An abstract formal basis for digital crowds

Marija Slavkovik · Louise A. Dennis ·
Michael Fisher

Published online: 13 November 2014

© The Author(s) 2014. This article is published with open access at Springerlink.com

Abstract Crowdsourcing, together with its related approaches, has become very popular in recent years. All crowdsourcing processes involve the participation of a digital crowd, a large number of people that access a single Internet platform or shared service. In this paper we explore the possibility of applying formal methods, typically used for the verification of software and hardware systems, in analysing the behavior of a digital crowd. More precisely, we provide a formal description language for specifying digital crowds. We represent digital crowds in which the agents do not directly communicate with each other. We further show how this specification can provide the basis for sophisticated formal methods, in particular formal verification.

Keywords Logical foundations · Crowd specification · Formal verification · Predictability

1 Introduction

Crowdsourcing is an umbrella term used to describe a wide range of activities coordinated within an Internet-based platform. It's original definition is by Howe:

Simply defined, crowdsourcing represents the act of a company or institution taking a function once performed by employees and outsourcing it to an undefined (and generally large) network of people in the form of an open call. This can

M. Slavkovik (✉)

Department of Information Science and Media Studies, University of Bergen, Bergen, Norway
e-mail: marija.slavkovik@infomedia.uib.no

L. A. Dennis · M. Fisher

Department of Computer Science, University of Liverpool, Liverpool, UK
e-mail: l.a.dennis@liverpool.ac.uk

M. Fisher

e-mail: mfisher@liverpool.ac.uk

take the form of peer-production (when the job is performed collaboratively), but is also often undertaken by sole individuals. The crucial prerequisite is the use of the open call format and the large network of potential laborers. [22]

In the most general sense, therefore, crowdsourcing can be defined as the business practice of a “seeker” entity, attempting to out-source activities, such as solving problems, performing tasks or raising funds, to a priori unidentified undertakers, namely the “crowd”. Crowdsourcing has a wide range of potential applications and is the basis for many successful businesses; as such it has also emerged as a new and vibrant research area [53].

The ubiquitous component of all the crowdsourcing processes is the *crowd* or, more specifically, the *digital crowd*. In English, the term “crowd” is used to refer to a large number of people that are in each other’s immediate presence [27]. To distinguish our terminology from this we use “digital crowd” to refer to a large number of agents that access a single Internet platform or shared service, i.e., a large number of agents that are in each other’s immediate e-presence. The agents can represent individual people, organisations, other web-based services, etc.

Crowdsourcing is, foremost, a business practice and the methods of managing and analysing it are “practical” rather than “formal”. To the best of our knowledge, there are no formal methods to verify various requirements, for example, to determine when crowdsourcing is an adequate approach to a problem, under which conditions crowdsourcing will be successful, how effective crowdsourcing can be on a particular problem, whether a digital crowd can be expected to solve a problem successfully, etc. Our aim is to explore whether formal methods can be applied to improving the understanding and effectiveness of digital crowds.

1.1 Varieties of digital crowds

Originally covered by the umbrella term, “crowdsourcing”, a range of applications involving digital crowds have been developed, incorporating many different agent structures and different interaction protocols. Some initial attempts to develop a taxonomy for the growing number of crowdsourcing varieties, have been made [12, 39].

We consider some, but by no means all, crowdsourcing practices, from the viewpoint of the interaction between the seeker and the digital crowd and the communication among the digital crowd participants. We would like to emphasise that because of the lack of a comprehensive common taxonomy, and since the terms for different processes partially overlap among different taxonomists, the names used for each process here may be considered imprecise by some or inadequate by others. We follow to some extent the crowdsourcing taxonomy developed in [12] as a recent approach that encompasses and builds upon earlier work.

Crowdcasting is possibly the earliest and best understood example of a crowdsourcing process. Here, a seeker advertises an open challenge to the crowd and offers a financial reward for one solution, e.g., the best or the first. Seekers can be people, businesses, or representatives of larger institutions, but each is represented by, and acts as, one agent. The challenges can range from small creative tasks such as the design of a logo, to

grand challenges for solving complicated technical problems. The crowd that responds to crowdcasting is usually comprised of people that have the adequate skills to solve the challenge, though no certification is required that participants possess such skills. People can solve the tasks individually, or work in teams. However, it is not the aim of the platform to facilitate the finding of team-mates.

The *99designs* (99designs.co.uk) system is a platform that enables seekers to advertise their design needs and proposed rewards. The platform users wishing to respond to a seeker do so by sending in their design proposals. The seeker chooses the preferred design solution and rewards the proponent of that solution. The respondents are typically individuals and all communication is done via the platform. In contrast, *Innocentive* (www.innocentive.com) is a platform where difficult technical challenges are posed, for which a solution may not even exist. Here, the respondents are typically teams of people who already know each other, with varied skill sets, that share the prize if awarded.

Crowdcomputing is a general term we use to name the many processes whereby a seeker defines an end goal that is to be accomplished by the crowd participants. Each of the participants contributes a fractional effort towards accomplishing that goal. One of these processes is microtasking, a process where a vast number of very simple tasks that take little time and effort to solve are created by the seeker. The task solvers are almost exclusively individuals, and the rewards are small. The rewards for the crowd might be monetary or an intrinsic motivation maybe induced in the participants to inspire them to contribute, e.g., altruism or entertainment. As in crowdcasting, the platforms do not offer any means for communication among members of the crowd. However, in contrast to crowdcasting one task is given to one solver, and thus all solutions are paid for. To ensure that the solutions are of a satisfactory standard, the solvers must be certified before they can be assigned tasks.

Representative examples of microtasking crowdcomputing platforms with financial rewards are Amazon's *Mechanical Turk* (www.mturk.com), and CrowdFlower (crowdflower.com). An example of a crowdcomputing process where the rewards are entertainment rather than monetary are when they are organised as a *game with a purpose*. Here, the game is hosted by the platform, the participants solve the tasks by playing the game and their reward is the pleasure derived from playing. In this type of crowdcomputing, pre-certification is not necessary. A typical example is the DNA puzzle game for multi-sequence alignment Phylo developed by the McGill Centre for Bioinformatics (phylo.cs.mcgill.ca).

Common to all these processes is the fact that although the crowd participants solve a problem jointly, they do not actually collaborate, are not aware of, nor depend on, the solutions of the other crowd participants.

