

By some measures, sound enforcement is a relatively weak requirement for adequate enforcement as it does not regulate the *extent* of the induced enforcement. More concretely, consider the case of enforcer  $m_s$  from Example 2. Although  $m_s$  manages to suppress the violating executions of system  $s_b$ , thereby bringing it in line with property  $\varphi_1$ , it needlessly modifies the behaviour of  $s_g$  (namely it prohibits it from producing any inputs with `req` on port numbers that are not  $j$ ), even though it satisfies  $\varphi_1$ . Thus, in addition to sound enforcement we require a *transparency* condition for adequate enforcement. The requirement dictates that whenever a system  $s$  already satisfies the property  $\varphi$ , the assigned enforcer  $m$  should not alter the behaviour of  $s$ . Put differently, the behaviour of the enforced system should be behaviourally equivalent to the original system.

► **Definition 6 (Transparent Enforcement).** An enforcer  $m$  is *transparent* when enforcing a formula  $\varphi$ , denoted as  $\text{tenf}(m, \varphi)$ , iff for *all*  $s \in \text{SYS}$ ,  $s \in \llbracket \varphi \rrbracket$  implies  $m[s] \sim s$ .

► **Example 7.** We have already argued – via the counter example  $s_g$ – why  $m_s$  does *not* transparently enforce  $\varphi_1$ . We can also argue easily why  $\neg \text{tenf}(m_r, \varphi_1)$  either: the simple system  $i?\text{req}.\text{nil}$  trivially satisfies  $\varphi_1$  but, clearly, we have the inequality  $m_r[i?\text{req}.\text{nil}] \not\sim i?\text{req}.\text{nil}$  since  $m_r[i?\text{req}.\text{nil}] \xrightarrow{j?\text{req}} m_r[\text{nil}]$  and  $i?\text{req}.\text{nil} \not\xrightarrow{j?\text{req}}$ .

It turns out that enforcer  $\text{tenf}(m_t, \varphi_1)$ , however. Although this property is not as easy to show – due to the universal quantification over all systems – we can get a fairly good intuition for why this is the case via the example  $s_g$ : it satisfies  $\varphi_1$  and  $m_t[s_g] \sim s_g$  holds.

► **Definition 8 (Enforcement).** A monitor  $m$  enforces property  $\varphi$  whenever it does so (*i*) soundly, Definition 4 and (*ii*) transparently, Definition 6.

For any reasonably expressive logic (such as  $\mu\text{HML}$ ), it is usually the case that *not* every formula can be enforced, as the following example informally illustrates.

$$\varphi, \psi \in \text{sHML} ::= \text{tt} \quad | \quad \text{ff} \quad | \quad \bigwedge_{i \in I} \varphi_i \quad | \quad \llbracket p, c \rrbracket \varphi \quad | \quad X \quad | \quad \max X. \varphi$$

■ **Figure 3** The syntax for the safety  $\mu\text{HML}$  fragment, sHML.

► **Example 9.** Consider the  $\mu\text{HML}$  property  $\varphi_{\text{ns}}$ , together with the two systems  $s_{\text{ra}}$  and  $s_r$ :

$$\varphi_{\text{ns}} \stackrel{\text{def}}{=} [i?\text{req}] \text{ff} \vee [i!\text{ans}] \text{ff} \qquad s_{\text{ra}} \stackrel{\text{def}}{=} i?\text{req}. \text{nil} + i!\text{ans}. \text{nil} \qquad s_r \stackrel{\text{def}}{=} i?\text{req}. \text{nil}$$

A system satisfies  $\varphi_{\text{ns}}$  if *either* it cannot produce action  $i?\text{req}$  *or* it cannot produce action  $i!\text{ans}$ . Clearly,  $s_{\text{ra}}$  violates this property as it can produce both. This system can only be enforced via action suppressions or replacements because insertions would immediately break transparency. Without loss of generality, assume that our monitors employ suppressions (the same argument applies for action replacement). The monitor  $m_r \stackrel{\text{def}}{=} \text{rec } y. (\{i?\text{req}, \tau\}.y + \{i!\text{ans}, \tau\}.y)$  would in fact be able to suppress the offending actions produced by  $s_{\text{ra}}$ , thus obtaining  $m_r[s_{\text{ra}}] \in \llbracket \varphi_{\text{ns}} \rrbracket$ . However, it would also suppress the sole action  $i?\text{req}$  produced by the system  $s_r$ , even though this system satisfies  $\varphi_{\text{ns}}$ . This would, in turn, violate the transparency criterion of Definition 6 since it needlessly suppresses  $s_r$ 's actions, i.e., although  $s_r \in \llbracket \varphi_{\text{ns}} \rrbracket$  we have  $m_r[s_r] \not\approx s_r$ . The intuitive reason for this problem is that a monitor cannot, in principle, look into the computation graph of a system, but is limited to the behaviour the system exhibits at runtime.

## 5 Synthesising Suppression Enforcers

Despite their merits, Definitions 3 and 8 are not easy to work with. The universal quantifications over all systems in Definitions 4 and 6 make it hard to establish that a monitor correctly enforces a property. Moreover, according to Definition 3, in order to determine whether a particular property is enforceable or not, one would need to show the existence of a monitor that correctly enforces it; put differently, showing that a property is *not* enforceable entails another universal quantification, this time showing that no monitor can possibly enforce the property. Lifting the question of enforceability to the level of a (sub)logic entails a further universal quantification, this time on all the logical formulas of the logic; this is often an infinite set. We address these problems in two ways. First, we identify a non-trivial syntactic subset of  $\mu\text{HML}$  that is *guaranteed to be enforceable*; in a multi-pronged approach to system verification, this could act as a guide for whether the property should be considered at a pre-deployment or post-deployment phase. Second, for *every* formula  $\varphi$  in this enforceable subset, we provide an *automated procedure* to *synthesize* a monitor  $m$  from it that correctly enforces  $\varphi$  when instrumented over arbitrary systems, according to Definition 8. This procedure can then be used as a basis for constructing tools that automate property enforcement.

