Preface to the Special Issue on engineering reliable multi-agent systems

Jürgen Dix¹ · Brian Logan² · Michael Winikoff³

Accepted: 23 June 2021 / Published online: 8 July 2021
© Springer Science+Business Media, LLC, part of Springer Nature 2021

There is increasing interest in the application of autonomous systems technology in areas such as driverless cars, UAVs, manufacturing, healthcare, and personal assistants. Indeed, Robotics and Autonomous Systems have been identified as one of the *Eight Great Technologies* [3] with the potential to revolutionise our economy and society. For example, it has been claimed that the “economic, cultural, environmental and social impacts and benefits [of autonomous systems] will be unprecedented” [2].

However, widespread deployment and the consequent benefits of autonomous systems will only be achieved if they can be shown to operate *reliably*. Reliable operation is essential for public and regulatory acceptance, as well as the myriad of societal changes necessary for their widespread deployment (e.g., liability insurance).

Autonomous systems can be viewed as a particular kind of (multi)agent system, where the focus is on the problem of achieving flexible intelligent behaviour in dynamic and unpredictable environments. Demonstrating that such a system will operate reliably is an extremely challenging problem. The potential ‘behaviour space’ of many systems (e.g., robots for care of the elderly) is vastly larger than that addressed by current approaches to engineering reliable systems [4, 5]. Multiagent/autonomous systems are implicitly expected to be able to ‘do the right thing’ in the face of conflicting objectives and in complex, ill-structured environments. Addressing these challenges cannot be achieved by incremental improvements to existing software engineering and verification methodologies, but will require step changes in how we specify, engineer, test and verify systems.

Achieving reliability poses two challenges: how to define what is a “good decision”, and how to check that software will make such decisions in all situations. The latter

✉ Jürgen Dix
dix@tu-clausthal.de

✉ Brian Logan
b.s.logan@uu.nl

✉ Michael Winikoff
michael.winikoff@vuw.ac.nz

¹ Clausthal University of Technology, Julius-Albert Strasse 4, 38678 Clausthal, Germany

² School of Computer Science, University of Nottingham, Wollaton Road, Nottingham NG8 1BB, UK

³ School of Information Management, Victoria University of Wellington, PO Box 600, Wellington 6140, New Zealand

challenge is the focus of formal verification. However, achieving reliability is not just about formal methods. Especially for multiagent/autonomous systems that are likely to be situated in a complex socio-technical context, defining what is a “good decision” is highly challenging to say the least. Factors such as social structures, ethical principles, and legal requirements may have to be taken into account.

This special issue is one of the outcomes from Dagstuhl seminar 19112 “Engineering Reliable Multiagent Systems”, which was held in March 2019 (<http://dagstuhl.de/19112>, [1]). The Dagstuhl workshop attracted 26 leading international researchers from a range of fields including theoretical computer science, engineering multiagent systems, machine learning, and ethics in artificial intelligence.

The call for papers for this special issue was circulated later in 2019, with the deadline for submissions being the start of 2020. All submissions were reviewed following the normal process and quality expectations of JAAMAS. The pandemic struck us hard in 2020 and resulted in an overly lengthy reviewing process at the end of which only two papers were accepted.

The first article, by Michael Fisher, Viviana Mascardi, Kristin Yvonne Rozier, Bernd-Holger Schlingloff, Michael Winikoff, and Neil Yorke-Smith is titled “Towards a Framework for Certification of Reliable Autonomous Systems”. This paper considers the challenge of how regulators should deal with autonomous systems. For instance, how might an Unmanned Aerial System be certified for use in civilian airspace? The paper reviews relevant standards, highlights issues, surveys the state of the art in verification of autonomous systems, and proposes a reference three layer autonomy framework, and a process for identifying requirements for certification. Finally, challenges to researchers, to engineers, and to regulators are articulated.

The second article, by Davide Calvaresi, Yashin Dicient Cid, Mauro Marinoni, Aldo Franco Dragoni, Amro Najjar, and Michael Ignaz Schumacher is titled “Real-Time Multi-Agent Systems: Formal Model and Empirical Results”. This paper considers the challenge of getting agent systems to operate in environments that require real-time responses: not only reasoning about time but *in time*. The authors suggest a formal mathematical model of Real-Time Multi-Agent Systems which they also implemented and evaluated on a number of tasks, using a multi-agent system simulator for general-purpose and real-time constraints. One of the observations is that when employing multi-agent systems in scenarios with strict timing constraints, one not only needs to adopt real-time theories and scheduling models, but also suitable protocols (e.g., RBN protocol), and a communication middle-ware with a bounded time delay.

We hope that this special issue takes a step forward in establishing a new research agenda for engineering reliable autonomous systems: clarifying the problem, its properties, and their implications for solutions.

The Dagstuhl workshop that led to this special issue was in March 2019. On the last day of the workshop (Friday 15th March) we awoke to news of the terrorist attack in Christchurch, New Zealand. The Dagstuhl seminar 19112 participants, like any academic gathering, was diverse. There were many religions, ethnicities, and nationalities. Inclusiveness and tolerance are key values of our community, and we condemn in the strongest possible terms not just the attack itself, but racism, extremism, intolerance, white supremacy (and all forms of supremacy), and the misguided beliefs that led to the Christchurch attack, and to other attacks.

References

1. Dix, J., Logan, B., & Winikoff, M. (2019). Engineering reliable multiagent systems (Dagstuhl Seminar 19112). *Dagstuhl Reports*, 9(3), 52–63.
2. Patchett, C. (2014). *Robotics and autonomous systems: Challenges and opportunities for the UK*. <https://royalsociety.org/~media/events/2015/11/robotics/Robotics>.
3. Willetts, D. (2013). *Eight great technologies*. Policy Exchange. <https://policyexchange.org.uk/wp-content/uploads/2016/09/eight-great-technologies.pdf>.
4. Winikoff, M., & Cranefield, S. (2014). On the testability of BDI agent systems. *Journal of Artificial Intelligence Research (JAIR)*, 51, 71–131. <https://doi.org/10.1613/jair.4458>.
5. Winikoff, M. (2017). BDI agent testability revisited. *Journal of Autonomous Agents and Multi-Agent Systems (JAAMAS)*, 31(5), 1094–1132. <https://doi.org/10.1007/s10458-016-9356-2>.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.