Crowdcontent, as described by [12] includes the processes in which a content is created, analysed or discovered by a digital crowd. The term “Wisdom of the Crowds” or “Collective Intelligence”, can also be considered to include these processes. The abilities, information and motivation of the digital crowd's participants are combined, resulting in the emergence of new abilities, information and intentionality. While to some extent this is true with crowdsourcing processes such as Mechanical Turk,

there is a notable difference; in Mechanical Turk the tasks, called “human intelligence tasks” require no particular cognitive effort to perform and all tasks requested by each individual seeker are similar. With collective intelligence, the tasks are not given to the crowd, but the members contribute cognitively more demanding efforts in response to what they perceive is necessary to complete the larger goal. Furthermore, while in the crowdcomputing processes the participants work independently, here they consider and build upon the work of others, though direct communication among the participants is not strictly necessary. In [12] the crowdcontent processes are further differentiated into crowd production, crowd searching, and crowdanalysing, the terms being self-explanatory.

While collective intelligence is a phenomenon that occurs spontaneously in some systems, platforms can be designed specifically with the task of promoting such collective efforts. In this case, the platform is the *seeker*, looking for information and abilities, as well as assigning tasks to the digital crowd. Possibly the best example of this crowdsourcing activity is *Wikipedia* (www.wikipedia.org).¹ Typically no material payment is associated with this subtype of crowdcomputing process, the reward being advancement of a common good.

Crowdvoting is a crowdsourcing process in which a platform only promotes or sells those goods and services that are supported by the majority of the digital crowd. In another variant, crowdvoting is a method of eliciting an opinion only from the community that is interested in a particular issue. An example of a crowdvoting platform is *Threadless* (beta.threadless.com). Threadless is a platform on which the seekers are *designers* who can submit garment illustration designs, while the crowd assigns approval to the designs that they like. Designs that achieve a certain number of votes are printed and offered for sale on the platform, with profits being transferred to the seekers.

An example that conceptually lies between crowdvoting and crowdcomputing is a collaborative effort such as writing a scenario for a movie together, as instigated by Paul Verhoeven² or drafting alterations to a constitution, as is the case with Iceland.³

The processes by which the crowd chooses an item, such as in *Threadless* (beta.threadless.com) is termed *crowdopinion* in [12], while the processes by which the crowd changes the good they are creating until there are no objections, as is the case with writing a scenario or a constitution, is termed *crowd storming*. What these processes have in common is that unlike the others mentioned earlier, crowd participants can appear to interact with each other directly, and even collaborate directly towards the creation of the end good. However, even in these processes, the communication is actually executed via the hosting platform, namely the participants do not send private messages to each other and all communication is accessible to the whole digital crowd.

¹ Though the lack of unanimity from the crowd makes Wikipedia an atypical example. We are grateful to the anonymous reviewer who pointed this out.

² <http://www.bbc.com/culture/story/20130807-the-public-cant-write>.

³ <http://www.wired.co.uk/news/archive/2012-10/23/iceland-crowdsourced-constitution>.

Crowdfunding is different from other crowdsourcing processes, and one of the least ambiguous. In crowdfunding the seeker does not out-source a task, but advertises a fund-raising project. The digital crowd does not compete, but participates, as they are donors that pledge their own funds and resources towards the accomplishment of the project. A seeker may directly advertise a project or it may do so through a moderator. As in crowdcasting, seekers and donors can be individuals, teams, or even companies, but crowd members are not expected to communicate with each other or collaborate. The seekers can look to raise funds to finance their business venture, micro-loans or charity offers for non-profit organisations.

A crowdfunding platform that matches seekers looking to finance a business ventures is *Kickstarter* (www.kickstarter.com). The donors' funds are given with no return on investment up to a certain value or with some minor goods in return over a certain value. *Kiva*⁴ (www.kiva.org) is a platform that manages micro-loans to private individuals. More precisely, Kiva first establishes contact with moderators, which are traditional loan-givers and repayment enforcers, and a seeker who is requesting a loan from the moderators. The stories and needs of the seekers are passed on to Kiva by the moderators. These stories are advertised by Kiva to the donors who then transfer money to Kiva, which distributes it to the moderators. The money is loaned to the seekers and, when returned, is repaid to the donors. Kiva's task is to manage the moderators to ensure that funds end up in the hands of the seekers. Platforms that handle traditional fund-raising projects are *Global Giving* (www.globalgiving.org) for charity and *Indiegogo* (www.indiegogo.com) for both charity and personal projects.

Smart Mobs represent a hybrid between a digital and a physical crowd [38]. Smart mobs are not a crowdsourcing process, however we do consider them here because they involve a digital crowd. Smart mobs do not utilise a crowdsourcing platform exclusively for their purpose, as these are one-off activities. A generic platform, such as an Internet forum, is used to coordinate the activities of an actual crowd, which as a result is formed deliberately, is purposeful and is efficient. A seeker, typically an instigator, advertises a cause and gives the specifications of the activity he or she wants to occur. Those from the digital crowd that wish to participate respond by following the activity specification and, as a result, an actual physical crowd may form.

A concept related to smart mobs is *Internet vigilantism*. As in a smart mob, a seeker posts a cause and a description of activity to a generic platform while responders participate by complying. A seeker may act as an individual in a single instance of vigilantism or as a complex entity that frequently takes on vigilantism activities e.g., *Anonymous* [30]. Seekers in the context of Internet vigilantism engage in a wider range of activities that may or may not result in physical crowd formation. For example, a seeker may describe a perceived injustice and solicit the help of the digital crowd to seek out the perpetrators and expose their identities to the authorities or to the public. Seekers may also advertise a smart mob activity but with the purpose of vigilantism in the cyber space, such as a coordinated *distributed denial-of-service attack* to an Internet service.

⁴ Although Kiva was created before crowdfunding became popular.

Common to all digital crowds we consider here is that the participants in the crowd do not communicate directly and privately with each other. In the rare instance when there is a need to communicate with another crowd participant, the participant directs this communication to the crowdsourcing platform that hosts the crowd, and it is the platform that relays the message. In the rest of this work we focus on an abstract digital crowd, not specific for any one crowdsourcing process, in which the communication between crowd and seeker, as well as among crowd participants, is executed via a mutual crowd hosting platform.

1.2 Formal methods

Formal methods represent a collection of techniques and tools for the specification, development and verification of software and hardware systems. These techniques are typically based on mathematical structures and formal logic and so provide clear and unambiguous descriptions of required behaviour. This logical basis then provides the opportunity for sophisticated analysis tools to be utilised, for example based on logical proof or exhaustive state-space exploration. Verification is the process of establishing that a designed system has its intended properties. Formal verification has become widespread, enabling deep and (semi) automated formal analysis of systems and so providing greater clarity concerning reliability and correctness. The behaviour of complex systems, even in a restricted environment, can easily become unpredictable making verifiability a very desirable quality in such systems.