In this paper, we limit our enforceability study to suppression monitors, transducers that are only allowed to intervene by dropping (observable) actions. Despite being more constrained, suppression monitors side-step problems associated with what data to use in a payload-carrying action generated by the enforcer, as in the case of insertion and replacement monitors: the notion of a default value for certain data domains is not always immediate. Moreover, suppression monitors are particularly useful for enforcing *safety* properties, as shown in [33, 10, 20]. Intuitively, a suppression monitor would suppress actions as soon as it becomes apparent that a violation is about to be committed by the SuS. Such an intervention intrinsically relies on the *detection* of a violation. To this effect, we use a prior result from

[25], which identified a maximally-expressive logical fragment of  $\mu\text{HML}$  that can be handled by violation-detecting (recogniser) monitors. We thus limit our enforceability study to this maximal safety fragment, called sHML, since a *transparent* suppression monitor cannot judiciously suppress actions without first detecting a (potential) violation. Figure 3 recalls the syntax for sHML. The logic is restricted to *truth* and *falsehood* ( $\text{tt}$  and  $\text{ff}$ ), conjunctions ( $\bigwedge_{i \in I} \varphi$ ), and necessity modalities ( $\{\{p, c\}\}\varphi$ ), while recursion may only be expressed through greatest fixpoints ( $\max X.\varphi$ ); the semantics follows that of Figure 1.

A standard way how to achieve our aims would be to (i) define a (total) synthesis function  $\langle - \rangle :: \text{sHML} \mapsto \text{TRN}$  from sHML formulas to suppression monitors and (ii) then show that for *any*  $\varphi \in \text{sHML}$ , the synthesised monitor  $\langle \varphi \rangle$  enforces  $\varphi$ . Moreover, we would also require the synthesis function to be compositional, whereby the definition of the enforcer for a composite formula is defined in terms of the enforcers obtained for the constituent subformulas. There are a number of reasons for this requirement. For one, it would simplify our analysis of the produced monitors and allow us to use standard inductive proof techniques to prove properties about the synthesis function, such as the aforementioned criteria (ii). However, a naive approach to such a scheme is bound to fail, as discussed in the next example.

► **Example 10.** Consider a semantically equivalent reformulation of  $\varphi_1$  from Example 1.

$$\varphi_2 \stackrel{\text{def}}{=} \max X.(\{\{d\}\}\text{req}, d \neq j\}\{\{d!\text{ans}, \text{true}\}\}X) \wedge (\{\{d\}\}\text{req}, d \neq j\}\{\{d?\text{req}, \text{true}\}\}\text{ff})$$

At an intuitive level, the suppression monitor that one would expect to obtain for the subformula  $\varphi_2' \stackrel{\text{def}}{=} \{\{d\}\}\text{req}, d \neq j\}\{\{d?\text{req}, \text{true}\}\}\text{ff}$  is  $\{\{d\}\}\text{req}, d \neq j\}.\text{rec } y.\{\{d?\text{req}, \tau\}\}.y$  (i.e., an enforcer that repeatedly drops any  $\text{req}$  inputs following a  $\text{req}$  input on the same port), whereas the monitor obtained for the subformula  $\varphi_2'' \stackrel{\text{def}}{=} \{\{d\}\}\text{req}, d \neq j\}\{\{d!\text{ans}, \text{true}\}\}X$  is  $\{\{d\}\}\text{req}, d \neq j\}.\{d!\text{ans}\}.x$  (assuming some variable mapping from  $X$  to  $x$ ). These monitors would then be combined in the synthesis for  $\max X.\varphi_2'' \wedge \varphi_2'$  as

$$m_{\mathbf{b}} \stackrel{\text{def}}{=} \text{rec } x.(\{\{d\}\}\text{req}, d \neq j\}.\{d!\text{ans}\}.x) + (\{\{d\}\}\text{req}, d \neq j\}.\text{rec } y.\{\{d?\text{req}, \tau\}\}.y)$$

One can easily see that  $m_{\mathbf{b}}$  does *not* behave deterministically, *nor* does it soundly enforce  $\varphi_2$ . For instance, for the violating system  $i?\text{req}.i?\text{req}.\text{nil} \notin \llbracket \varphi_2 \rrbracket (= \llbracket \varphi_1 \rrbracket)$  we can observe the transition sequence  $m_{\mathbf{b}}[i?\text{req}.i?\text{req}.\text{nil}] \xrightarrow{i?\text{req}} \{i!\text{ans}\}.m_{\mathbf{b}}[i?\text{req}.\text{nil}] \xrightarrow{i?\text{req}} \text{id}[\text{nil}]$ .

Instead of complicating our synthesis function to cater for anomalies such as those presented in Example 10 – also making it *less* compositional in the process – we opted for a two stage synthesis procedure. First, we consider a *normalised* subset for sHML formulas which is amenable to a (straightforward) synthesis function definition that is compositional. This also facilitates the proofs for the conditions required by Definition 8 for any synthesised enforcer. Second, we show that every sHML formula can be reformulated in this normalised form without affecting its semantic meaning. We can then show that our two-stage approach is expressive enough to show the enforceability for all of sHML.