Our long-term aim is to use an automated formal verification technique called *model-checking* [4] to analyse digital crowd systems. However, in order to achieve this a strong and logically precise logical basis for crowd systems must be developed; this is the issue we address here. Model-checking has previously been applied to the analysis of *multi agent systems* [8, 34] so it is a plausible technology to target for the analysis of digital crowds. A Multi Agent System (MAS) is an artificial system containing (possibly among other entities) a set of agents and managing the communication between them. Although a digital crowd is a system of agents, it cannot be considered to be straight-forward MAS in the traditional sense of this paradigm (we elaborate on this in Sect. 2).

1.3 Multi-agent systems (MAS)

The definition of what constitutes an ‘agent’ varies among disciplines. In computer science and artificial intelligence, an agent is typically defined as

an encapsulated computer system that is situated in some environment, and that is capable of flexible, autonomous action in that environment in order to meet its design objectives [24].

Further, an agent is usually considered to be responsive with respect to other agents and its environment, pro-active in its activities, and able to interact with other agents where necessary. Analysis and construction of MASs is central to issues such as problem solving, coordination, cooperation and control with, and of, multiple agents.

The agent paradigm, and the ensuing methods for modelling and representing agents and systems, are geared towards supporting these research aims.

The *specification* of a MAS means building a representation of the agents, and possibly the enclosing environment, that constitute the system. A common route to specifying more sophisticated agents involves attributing mental states to them. The well known *Belief–Desire–Intention* (BDI) model [37] uses:

- *beliefs* to represent the agent’s (possibly incomplete and incorrect) information about itself, other agents, and its environment;
- *desires* to represent the agent’s long-term aims; and
- *intentions* to capture aims that are actively being pursued.

The agent then *deliberates* over what desires to target, how to tackle them and what actions to take. The BDI model has been very influential and has consequently been altered and extended with a range of other mental states such as actions, obligations, abilities etc, and various logics, such as epistemic, doxastic, deontic and temporal logics, have been developed and successfully applied to represent these states.

Unlike typical MASs, digital crowds are not *deliberate* systems. In a MAS each agent is an individual that is fully identified by her mental state. Illustratively, the agent is the entity of interest and we specify the environment and the other agents, including the MAS, as seen through her “mental eye”. Even when a MAS of several agents is specified, all agents in the system are specified individually one by one by representing their mental states.

In a digital crowd, no part of the participants’ mental states may be made available to other agents as part of the process in which they are involved. Therefore we cannot represent the crowd participants one by one, as none of their mental states may be known.

What is known of the digital crowd participants is the messages they exchange within the crowd. It is the information exchanged between the agents that are the focus of our interest, because this is where interesting phenomena can emerge. We propose that such exchanges are what need to be specified in order to specify a digital crowd.

Agent-oriented software engineering [23,51] is a research area specifically concerned with developing methods for engineering applications conceptualised as MASs [28]. Within agent-oriented software engineering, logic-based formal methods are used to specify systems, program systems, and verify systems [9]. The agent specification languages that drive current agent model-checking systems are inadequate for specifying the structural complexity of digital crowds (see our discussion in Sect. 2). In this paper, we develop an improved logic-based agent specification language that overcomes the shortcomings. Communication among agents in a MAS is considered to be an act and implemented as such, see e.g., [5,29,50]. Here we are not concerned with how the agents exchange messages, but with the content of such messages as a special, public part of the mental state of the agent in a social setting.

The paper is structured as follows. To distinguish the requirements for an improved agent specification language for digital crowds, we first discuss in more detail the difference between the “typical” MAS and digital crowds and the shortcomings of

existing agent logics-based languages and logics in Sect. 2. In Sect. 3 we propose an extension to the BDI agent paradigm, and in Sect. 4 we propose an agent specification language. In Sect. 5 we present examples of verification routes for crowdsourcing using our specification. In Sect. 6 we discuss conclusions and future work.

2 Problem analysis

Within crowdsourcing platforms, a digital crowd is typically considered as a single entity (not an agent structure) whose intentions, beliefs or goals cannot be known, but whose incentives can be foreseen. So, while being fully aware that there could be vast numbers of agents accessing the platform, and that the content of the communication is visible to all these agents, the *seekers* communicate with the digital crowd as they would with one particular agent. This one agent is an embodiment of the crowd, not its leader, but rather a representative agent; a typical agent in the crowd that embodies all the properties that the seeker finds desirable in the agent(s) he or she would ideally task with their problem.

Another common facet of crowdsourcing is that agents within the digital crowds are not expected to interact with each other privately, or even directly communicate with each other (some platforms that host crowdsourcing processes do not even provide the means for such interactions). Nevertheless, the participants are encouraged to communicate with other agents outside the digital crowd by sharing information involving the crowdsourcing process, and this type of communication is facilitated by the platform. Consider for example, the “share on social network” and “email to a friend” buttons embedded in many such platforms.

Although an agent may not be allowed to join a digital crowd if it lacks certain skills or characteristics, within the digital crowd no informational or motivational attitudes are shared. Even in the case of smart mobs, the motivations and beliefs driving each agent to participate in the mob may differ. If enough agents happen to have similar states of mind, a mob will occur. Excluding collaboration and the sharing of mental states from the agent structure makes a digital crowd different from a “traditional” MAS.

Horling et al. [21] give an extensive overview of multi-agent structures such as coalitions, teams, societies etc. We provide a brief overview to compare these structures with digital crowds. We first distinguish between a *shared attitude* and a *we-attitude*. In the literature the we-attitude is called a *joint attitude*, but we here use we-attitudes to distinguish better from the shared attitudes. The difference is subtle: a shared attitude is an attitude that conditions participation in a structure, e.g., if your goal is to achieve φ , then you are part of the agent structure. Illustratively, since your colleague John wants to play football, he is part of the university football team. In contrast, a we-attitude is an attitude that characterises participation in a structure, e.g., we (as an organisation) have the goal to achieve φ , so if you are part of the structure then your goal is φ as well. Illustratively, we are married therefore we-intend to abstain from intimate relations with other agents.

A *coalition* is an agent structure in which each agent is committed to pursuing a we-goal. An agent *institution*, or society, is a structure in which members have certain we-beliefs and we-obligations of compliance with certain norms. This does

not mean that if a norm is violated the entire institution is subject to punitive measures, rather that if an agent is a member of the institution, then he or she is subject to the institutional norms and sanctions. Collaboration is not essential, either in coalitions or in institutions. A *team* is a structure of agents that collaborate, by assuming different roles, towards accomplishing a shared goal.

Digital crowds, and crowds in general, exhibit a dual nature: both that of a single agent and a structure of agents. On one hand, the whole crowd can be viewed as one rational agent. The crowd can be given a task, can be considered accountable for transgressions e.g., in the case of *smart mobs* [38], where the crowd exhibits behaviour emergence and abilities not pertinent to any one participant. On the other hand, the participants in the crowd are individual rational agents, in the sense that they are fully autonomous, undertaking only those tasks they choose to [53].

Comparing how agent structures are specified, we can observe that in coalitions, teams, even institutions, the agents either share mental attitudes or adopt we-attitudes. Consequently, it is possible to represent the agent structure bottom up, by representing all the constituent agents' shared and/or we-attitudes. However, neither crowds nor crowd participants can be modelled in this manner since, in a digital crowd, agents neither share attitudes nor adopt we-attitudes. Therefore, we cannot obtain a digital crowd specification using the bottom up approach of specifying each individual agent of the crowd. Furthermore, specifying each crowd agent is both unnecessary and difficult, as the number of agents in the crowd is potentially vast and constantly changing.

To specify a crowd, we must encapsulate (hide) the mental states of the agent in the representation. This allows us to represent agent structures, such as digital crowds, as single agents within larger structures, such as platforms. We can represent each agent with the attitudes it decides to reveal to its environment or the structures with the attitudes that all its members have in common. By introducing encapsulation, we effectively construct structures top down: first the agent structure is represented, then agents are distinguished by adding more detailed information about their mental states.

Agent specification languages do not represent communication. Collaboration in agent structures is formally represented as a we-mental attitude. Eventually the act, not the content, of communicating is represented as an action. How can we specify an agent in an agent structure without specifying his/her mental attitudes?

Consider an imageboard, an Internet forum where members can post images available to be viewed by other members or visitors without restriction. The forum is a context that contains a crowd of unspecified agents that pop in and out of existence. This forum context exhibits the characteristics of an agent and, clearly, the members and visitors are also agents. How can we represent this system of agents without peering into the mental states of the agents involved? Mental states aim to capture the behavior of an agent. Publicly accessible messages passed between an agent and a structure, e.g., crowdsourcing platform, can be used to capture the behavior of a complex system of agents. Such messages can be either sent *to* the agent or sent *by* the agent. When accessing a forum either as an administrator, visitor or member, the messages posted on the forum are what can be seen and are what is used to make deductions about the system as a whole and individual agents within it. We need to be

able to express not only the fact that a message has occurred but also what the content of the message was, so that we are able to represent the information that agents wish to disclose to their environment, to give incentives to digital crowds, and allow agents to reason with them.

A clear distinction between agent specification and the field of agent communication is needed. Agent communication is a vast area of research within engineering multi agent systems, with considerable work having taken place on developing communication protocols [5,29,32,50]. Agent communication languages [43], such as the FIPA-ACL proposed by FIPA (<http://www.fipa.org>), and increasingly captured as OWL ontologies [17], are languages designed to be used by agents to communicate among themselves within a MAS.

Within the research in agent communications, the communication between agents is represented as an act, in the sense that for example, an autonomous vacuum cleaner is capable of vacuuming dust and communicating its location to its docking station. While consideration is given to the allowed content of the messages exchanged, the agent does not internalise, i.e., represent as a mental state, the fact that a message happened. This is typical of all acts when MAS are implemented. For example, the fact that dirt is collected by the vacuum cleaner is not internalised after the cleaner vacuums. The agent can detect a change in his environment (its dirt container is now non empty) as an effect of having successfully executed the action of vacuuming, but no information state is added that represents the fact that vacuuming has taken place. The autonomous vacuum cleaner will receive a message with content “stop” from the owner, and as a result of that message it will stop its actions, but the fact that a message with content stop has been received is not part of the informational states of the cleaner.

We are here not concerned with communication protocols or the implementation of the communication acts among agents. Instead, we are exclusively concerned with representing, as information states, the information that a message has been sent, or received, and the content of that message. So we represent the intermediary step, the addition *to* the information state of the agent of the fact that communication happened. This happens *after* the communication *act* was executed.

We have to further make clear the difference between a specification and communication language. A specification language is designed to describe the overall behaviour of a system, such as e.g., a multi agent system. We specify what is said to, and by, an agent, but this may not be the language that the agent actually uses to execute communication. Consider an illustrative example. An example of agents communicating using English is when John says to Mary: “I will go to play football.” An example of English used as a specification language is when we as the designers of this example describe the John-Mary interaction: John has informed Mary of his intention to engage in the football play activity. English, as logic, can be used for both communication and specification, but we are here concerned only about the specification aspect. For the challenges regarding agent communication languages in social settings, which do have an overlap with the challenges for specification languages for systems of agents, one can see for example [43].

Fig. 1 Structure of an extended agent

3 From agent to agent structure

A *rational* agent takes action in order to achieve its goals based on the beliefs it holds about its environment and other agents [52]. A rational agent is typically described by representing its dynamic, informational and motivation aspects, i.e., its mental attitudes. This often conforms to the BDI model [36,37]; recall that “BDI” denotes *Beliefs*, *Desires*, and *Intentions*. There are *many* different agent programming languages and agent platforms based, at least in part, on the BDI approach. Particular languages developed for programming *rational* agents in a BDI-like way include AgentSpeak [35], Jason [3], 3APL [19], Jadex [33], GOAL [6], SAAPL [48], and GWENDOLEN [7].

In Fig. 1 we present a diagram of an extended agent that will also be used to specify agent structures. The mental attitudes of the agent are encapsulated in the *private* section of the agent, while the *public* section contains the information that has been sent to the agent, and the information sent by the agent to others. In addition to the public and private sections, there is reference to the other agents (specified in the same manner) that are contained in (respectively, contain) this agent. So, if a single agent is being specified, the list of agents contained within it is empty.

We can see, from Fig. 1 that there is a connection between the private and public sections. This is relevant once the agent decides to internalize some of the information it receives into its mental states, or when it wants to share some aspect of its mental states with its environment.

3.1 Encapsulation

An agent is typically considered to be an autonomous software or hardware entity capable of perceiving its environment, pursuing its own goals and interacting with the environment and with other agents. Initially, it might appear that the agent paradigm

Fig. 2 Representing a team using the agent structure [15]

is not appropriate for modelling digital crowd participants, since it applies to software and hardware entities, while digital crowds are comprised almost exclusively of people. However, people do not directly participate in the digital crowd—their participation is almost exclusively by medium of a user account, effectively their software avatar. We can build a representation of the avatar and extend it with a representation of the expected behaviour of the agent.