► **Definition 11** (sHML normal form). The set of normalised sHML formulas is defined as:

$$\varphi, \psi \in \text{sHML}_{\text{nf}} ::= \text{tt} \quad | \quad \text{ff} \quad | \quad \bigwedge_{i \in I} \{\{p_i, c_i\}\}\varphi_i \quad | \quad X \quad | \quad \max X.\varphi.$$

The above grammar combines necessity operators with conjunctions into one construct  $\bigwedge_{i \in I} \{\{p_i, c_i\}\}\varphi_i$ . Normalised sHML formulas are required to satisfy two further conditions:

1. For every  $\bigwedge_{i \in I} \{\{p_i, c_i\}\}\varphi_i$ , for all  $j, h \in I$  where  $j \neq h$  we have  $\llbracket \{p_j, c_j\}\} \cap \llbracket \{p_h, c_h\}\} = \emptyset$ .
2. For every  $\max X.\varphi$  we have  $X \in \text{fv}(\varphi)$ .

In a (closed) normalised sHML formula, the basic terms  $\text{tt}$  and  $\text{ff}$  can never appear unguarded unless they are at the top level (e.g., we can never have  $\varphi \wedge \text{ff}$  or  $\max X_0. \dots \max X_n. \text{ff}$ ). Moreover, in any conjunction of necessity subformulas,  $\bigwedge_{i \in I} [\{p_i, c_i\}] \varphi_i$ , the necessity guards are *disjoint* and *at most one* necessity guard can satisfy any particular action.

► **Definition 12.** The synthesis function  $\llbracket - \rrbracket : \text{SHML}_{\text{nf}} \mapsto \text{TRN}$  is defined inductively as:

$$\begin{aligned} \llbracket X \rrbracket &\stackrel{\text{def}}{=} x & \llbracket \text{tt} \rrbracket &\stackrel{\text{def}}{=} \llbracket \text{ff} \rrbracket \stackrel{\text{def}}{=} \text{id} & \llbracket \max X. \varphi \rrbracket &\stackrel{\text{def}}{=} \text{rec } x. \llbracket \varphi \rrbracket \\ \llbracket \bigwedge_{i \in I} [\{p_i, c_i\}] \varphi_i \rrbracket &\stackrel{\text{def}}{=} \text{rec } y. \sum_{i \in I} \begin{cases} \{p_i, c_i, \tau\}. y & \text{if } \varphi_i = \text{ff} \\ \{p_i, c_i, \underline{p}_i\}. \llbracket \varphi_i \rrbracket & \text{otherwise} \end{cases} \end{aligned}$$

The synthesis function is compositional. It assumes a bijective mapping between formula variables and monitor recursion variables and converts logical variables  $X$  accordingly, whereas maximal fixpoints,  $\max X. \varphi$ , are converted into the corresponding recursive enforcer. The synthesis also converts truth and falsehood formulas,  $\text{tt}$  and  $\text{ff}$ , into the identity enforcer  $\text{id}$ . Normalized conjunctions,  $\bigwedge_{i \in I} [\{p_i, c_i\}] \varphi_i$ , are synthesised into a *recursive summation* of enforcers, i.e.,  $\text{rec } y. m_i$ , where  $y$  is fresh, and every branch  $m_i$  can be either of the following:

- (i) when  $m_i$  is derived from a branch of the form  $[\{p_i, c_i\}] \varphi_i$  where  $\varphi_i \neq \text{ff}$ , the synthesis produces an enforcer with the *identity transformation* prefix,  $\{p_i, c_i, \underline{p}_i\}$ , followed by the enforcer synthesised from the continuation  $\varphi_i$ , i.e.,  $[\{p_i, c_i\}] \varphi_i$  is synthesised as  $\{p_i, c_i, \underline{p}_i\}. \llbracket \varphi_i \rrbracket$ ;
- (ii) when  $m_i$  is derived from a branch of the form  $[\{p_i, c_i\}] \text{ff}$ , the synthesis produces a *suppression transformation*,  $\{p_i, c_i, \tau\}$ , that drops every concrete action matching the symbolic action  $\{p_i, c_i\}$ , followed by the recursive variable of the branch  $y$ , i.e., a branch of the form  $[\{p_i, c_i\}] \text{ff}$  is translated into  $\{p_i, c_i, \tau\}. y$ .

► **Example 13.** Recall formula  $\varphi_1$  from Example 1, recast in term of  $\text{SHML}_{\text{nf}}$ 's grammar:

$$\varphi_1 \stackrel{\text{def}}{=} \max X. \bigwedge ( [\{(d)?\text{req}, d \neq j\}] ( [\{d!\text{ans}, \text{true}\}] X \wedge [\{d?\text{req}, \text{true}\}] \text{ff} ) )$$

Using the synthesis function defined in Definition 12, we can generate the enforcer

$$\llbracket \varphi_1 \rrbracket = \text{rec } x. \text{rec } z. \sum ( [\{(d)?\text{req}, d \neq j\}]. \text{rec } y. (\{d!\text{ans}, \text{true}\}. x + \{d?\text{req}, \text{true}, \tau\}. y) )$$

which can be optimized by removing redundant recursive constructs (e.g.,  $\text{rec } z. \_$ ), obtaining:

$$= \text{rec } x. \{d?\text{req}, d \neq j\}. \text{rec } y. (\{d!\text{ans}, \text{true}\}. x + \{d?\text{req}, \text{true}, \tau\}. y) = m_{\mathbf{t}}$$

We now present the first main result to the paper.