The formal description of flexible agent structures was developed in [16] in the context of executable agent specifications. There, to the formal specification of the agent's mental attitudes, two additional sets are added: *content* and *context*. These correspond to the two sets given in Fig. 1, and allow groups, teams or organizations of agents to be dually treated as individual agents. For example, as the environment is described as a set of agents within the context then, if a message is to be sent into the environment, it is just broadcast to the context agents. Crucially, there is no distinction between an agent that contains others and one that does not; effectively, organisations such as crowds can be themselves viewed as agents.

In [15, 18] this agent specification approach is extended further to include an agent's specification that is visible, or accessible, to the agent's content or context respectively. Effectively this means that one part of the agent is visible to the agent's content and another part, possibly the same, is visible to the agent's context. This extension allows for the “typical” agent structures, in which the agents have joint or shared attitudes, to be represented. We give a simple example of a team using this representation in Fig. 2. The agent *team* has, in its content, all the agents T_1 , T_2 , T_3 . The team agent has we-attitudes visible to its content, in this example it is the intention to “move the sofa”, but nothing in its private specification. The agents T_i have their own shared attitudes, in this example the belief that the agent in question can “lift” (the sofa), visible to the context and in the private specification, the rest of their mental attitudes.

The *context/content* extension enables an agent structure to be directly represented but, regardless of whether the agent specification is private or public, it is still in a typical BDI specification. What all BDI-oriented agent specification languages have in common is that they express what the agent “thinks” and how it “reasons”. Although the agent might choose to reveal some of its mental attitudes, this is still internal information. The information that the agent wants to share with the environment and other

agents might not correspond to, or even be consistent with, its formal specification. To model a digital crowd participant we also need to model what the agent “hears” and “says”, potentially hiding what the agent “thinks” from other agents. Thus, to distinguish what the agent “was told” and “has said” from its own private “thought” behaviour we extend the above approach and explicitly represent messages that are sent to (and received from) the agent’s content or context. By having this information explicitly represented, the agent can reason about it and react to it accordingly.

3.2 Communication as an informational state

A formal specification of an agent is a logic formula that describes its behaviour, including the information it has about the world, the motivations that drive it, and the methods that define how it reacts in response to messages. Logics developed for representing different aspects of agent behaviour include modal logics of belief (B), goals (G), wishes (W), desires (D), intentions (I), actions ($\llbracket A \rrbracket$), abilities (A), knowledge (K), etc, all with an underlying temporal and/or probabilistic basis [14, 40]. Following this tradition, we introduce a new modal logic operator “ M ” to handle message passing, representing messages as $M\varphi$, where φ describes the message content. We here consider that the message content is a well formed formula for the language used to represent the mental states of the agent.

There are two basic qualifiers that must be attributed to a message to distinguish between messages that have been received and messages that have been sent. Additionally, it is also necessary to identify the sender or recipient, respectively. To accomplish this distinction we modify the M operator using a superscript $\uparrow j$ denoting messages sent to agent j and $\downarrow j$ denoting messages received from agent j . Thus, given an agent communication specification, including $M^{\uparrow j}\varphi$ in the specification requires that the agent has sent a message with content φ to agent j . Similarly, $M^{\downarrow j}\varphi$ represents that the specified agent has received a message with content φ from agent j . However, the modifiers $\uparrow j$ and $\downarrow j$ are insufficient to express complex messages such as incentives.

We further modify the M operator with a *type*, denoted as a subscript e.g., $M_{type}^{\uparrow j}\varphi$, clarifying the intended meaning of the message. We distinguish between four types.

tell denotes information passing, implying the sender informs the receiver of what its mental state is; for example, $M_{tell}^{\uparrow j}I\varphi$ is a message with which the specified agent informs agent j of its intention to actively pursue φ (i.e., it has φ as an intention).

ask denotes a request for confirmation of state; for example, $M_{ask}^{\downarrow j}I\varphi$ represents the fact that the agent has received an inquiry from j who wants to know whether the agent intends to pursue φ .

do denotes delegation, extension and transfer of attitudes; for example, $M_{do}^{\downarrow j}I\varphi$ is a message in which j delegates to the recipient the active pursuit of φ , while $M_{do}^{\uparrow j}B\varphi$ is a message in which the sender extends the requirement of adopting the belief φ to agent j , etc. (Note that, to send or receive a message of delegation does not mean that the recipient is obliged to adopt the content.)

Table 1 Intuitive interpretation of the messages in the specification of agent i

<i>tell</i>	
$M_{tell}^{\downarrow j} \varphi$ i is told that φ holds for j	$M_{tell}^{\uparrow j} \varphi$ j is told that φ holds for i
<i>ask</i>	
$M_{ask}^{\downarrow j} \varphi$ j asks: does φ holds for i ?	$M_{ask}^{\uparrow j} \varphi$ i asks: does φ holds for j ?
<i>do</i>	
$M_{do}^{\downarrow j} \varphi$ j says φ should hold for i	$M_{do}^{\uparrow j} \varphi$ i says φ should hold for j
<i>adv</i>	
$M_{adv}^{\downarrow j} \varphi$ φ should hold for i , if i were in the content of j	$M_{adv}^{\uparrow j} \varphi$ φ should hold for j , if j were in the content of i

Agent j is an agent different from agent i

adv denotes announcement, promise and advertisement, and allows agents to inform others of incentives or constraints and other mental attitudes that are in force for prospective members; for example, $M_{adv}^{\uparrow j} I \varphi$ denotes that if the recipient j joins the content of the sender, then i will be asked to adopt the intention φ , while $M_{adv}^{\downarrow j} I \varphi$ denotes that if j is added to the content of the recipient, then j will adopt the intention φ .

Table 1 gives the intuitive interpretation of such communication types in the specification of agent i .

Notation In the rest of the text, we use the notation $M_{[\cdot]}^{\uparrow}$ and $M_{[\cdot]}^{\downarrow}$, respectively to denote messages of any type, and we use the $\Downarrow$ symbol when the orientation of the message (incoming or outgoing) is irrelevant.