► **Theorem 14 (Enforcement).** *The (sub)logic  $\text{SHML}_{\text{nf}}$  is enforceable.*

**Proof.** By Definition 3, the result follows if we show that for all  $\varphi \in \text{SHML}_{\text{nf}}$ ,  $\llbracket \varphi \rrbracket$  enforces  $\varphi$ . By Definition 8, this is a corollary following from Propositions 15 and 16 stated below. ◀

► **Proposition 15 (Enforcement Soundness).** *For every system  $s \in \text{SYS}$  and  $\varphi \in \text{SHML}_{\text{nf}}$  then  $\varphi \in \text{SAT}$  implies  $\llbracket \varphi \rrbracket[s] \in \llbracket \varphi \rrbracket$ .*

► **Proposition 16 (Enforcement Transparency).** *For every system  $s \in \text{SYS}$  and  $\varphi \in \text{SHML}_{\text{nf}}$  then  $s \in \llbracket \varphi \rrbracket$  implies  $\llbracket \varphi \rrbracket[s] \sim s$ .*



Following Theorem 14, to show that sHML is an enforceable logic, we only need to show that for every  $\varphi \in \text{sHML}$  there exists a corresponding  $\psi \in \text{sHML}_{\text{nf}}$  with the same semantic meaning, i.e.,  $\llbracket \varphi \rrbracket = \llbracket \psi \rrbracket$ . In fact, we go a step further and provide a constructive proof using a transformation  $\langle\langle - \rangle\rangle : \text{sHML} \mapsto \text{sHML}_{\text{nf}}$  that derives a semantically equivalent sHML<sub>nf</sub> formula from a standard sHML formula. As a result, from an arbitrary sHML formula  $\varphi$  we can then automatically synthesise a correct enforcer using  $\langle\langle \varphi \rangle\rangle$  which is useful for tool construction.

Our transformation  $\langle\langle \varphi \rangle\rangle$  relies on a number of steps; here we provide an outline of these steps. First, we assume sHML formulas that only use symbolic actions with *normalised* patterns  $p$ , i.e., patterns that do not use any data or free data variables (but they may use bound data variables). In fact, any symbolic action  $\{p, c\}$  can be easily converted into a corresponding one using normalised patterns as shown in the next example.

► **Example 17.** Consider the symbolic action  $\{d!ans, d \neq j\}$ . It may be converted to a corresponding normalised symbolic action by replacing every occurrence of a data or free data variable in the pattern by a fresh bound variable, and then add an equality constraint between the fresh variable and the data or data variable it replaces in the pattern condition. In our case, we would obtain  $\{(e)!(f), d \neq j \wedge e=d \wedge f=ans\}$ .

Our algorithm for converting sHML formulas (with normalised patterns) to sHML<sub>nf</sub> formulas,  $\langle\langle - \rangle\rangle$ , is based on Rabinovich’s work [41] for determinising systems of equations which, in turn relies on the standard powerset construction for converting NFAs into DFAs. It consists in the following six stages that we outline below:

1. We unfold each recursive construct in the formula, to push recursive definitions inside the formula body. E.g., the formula  $\max X. (\{p_1, c_1\}X \wedge \{p_2, c_2\}\text{ff})$  is expanded to the formula  $\{p_1, c_1\}(\max X. \{p_1, c_1\}X \wedge \{p_2, c_2\}\text{ff}) \wedge \{p_2, c_2\}\text{ff}$ .
2. The formula is converted into a system of equations. E.g., the expanded formula from the previous stage is converted into the set  $\{X_0 = \{p_1, c_1\}X_0 \wedge \{p_2, c_2\}X_1, X_1 = \text{ff}\}$ .
3. For every equation, the symbolic actions in the right hand side that are of the same kind are alpha-converted so that their bound variables match. E.g., Consider  $X_0 = \{p_1, c_1\}X_0 \wedge \{p_2, c_2\}X_1$  from the previous stage where, for the sake of the example,  $p_1 = (d_1)?(d_2)$  and  $p_2 = (d_3)?(d_4)$ . The patterns in the symbolic actions are made syntactically equivalent by renaming  $d_3$  and  $d_4$  in  $\{p_2, c_2\}$  into  $d_1$  and  $d_2$  respectively.
4. For equations with matching patterns in the symbolic actions, we create a variant that symbolically covers all the (satisfiable) permutations on the symbolic action conditions. E.g., Consider  $X_0 = \{p_1, c_1\}X_0 \wedge \{p_1, c_3\}X_1$  from the previous stage. We expand this to  $X_0 = \{p_1, c_1 \wedge c_3\}X_0 \wedge \{p_1, c_1 \wedge c_3\}X_1 \wedge \{p_1, c_1 \wedge \neg(c_3)\}X_0 \wedge \{p_1, \neg(c_1) \wedge c_3\}X_1$ .
5. For equations with branches having *syntactically equivalent* symbolic actions, we carry out a unification procedure akin to standard powerset constructions. E.g., we convert the equation from the previous step to  $X_{\{0\}} = \{p_1, c_1 \wedge c_3\}X_{\{0,1\}} \wedge \{p_1, c_1 \wedge \neg(c_3)\}X_{\{0\}} \wedge \{p_1, \neg(c_1) \wedge c_3\}X_{\{1\}}$  using the (unified) fresh variables  $X_{\{0\}}, X_{\{1\}}$  and  $X_{\{0,1\}}$ .
6. From the unified set of equations we generate again the sHML formula starting from  $X_{\{0\}}$ . This procedure may generate redundant recursion binders, i.e.,  $\max X. \varphi$  where  $X \notin \text{fv}(\varphi)$ , and we filter these out in a subsequent pass.

We now state the second main result of the paper.

► **Theorem 18 (Normalisation).** *For any  $\varphi \in \text{sHML}$  there exists  $\psi \in \text{sHML}_{\text{nf}}$  s.t.  $\llbracket \varphi \rrbracket = \llbracket \psi \rrbracket$ .*

**Proof.** The witness formula in normal form is  $\langle\langle \varphi \rangle\rangle$ , where we show that each and every stage in the translation procedure preserves semantic equivalence. ◀

## 6 Alternative Transparency Enforcement

Transparency for a property  $\varphi$ , Definition 6, only restricts enforcers from modifying the behaviour of satisfying systems, i.e., when  $s \in \llbracket \varphi \rrbracket$ , but fails to specify any enforcement behaviour for the cases when the SuS violates the property  $s \notin \llbracket \varphi \rrbracket$ . In this section, we consider an alternative transparency requirement for a property  $\varphi$  that incorporates the expected enforcement behaviour for *both* satisfying and violating systems. More concretely, in the case of safety languages such as SHML, a system typically violates a property along a specific set of execution traces; in the case of a satisfying system this set of “violating traces” is *empty*. However, not every behaviour of a violating system would be part of this set of violating traces and, in such cases, the respective enforcer should be required to leave the generated behaviour unaffected.

► **Definition 19** (Violating-Trace Semantics). A logic  $\mathcal{L}$  with an interpretation over systems  $\llbracket - \rrbracket : \mathcal{L} \mapsto \mathcal{P}(\text{SYS})$  has a violating-trace semantics whenever it has a secondary interpretation  $\llbracket - \rrbracket_v : \mathcal{L} \mapsto \mathcal{P}(\text{SYS} \times \text{ACT}^*)$  satisfying the following conditions for all  $\varphi \in \mathcal{L}$ :

1.  $(s, t) \in \llbracket \varphi \rrbracket_v$  implies  $s \notin \llbracket \varphi \rrbracket$  and  $s \xrightarrow{t}$ ,
2.  $s \notin \llbracket \varphi \rrbracket$  implies  $\exists t \cdot (s, t) \in \llbracket \varphi \rrbracket_v$ .

We adapt the work in [26] to give SHML a violating-trace semantics. Intuitively, the judgement  $(s, t) \in \llbracket \varphi \rrbracket_v$  according to Definition 20 below, denotes the fact that  $s$  violates the SHML property  $\varphi$  along trace  $t$ .

► **Definition 20** (Alternative Semantics for SHML [26]). The forcing relation  $\vdash_v \subseteq (\text{SYS} \times \text{ACT}^* \times \text{SHML})$  is the least relation satisfying the following rules:

$$\begin{array}{ll}
 (s, \epsilon, \text{ff}) \in \mathcal{R} & \text{always} \\
 (s, t, \bigwedge_{i \in I} \varphi_i) \in \mathcal{R} & \text{if } \exists j \in I \text{ such that } (s, t, \varphi_j) \in \mathcal{R} \\
 (s, \alpha t, \llbracket [p, c] \varphi \rrbracket) \in \mathcal{R} & \text{if } \text{mtch}(p, \alpha) = \sigma, c\sigma \Downarrow \text{true} \text{ and } s \xrightarrow{\alpha} s' \text{ and } (s', t, \varphi\sigma) \in \mathcal{R} \\
 (s, t, \max X.\varphi) \in \mathcal{R} & \text{if } (s, t, \varphi\{\max X.\varphi/X\}) \in \mathcal{R}.
 \end{array}$$

We write  $s, t \vdash_v \varphi$  (or  $(s, t) \in \llbracket \varphi \rrbracket_v$ ) in lieu of  $(s, t, \varphi) \in \mathcal{R}$ . We say that trace  $t$  is a *violating trace* for  $s$  with respect to  $\varphi$  whenever  $s, t \vdash_v \varphi$ . Dually,  $t$  is a *non-violating trace* for  $\varphi$  whenever there does *not* exist a system  $s$  such that  $s, t \vdash_v \varphi$ .

► **Example 21.** Recall  $\varphi_1, s_{\mathbf{b}}$  from Example 1 where  $\varphi_1 \in \text{SHML}$ , and also  $m_{\mathbf{t}}$  from Example 5 where we argued in Example 13 that  $\llbracket \varphi_1 \rrbracket = m_{\mathbf{t}}$  (modulo cosmetic optimisations). Even though  $s_{\mathbf{b}} \notin \llbracket \varphi_1 \rrbracket$ , not all of its exhibited behaviours constitute violating traces: for instance,  $s_{\mathbf{b}} \xrightarrow{i? \text{req} \cdot i! \text{ans}} s_{\mathbf{b}}$  is not a violating trace according to Definition 20. Correspondingly, we also have  $m_{\mathbf{t}}[s_{\mathbf{b}}] \xrightarrow{i? \text{req} \cdot i! \text{ans}} m_{\mathbf{t}}[s_{\mathbf{b}}]$ .

► **Theorem 22** (Adapted and extended from [26]). *The alternative interpretation  $\llbracket - \rrbracket_v$  of Definition 20 is a violating-trace semantics for SHML (with  $\llbracket - \rrbracket$  from Figure 1) in the sense of Definition 19.*

Equipped with Definition 20 we can define an alternative definition for transparency that concerns itself with preserving exhibited traces that are non-violating. We can then show that the monitor synthesis for SHML of Definition 12 observes non-violating trace transparency.

► **Definition 23** (Non-Violating Trace Transparency). An enforcer  $m$  is *transparent* with respect to the non-violating traces of a formula  $\varphi$ , denoted as  $\text{nvtenf}(m, \varphi)$ , iff for all  $s \in \text{Sys}$  and  $t \in \text{ACT}^*$ , when  $s, t \not\vdash_v \varphi$  then

- $s \xrightarrow{t} s'$  implies  $m[s] \xrightarrow{t} m'[s']$  for some  $m'$ , and
- $m[s] \xrightarrow{t} m'[s']$  implies  $s \xrightarrow{t} s'$ .