The above message type modifiers are inspired by the illocutionary (speech) act types of [41]. Searle distinguishes among assertives, directives, commissives, expressives and declarations. We can draw correspondences between our message types and Searle's illocutionary types: expressives and messages of type *tell*, assertives and messages of type *ask*, directives and messages of type *do*, and commissives and messages of type *adv*. We do not have a message type that corresponds to declarations, which are speech acts that permanently change the world, such as pronouncing someone as being married, employed etc, because we do not consider these types necessary for representing digital crowds at present.

Within MAS research, agent languages for the coordination of agents, based on speech acts, have been developed, e.g., [1,47]. However, communicating using these languages, the agents use speech acts, with vocabulary and interpretation of the received messages depending on the specific language used. As discussed in Sect. 2, our messages are not speech acts, rather, they are the agent's attitudes and the illocutionary act types are used to modify the intended meaning of the communication, the reason why the message content is sent, or how the message content is to be interpreted.

Logics have been developed that are concerned with exchange of knowledge among agents, probably the most notable being *dynamic epistemic logic* [46]. It is important to outline the difference between our M operator and the communication represented in this logic. Within dynamic epistemic logic, one can specify information that is communicated to all the agents, called an *epistemic update*, as well as knowledge that all the agents possess, i.e., common knowledge. Once an update φ occurs, the agent that receives it necessarily changes her mental states to accommodate φ . As a result, from the moment of the update onwards the agent believes that φ holds. However, when an agent receives a message $M\varphi$, she has a choice of whether to accommodate φ in her mental states or not. In other words, after a message $M_{tell}^{\downarrow i} B\varphi$ or even after a message $M_{do}^{\downarrow i} B\varphi$ is received, it is not necessary that the agent will have $B\varphi$ as part of her specification. Explicitly, we do not presume that the agent tells the truth in the content of a “tell” message. Thus, $M_{tell}^{\uparrow i} B\varphi$ and $B\neg\varphi$ are mutually consistent.

4 Logically specifying agents

Let Agt be a set of unique agent identifiers, let $Prop$ be a set of atomic propositions and constants, and $Pred$ be a set of a first-order predicates of arbitrary arity. We begin by defining a language $\mathcal{L}_p$ to be a set of grounded first order logic formulas without function symbols, namely the set of all φ_p such that

$$\varphi_p ::= p \mid \neg\varphi_p \mid \varphi_p \wedge \varphi_p \mid P(x_1, \dots, x_m)$$

where $p \in Prop$, $P \in Pred$ and $x_1, \dots, x_m \in Agt$.

We next consider a *BDI* agent language. Depending on the specific needs for a specification, different *BDI* operators can be used but, for demonstrating our specification approach, we use the modal operators B , G and I , to denote agent's beliefs, long term goals and actively pursued goals (or intentions), respectively. We also use an operator, A , to denote that an agent has an ability, $A\varphi$ indicating that the agent is able to accomplish φ . Ability is a more complex mental attitude and several formalizations are possible, e.g., “ways” [2,10,45,49], depending on the precise interpretation of ability. To avoid an in-depth philosophical analysis of the logics of abilities, which is outside of the scope of this work, and to simplify our modelling language by avoiding the use of actions as logical primitives, we resolve to use $A\varphi$ for denoting both procedural knowledge of which action sequence brings about φ and the actual ability to perform an action sequence that brings about φ . The language $\mathcal{L}_{BDI}$ is then the set of

all formulas φ such that

$$\varphi ::= \varphi_p \mid \neg\varphi \mid \varphi \wedge \varphi \mid B\varphi_p \mid I\varphi_p \mid G\varphi_p \mid A\varphi_p,$$

where $\varphi_p \in \mathcal{L}_p$.

Finally, we define the language for specifying communication among agents, $\mathcal{L}_M$. For this language, temporal logic operators should be specified depending on the needs of the particular system specified. We use *LTL* operators in our examples [14]. Let T be the set of message types $T = \{tell, ask, do, adv\}$ and Agt be a set of unique agent identifiers. The language $\mathcal{L}_M$ is the set of all formulas θ such that

$$\theta ::= \varphi \mid \neg\theta \mid \theta \wedge \theta \mid \varphi U \varphi \mid \bigcirc\varphi \mid \Diamond\varphi \mid M_x^{\downarrow j} \theta \mid M_x^{\uparrow j} \theta,$$

where $\varphi \in \mathcal{L}_{BDI}$ and $x \in T$. In the intuitive interpretation of temporal operators: pUq means that p is continuously true up until the point when q becomes true; $\bigcirc r$ means that r is true in the next moment in time; while $\Diamond s$ means that s will be true at some moment in the future.

The messages are sent to an agent j , however either the *context* set CX or the *content* set CN as a whole can be the target of message broadcast. We use the shorthand⁵

$$M_{[\cdot]}^{\uparrow CN} \varphi \equiv \bigwedge_{j \in CN} M_{[\cdot]}^{\uparrow j} \varphi, \quad M_{[\cdot]}^{\uparrow CX} \varphi \equiv \bigwedge_{j \in CX} M_{[\cdot]}^{\uparrow j} \varphi.$$

The language $\mathcal{L}_{BDI}$ restricts the nesting of modal operators, while $\mathcal{L}_M$ forbids the use of *BDI* and temporal operators outside of the scope of a message operator. Thus the agents do not have mental attitudes about the future, e.g., $B\Diamond\varphi$ “I believe that sometimes φ is true” nor $BG\varphi$ “I believe I have the goal to accomplish φ ”.

Nested messages express meta communication, allowing agents to communicate about what was communicated to them or by them. However not all nesting is meaningful. We state constraints as shown in (1) in order to apply restrictions mandating that all but the innermost and outermost message operators are dropped, the orientation of the inner most and outermost messages is retained, as is the type of the innermost operator, while the type of the outermost operator must be *tell*.

$$\begin{aligned} M_{[\cdot]}^{\downarrow i} M_{[\cdot]}^{\uparrow i} \dots M_{[\cdot]}^{\uparrow i} M_{[x]}^{\downarrow i} \varphi &\leftrightarrow M_{[tell]}^{\downarrow i} M_{[x]}^{\downarrow i} \varphi \\ M_{[\cdot]}^{\downarrow i} M_{[\cdot]}^{\uparrow i} \dots M_{[\cdot]}^{\uparrow i} M_{[x]}^{\uparrow i} \varphi &\leftrightarrow M_{[tell]}^{\uparrow i} M_{[x]}^{\uparrow i} \varphi \\ M_{[\cdot]}^{\uparrow i} M_{[\cdot]}^{\uparrow i} \dots M_{[\cdot]}^{\uparrow i} M_{[x]}^{\downarrow i} \varphi &\leftrightarrow M_{[tell]}^{\uparrow i} M_{[x]}^{\downarrow i} \varphi \\ M_{[\cdot]}^{\uparrow i} M_{[\cdot]}^{\uparrow i} \dots M_{[\cdot]}^{\uparrow i} M_{[x]}^{\uparrow i} \varphi &\leftrightarrow M_{[tell]}^{\uparrow i} M_{[x]}^{\uparrow i} \varphi \end{aligned} \tag{1}$$

⁵ We define the messages with individual agents, not sets as in [15,16,18], because a message can be broadcast to many agents, but it can be sent from one agent, otherwise the sender is unknown, which cannot happen here—if your contexts sends you a message it is from exactly one context.

We assume that (some of) the agent abilities are transferable and thus messages such as $M_{do}^{\downarrow i} A\varphi$ and $M_{do}^{\uparrow i} A\varphi$ represent the transfer of abilities from agent i to the sender and vice versa, respectively. Messages such as $M_{adv}^{\uparrow i} A\varphi$ represent transfer of abilities after the agent i has joined the content of the sender.

We can now give the following definition of an agent.

Definition 1 Let Agt be a set of unique agent identifiers. An agent is a tuple $\langle ID, Bel, Int, Goal, Ablt, Com, CN, CX \rangle$, where $ID \in Agt$ is a unique agent identifier, $Bel \subset \mathcal{L}_p$ is the set of beliefs the agent holds about the world, $Int \subset \mathcal{L}_p$ is the set of the agent's intentions, $Goal \subset \mathcal{L}_p$ is the set of the agent's goals, $Ablt \subset \mathcal{L}_p$ is the set of the agent's abilities, $Com \subset \mathcal{L}_M$ is the set of messages the agent has received and sent, $CN \subset \mathcal{P}(Agt \setminus \{ID\})$ is the set of agents contained and lastly $CX \subset \mathcal{P}(Agt \setminus \{ID\})$ is the set of agents in which the agent is contained, i.e., its set of contexts. Each of the sets Bel , Int , $Goal$ and $Ablt$ are consistent and simplified.

Given an agent $i \in Agt$, an agent specification is a set $SPEC(i) \subset \mathcal{L}_M$, where $B\varphi$ is true iff $\varphi \in Bel$, $G\varphi$ is true iff $\varphi \in Goal$, $I\varphi$ is true iff $\varphi \in Int$, $A\varphi$ is true iff $\varphi \in Ablt$, $cn(j)$ is true iff $j \in CN$, $cx(j)$ is true iff $j \in CX$ and $M_{[\cdot]}^{\uparrow i} \varphi$ is true iff $M_{[\cdot]}^{\uparrow i} \varphi \in Com$.

Note that, to be able to reason about contents and contexts, we introduce special formulas in $SPEC(i)$, namely $cn(j)$ and $cx(j)$. The formula $G\ cx(j)$ within $SPEC(i)$ expresses i 's goal to have j as a context agent, and similarly $G\ cn(j)$ within $SPEC(i)$ denotes the goal to include j in his content. Symmetrically $G\ \neg cx(j)$ and $G\ \neg cn(j)$ express the goal to remove j from ones own context, or content, respectively.

Lastly, we assume, via (2), that if a message is sent then it will eventually be received. This is a property of communication among agents that should hold in the environment, for communication to be meaningful.

$$\exists i, M_{[\cdot]}^{\uparrow j} \varphi \in SPEC(i) \Rightarrow \exists j, \Diamond M_{[\cdot]}^{\downarrow i} \varphi \in SPEC(j) \quad (2)$$

Note that we do not develop an axiomatisation for $\mathcal{L}_M$ and do not intend to prove soundness for this language, because we intend ultimately to use it to create specifications for model checking, where soundness is not necessary. The above, together with standard modal and temporal logic semantic structures [44], provides a formal basis for describing digital crowd structures, communication and, hence, behaviour.

5 Specification and verification of digital crowds

We demonstrate the applicability of the specification notation above by considering two digital crowd examples, and describing the verification processes that can be deployed.

Throughout, we assume that verification tools for individual agents, such as those available to analyse the *BDI* properties of Java-based agents [8], can be utilised to assess individual agents. We here show, using examples, how the specification language we have developed is appropriate for describing and reasoning about the

crowd behaviour, and that properties we might wish to establish of the crowd can be built up from the properties of individual agents. Note also that, while the proofs we outline are provided by us, we would expect automated (or, at least, semi-automated) provers to be able to generate these in a straight-forward way.

We begin by discussing some common aspects of digital crowds.

In a (digital) crowd, the agents are highly autonomous with goals, and attitudes in general, beyond a seeker's control. As a consequence, agents cannot be given a goal, they can only be inspired, or be provided with incentives, to adopt a goal themselves. In economics, an *incentive* is considered to be a cost or benefit that motivates decisions or actions of agents. As observed in the multitude of work concerned with crowdsourcing, e.g., [11,20,25,39], one of the key features of the crowdsourcing process is the provision of incentives, by the seeker, in the form of benefits to the agents that (successfully) participate in the crowdsourcing process. We can specify, as an advertisement, a conditional future possibility to benefit with the schema (3).

$$[INC]: M_{adv}^{\uparrow i}(conditions \rightarrow \Diamond reward). \quad (3)$$

The incentive (3) is a message of type “promise”. The premise *conditions* of the message is a formula describing the required skills, the pursuit of required goals, etc., that could lead to benefit in the content of the sender. The *reward* is the benefit that an agent can obtain, such as money, goods, recognition, status, etc. In our examples we use the formula “*A earn*” as a general representation for an earned reward.

In the same manner as specifying positive incentives, those that promise gain, we can also specify negative incentives that promise a penalty. In this case, the conditions specify what the agent should refrain from, e.g., intentions not to be upheld, and instead of $\Diamond reward$ as a consequent we would have $\Diamond penalty$. While penalties are of interest in agent structures such as institutions, in crowdsourcing, penalties are not commonly utilized; therefore we investigate them no further here.

We cannot, nor do we want to, access the mental states of each of the agents of the crowd, but we can look at the exchanged messages and, based on assumptions about how an agent interprets those messages, we can establish certain properties of a crowd.

5.1 Information retrieval

We next consider an example in which the retrieval of a particular item of information is *crowdsourced*, such as finding a suspect's name (in our case, *Nemo*) from an image in a criminal investigation. We want to formally ascertain that Nemo will be found under the assumption that there exists an agent in the crowd who knows where Nemo is, as long as some simple assumptions for the relations among the agents in the crowd hold as well.