► **Proposition 24** (Non-Violating Trace Transparency). For all  $\varphi \in \text{SHML}$ ,  $s \in \text{Sys}$  and  $t \in \text{ACT}^*$ , when  $s, t \not\vdash_v \varphi$  then

- $s \xrightarrow{t} s'$  implies  $(\llbracket \varphi \rrbracket)[s] \xrightarrow{t} m'[s']$ , and
- $(\llbracket \varphi \rrbracket)[s] \xrightarrow{t} m'[s']$  implies  $s \xrightarrow{t} s'$ .

We can thus obtain a new definition for “ $m$  enforces  $\varphi$ ” instead of Definition 8 by requiring sound enforcement, Definition 6, and non-violating trace transparency, Definition 23 (instead of the transparent enforcement of Definition 6). This in turn gives us a new definition for enforceability for a logic, akin to Definition 3. Using Propositions 15 and 24, one can show that SHML is also enforceable with respect to the new definition as well.

## 7 Conclusion

This paper presents a preliminary investigation of the enforceability of properties expressed in a process logic. We have focussed on a highly expressive and standard logic,  $\mu\text{HML}$ , and studied the ability to enforce  $\mu\text{HML}$  properties via a specific kind of monitor that performs suppression-based enforcement. We concluded that SHML, identified in earlier work as a maximally expressive safety fragment of  $\mu\text{HML}$ , is also an enforceable logic. To show this, we first defined enforceability for logics and system descriptions interpreted over labelled transition systems. Although enforceability builds upon soundness and transparency requirements that have been considered in other work, our branching-time framework allowed us to consider novel definitions for these requirements. We also contend that the definitions that we develop for the enforcement framework are fairly modular: e.g., the instrumentation relation is independent of the specific language constructs defining our transducer monitors and it functions as expected as long as the transition semantics of the transducer and the system are in agreement. Based on this notion of enforcement, we devise a two-phase procedure to synthesise correct enforcement monitors. We first identify a syntactic subset of our target logic SHML that affords certain structural properties and permits a compositional definition of the synthesis function. We then show that, by augmenting existing rewriting techniques to our setting, we can convert any SHML formula into this syntactic subset.

## Related Work

In his seminal work [44], Schneider regards a property (in a linear-time setting) to be enforceable if its *violation* can be *detected* by a *truncation automaton*, and prevents its occurrence via system termination; by preventing misbehaviour, these enforcers can only enforce safety properties. Ligatti et al. in [33] extended this work via *edit automata* – an enforcement mechanism capable of *suppressing* and *inserting* system actions. A property is thus enforceable if it can be expressed as an edit automaton that *transforms* invalid executions into valid ones via suppressions and insertions. Edit automata are capable of enforcing instances of safety and liveness properties, along with other properties such as infinite renewal properties [33, 10]. As a means to assess the correctness of these automata, the authors introduced *soundness* and *transparency*. In both of these settings, there is no

clear separation between the specification and the enforcement mechanism, and properties are encoded in terms of the languages accepted by the enforcement model itself, i.e., as edit/truncation automata. By contrast, we keep the specification and verification aspects of the logic separate.

Bielova et al. [10, 11] remark that soundness and transparency do not specify to what extent a transducer should modify an invalid execution. They thus introduce a *predictability* criterion to prevent transducers from transforming invalid executions arbitrarily. More concretely, a transducer is *predictable* if one can predict the number of transformations that it will apply in order to transform an invalid execution into a valid one, thereby preventing enforcers from applying unnecessary transformations over an invalid execution. Using this notion, Bielova et al. thus devise a more stringent notion of enforceability. Although we do not explore this avenue, Definition 23 may be viewed as an attempt to constrain transformations of violating systems in a branching-time setup, and should be complementary to these predictability requirements.

Könighofer et al. in [29] present a synthesis algorithm that produces action replacement transducers called *shields* from safety properties encoded as automata-based specifications. Shields analyse the inputs and outputs of a reactive systems and enforce properties by modifying the least amount of output actions whenever the system deviates from the specified behaviour. By definition, shields should adhere to two desired properties, namely correctness and minimum deviation which are, in some sense, analogous to soundness and transparency respectively. Falcone et al. in [19, 21, 20], also propose synthesis procedures to translate properties – expressed as Streett automata – into the resp., enforcers. The authors show that most of the property classes defined within the *Safety-Progress hierarchy* [35] are enforceable, as they can be encoded as Streett automata and subsequently converted into enforcement automata. As opposed to Ligatti et al., both Könighofer et al. and Falcone et al. separate the specification of the property from the enforcement mechanism, but unlike our work they do not study the enforceability of a branching time logic.

To the best of our knowledge, the only other work that tackles enforceability for the modal  $\mu$ -calculus [30] (a reformulation of  $\mu$ HML) is that of Martinelli et al. in [37, 38]. Their approach is, however, different from ours. In addition to the  $\mu$ -calculus formula to enforce, their synthesis function also takes a “witness” system satisfying the formula as a parameter. This witness system is then used as the behaviour that is mimicked by the instrumentation via suppression, insertion or replacement mechanisms. Although the authors do not explore automated correctness criteria such as the ones we study in this work, it would be interesting to explore the applicability of our methods to their setting.

Bocchi et al. [12] adopt *multi-party session types* to project the global protocol specifications of distributed networks to *local types* defining a local protocol for every process in the network that are then either verified statically via typechecking or enforced dynamically via suppression monitors. To implement this enforcement strategy, the authors define a dynamic monitoring semantics for the local types that suppress process interactions so as to conform to the assigned local specification. They prove local soundness and transparency for monitored processes that, in turn, imply global soundness and transparency by construction. Their local enforcement is closely related to the suppression enforcement studied in our work with the following key differences: (i) well-formed branches in a session type are, by construction, *explicitly disjoint* via the use of distinct choice labels (i.e., similar to our normalised subset sHML<sub>nf</sub>), whereas we can synthesise enforcers for *every* sHML formula using a normalisation procedure; (ii) they give an LTS semantics to their local specifications (which are session types) which allows them to state that a process satisfies a specification when its behaviour is bisimilar to the operational semantics of the local specification – we do

not change the semantics of our formulas, which is left in its original denotational form; (iii) they do not provide transparency guarantees for processes that violate a specification, along the lines of Definition 23; (iv) Our monitor descriptions sit at a lower level of abstraction than theirs using a dedicated language, whereas theirs have a session-type syntax with an LTS semantics (e.g., repeated suppressions have to be encoded in our case using the recursion construct while this is handled by their high-level instrumentation semantics).

In [14], Castellani et al. adopt session types to define reading and writing privileges amongst processes in a network as global types for information flow purposes. These global types are projected into local monitors capable of preventing read and write violations by adapting certain aspects of the network. Although their work is pitched towards adaptation [24, 13], rather than enforcement, in certain instances they adapt the network by suppressing messages or by replacing messages with messages carrying a default nonce value. It would be worthwhile investigating whether our monitor correctness criteria could be adapted or extended to this information-flow setting.

### Future Work

We plan to extend this work along two different avenues. On the one hand, we will attempt to extend the enforceable fragment of  $\mu\text{HML}$ . For a start, we intend to investigate maximality results for suppression monitors, along the lines of [25, 2]. We also plan to consider more expressive enforcement mechanisms such as insertion and replacement actions. Finally, we will also investigate more elaborate instrumentation setups, such as the ones explored in [1], that can reveal refusals in addition to the actions performed by the system.

On the other hand, we also plan to study the implementability and feasibility of our framework. We will consider target languages for our monitor descriptions that are closer to an actual implementation (e.g., an actor-based language along the lines of [26]). We could then employ refinement analysis techniques and use our existing monitor descriptions as the abstract specifications that are refined by the concrete monitor descriptions. The more concrete synthesis can then be used for the construction of tools that are more amenable towards showing correctness guarantees.

---

### References

- 1 Luca Aceto, Antonis Achilleos, Adrian Francalanza, and Anna Ingólfssdóttir. A framework for parameterized monitorability. In *Foundations of Software Science and Computation Structures*, pages 203–220, Cham, 2018. Springer International Publishing.
- 2 Luca Aceto, Antonis Achilleos, Adrian Francalanza, and Anna Ingólfssdóttir. Monitoring for silent actions. In Satya Lokam and R. Ramanujam, editors, *FSTTCS 2017: Foundations of Software Technology and Theoretical Computer Science*, volume 93 of *LIPICs*, pages 7:1–7:14, Dagstuhl, Germany, 2018. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.
- 3 Luca Aceto, Anna Ingólfssdóttir, Kim Guldstrand Larsen, and Jiri Srba. *Reactive Systems: Modelling, Specification and Verification*. Cambridge University Press, New York, NY, USA, 2007.
- 4 Rajeev Alur and Pavol Černý. Streaming transducers for algorithmic verification of single-pass list-processing programs. In *Proceedings of the 38th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 599–610. ACM, 2011.
- 5 Henrik Reif Andersen. Partial model checking. In *Proceedings of Tenth Annual IEEE Symposium on Logic in Computer Science*, pages 398–407. IEEE, 1995.
- 6 Cyrille Artho, Howard Barringer, Allen Goldberg, Klaus Havelund, Sarfraz Khurshid, Michael R. Lowry, Corina S. Pasareanu, Grigore Rosu, Koushik Sen, Willem Visser, and

- Richard Washington. Combining test case generation and runtime verification. *Theoretical Computer Science*, 336(2-3):209–234, 2005.
- 7 Duncan Paul Attard, Ian Cassar, Adrian Francalanza, Luca Aceto, and Anna Ingólfssdóttir. *A Runtime Monitoring Tool for Actor-Based Systems.*, chapter 3, pages 49–74. River Publishers, 2017.
  - 8 Duncan Paul Attard and Adrian Francalanza. A monitoring tool for a branching-time logic. In *Runtime Verification*, pages 473–481, Cham, 2016. Springer International Publishing.
  - 9 Jean Berstel and Luc Boasson. Transductions and context-free languages. *Ed. Teubner*, pages 1–278, 1979.
  - 10 Nataliia Bielova. *A theory of constructive and predictable runtime enforcement mechanisms.* PhD thesis, University of Trento, 2011.
  - 11 Nataliia Bielova and Fabio Massacci. Predictability of enforcement. In *International Symposium on Engineering Secure Software and Systems*, pages 73–86. Springer, 2011.
  - 12 Laura Bocchi, Tzu-Chun Chen, Romain Demangeon, Kohei Honda, and Nobuko Yoshida. Monitoring networks through multiparty session types. *Theoretical Computer Science*, 669:33–58, 2017.
  - 13 Ian Cassar and Adrian Francalanza. On implementing a monitor-oriented programming framework for actor systems. In *International Conference on Integrated Formal Methods*, pages 176–192. Springer, 2016.
  - 14 Ilaria Castellani, Mariangiola Dezani-Ciancaglini, and Jorge A. Pérez. Self-adaptation and secure information flow in multiparty communications. *Formal Aspects of Computing*, 28(4):669–696, July 2016.
  - 15 Edward Chang, Zohar Manna, and Amir Pnueli. The safety-progress classification. In *Logic and Algebra of Specification*, pages 143–202. Springer, 1993.
  - 16 Clare Cini and Adrian Francalanza. An LTL proof system for runtime verification. In *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, pages 581–595. Springer, 2015.
  - 17 Edmund M Clarke and E Allen Emerson. Design and synthesis of synchronization skeletons using branching time temporal logic. In *25 Years of Model Checking*, pages 196–215. Springer, 2008.
  - 18 Ankush Desai, Tommaso Dreossi, and Sanjit A. Seshia. Combining model checking and runtime verification for safe robotics. In *Runtime Verification (RV)*, LNCS, pages 172–189, Cham, 2017. Springer International Publishing.
  - 19 Yliès Falcone. You should better enforce than verify. In *Runtime Verification*, pages 89–105. Springer Berlin Heidelberg, 2010.
  - 20 Yliès Falcone, Jean-Claude Fernandez, and Laurent Mounier. What can you verify and enforce at runtime? *International Journal on Software Tools for Technology Transfer*, 14(3):349, jun 2012.
  - 21 Yliès Falcone, Laurent Mounier, Jean-Claude Fernandez, and Jean-Luc Richier. Runtime enforcement monitors: composition, synthesis, and enforcement abilities. *Formal Methods in System Design*, 38(3):223–262, jun 2011.
  - 22 Adrian Francalanza. A Theory of Monitors. In *International Conference on Foundations of Software Science and Computation Structures*, pages 145–161. Springer, 2016.
  - 23 Adrian Francalanza. Consistently-Detecting Monitors. In *28th International Conference on Concurrency Theory (CONCUR 2017)*, volume 85 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 8:1–8:19, Dagstuhl, Germany, 2017. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.
  - 24 Adrian Francalanza, Luca Aceto, Antonis Achilleos, Duncan Paul Attard, Ian Cassar, Dario Della Monica, and Anna Ingólfssdóttir. A foundation for runtime monitoring. In *Runtime Verification*, pages 8–29, Cham, 2017. Springer International Publishing.

- 25 Adrian Francalanza, Luca Aceto, and Anna Ingólfssdóttir. Monitorability for the Hennessy-Milner logic with recursion. *Formal Methods in System Design*, 51(1):87–116, 2017.
- 26 Adrian Francalanza and Aldrin Seychell. Synthesising correct concurrent runtime monitors. *Formal Methods in System Design*, 46(3):226–261, 2015.
- 27 Limin Jia, Hannah Gommerstadt, and Frank Pfenning. Monitors and blame assignment for higher-order session types. In *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 582–594, New York, NY, USA, 2016. ACM.
- 28 Katarína Kejstová, Petr Ročkal, and Jiří Barnat. From Model Checking to Runtime Verification and Back. In *RV*. Springer, 2017.
- 29 Bettina Könighofer, Mohammed Alshiekh, Roderick Bloem, Laura Humphrey, Robert Könighofer, Ufuk Topcu, and Chao Wang. Shield synthesis. *Formal Methods in System Design*, 51(2):332–361, Nov 2017.
- 30 Dexter C. Kozen. Results on the propositional  $\mu$ -calculus. *Theoretical Computer Science*, 27:333–354, 1983.
- 31 Frédéric Lang and Radu Mateescu. Partial model checking using networks of labelled transition systems and boolean equation systems. In Cormac Flanagan and Barbara König, editors, *TACAS*, pages 141–156, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg.
- 32 Kim G Larsen. Proof systems for satisfiability in Hennessy-Milner logic with recursion. *Theoretical Computer Science*, 72(2):265–288, 1990.
- 33 Jay Ligatti, Lujo Bauer, and David Walker. Edit automata: enforcement mechanisms for run-time security policies. *International Journal of Information Security*, 4(1):2–16, Feb 2005.
- 34 Jay Ligatti and Srikar Reddy. A theory of runtime enforcement, with results. In *CESORICS*, pages 87–100, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg.
- 35 Zohar Manna and Amir Pnueli. A hierarchy of temporal properties. In Cynthia Dwork, editor, *Proceedings of the Ninth Annual ACM Symposium on Principles of Distributed Computing*, pages 377–410. ACM, 1990. doi:10.1145/93385.93442.
- 36 Fabio Martinelli and Ilaria Matteucci. Partial model checking, process algebra operators and satisfiability procedures for (automatically) enforcing security properties. In *Foundations of Computer Security*, pages 133–144. Citeseer, 2005.
- 37 Fabio Martinelli and Ilaria Matteucci. Through modeling to synthesis of security automata. *Electronic Notes in Theoretical Computer Science*, 179:31–46, 2006.
- 38 Fabio Martinelli and Ilaria Matteucci. An approach for the specification, verification and synthesis of secure systems. *Electronic Notes in Theoretical Computer Science*, 168:29–43, 2007.
- 39 Robin Milner, Joachim Parrow, and David Walker. A calculus of mobile processes, I. *Information and computation*, 100(1):1–40, 1992.
- 40 Amir Pnueli and Aleksandr Zaks. PSL model checking and run-time verification via testers. In Jayadev Misra, Tobias Nipkow, and Emil Sekerinski, editors, *International Symposium on Formal Methods*, pages 573–586. Springer Berlin Heidelberg, 2006.
- 41 Alexander Moshe Rabinovich. A complete axiomatisation for trace congruence of finite state behaviors. In *Proceedings of the 9th International Conference on Mathematical Foundations of Programming Semantics*, pages 530–543, London, UK, UK, 1994. Springer-Verlag.
- 42 Jacques Sakarovitch. *Elements of Automata Theory*. Cambridge University Press, New York, NY, USA, 2009.
- 43 Davide Sangiorgi. *Introduction to Bisimulation and Coinduction*. Cambridge University Press, New York, NY, USA, 2011.
- 44 Fred B Schneider. Enforceable security policies. *ACM Transactions on Information and System Security (TISSEC)*, 3(1):30–50, 2000.