Note that, making assumptions about a system and then verifying that under these assumptions certain properties hold is not unusual in formal verification. While the exact state of a system cannot always be known, assuming that the system is in the required state, we can ensure that a desired state can be reached. This allows us to be certain that, when operational conditions allow, correct behaviour will ensue; it

also allows us to explore assumptions concerning failure in any of these operational conditions—if the pre-requisites we expect are not actually present, then what can the behaviour of the system be?

In order to consider communication among crowds, let us define the concept of *reachability* between two agents i and j . The agent i can reach agent j if, and only if, a message sent from i is eventually forwarded to j , under the assumption that the relevant context relay messages from one of their content agents to the rest of the content. We first define *relaying contexts*. Intuitively, a relaying context is an agent which broadcasts to all its content agents all messages received from one of his content agents.

Definition 2 Let i be an agent s.t. $CN(i) \neq \emptyset$. Agent $k \in CX(i)$ is a *relaying context*, and $REL(k)$ is true, when all the messages sent to k are sent on to all of the content agents of k :

$$((CN(i) \vee CX(i)) \wedge M_{tell}^{\downarrow i} \varphi) \rightarrow M_{tell}^{\uparrow CN} \varphi \in SPEC(k)$$

Clearly, there are messages that an agent sends to a content that are not to be shared, and should be kept private, but these can be specifically designated as such. The reachability between two agents is now defined recursively.

Definition 3 Agent j is *directly reachable* for agent i if at least one of the following is true:

- $\exists k \in CX(i) \cap CX(j)$ s.t. $REL(k)$ holds.
- $\exists k_1 \in CX(i)$ and $\exists k_2 \in CX(j)$ s.t. $CN(k_1) \cap CN(k_2) \neq \emptyset$ and $REL(k_1) \wedge REL(k_2)$ holds.

Agent j is *reachable* for agent i if at least one of the following is true:

- j is *directly reachable* for agent i .
- $\exists k \in Agt$ s.t. k is reachable for i and j is reachable for k .

In Fig. 3, we give an example of directly reachable (Fig. 3a) and reachable (Fig. 3b) agents.

The set of all agents that are reachable for i is called the neighbourhood of i , $NGH(i)$. Determining whether $j \in NGH(i)$, under the assumption that all context agents are relaying contexts, can be solved as an ST-connectivity problem [31].

There are always certain assumptions that have to be made concerning the behaviour of members of a digital crowd. We will make the assumptions that the following formulas are included in the *SPEC* of every agent in *Agt*:

- An agent is always pursuing the goal to earn (increase her utility):

$$G \text{ earn.} \quad (4)$$

- An agent always shares information with its content and context about any possibilities to earn, namely incentives are forwarded:

$$M_{adv}^{\downarrow j} INC \rightarrow M_{tell}^{\uparrow CN} M_{adv}^{\downarrow j} INC. \quad (5)$$

$$M_{adv}^{\downarrow j} INC \rightarrow M_{tell}^{\uparrow CX} M_{adv}^{\downarrow j} INC. \quad (6)$$

- A forwarded incentive is treated the same as directly received incentive:

$$M_{tell}^{\downarrow k} M_{adv}^{\downarrow j} INC \rightarrow M_{adv}^{\downarrow j} INC. \quad (7)$$

- An agent always responds to an incentive whose conditions it can satisfy:

$$\begin{aligned} & (conditions \wedge Greward \wedge M_{adv}^{\downarrow j} (conditions \rightarrow \Diamond reward)) \\ & \rightarrow M_{tell}^{\uparrow j} conditions. \end{aligned} \quad (8)$$

- An agent is required to have the abilities it claims to have, hence it will tell someone that it has a certain ability only if it really has it:

$$M_{tell}^{\uparrow j} A\varphi \rightarrow A\varphi. \quad (9)$$

- An agent that actively pursues a goal, and has the ability to accomplish its goal, will eventually (believe that she will) have accomplished his goal:

$$(A\varphi \wedge I\varphi) \rightarrow \Diamond B\varphi. \quad (10)$$

Assumptions about the behaviour of the crowd can be based on a statistical analysis of the agents that access a platform of interest, or by only allowing agents with desirable properties to access the platform. Thus, a condition that all contexts are relaying messages can be established by ensuring a platform in which all communication among the members of a structure is visible to all others (consider as an example a *Facebook wall*). It is important to note that, if we move to a different infrastructure, we might modify or even remove some of these assumptions. The key point is that the specification formalism is appropriate for describing these.

In addition to the general assumptions (4)–(10) we make the following specific assumptions for *SPEC(seeker)* within this particular information retrieval scenario.

Fig. 3 Examples of direct reachability and reachability

- If an incentive was advertised that required the ability $find_Nemo$ and an agent tells us that it has such an ability, then delegate the task of $find_Nemo$ to that agent:

$$(M_{adv}^{\uparrow CX} (A \text{ find_Nemo} \rightarrow \Diamond A \text{ earn}) \wedge M_{tell}^{\downarrow k} A \text{ find_Nemo}) \rightarrow M_{do}^{\uparrow k} I \text{ find_Nemo}. \quad (11)$$

- If a task to $find_Nemo$ was delegated to an agent k , and a message was received from k that Nemo is found, then the belief that Nemo is found is adopted:

$$(M_{do}^{\uparrow k} I \text{ find_Nemo} \wedge M_{tell}^{\downarrow k} B \text{ find_Nemo}) \rightarrow B \text{ find_Nemo}. \quad (12)$$

Note that, because of (11), all agents that have reported an ability to $find_Nemo$ will be assigned the task to $find_Nemo$. From (12), it follows that as soon as one of the delegated agents reports that Nemo is found, the respective belief will be adopted.

We also make additional assumptions for the *SPEC* of the agents in *Agt*.

- If we had told an agent that we have the ability to $find_Nemo$, and that agent delegated to us the finding of Nemo, then we will indeed adopt the intention to $find_Nemo$:

$$(M_{tell}^{\uparrow i} A \text{ find_Nemo} \wedge M_{do}^{\downarrow i} I \text{ find_Nemo}) \rightarrow I \text{ find_Nemo}. \quad (13)$$

- If we believe that the $find_Nemo$ task has been accomplished, and we were delegated by an agent to find Nemo, then we tell that agent that we believe Nemo is found:

$$(B \text{ find_Nemo} \wedge M_{do}^{\downarrow i} I \text{ find_Nemo}) \rightarrow M_{tell}^{\uparrow i} B \text{ find_Nemo}. \quad (14)$$

We aim to verify whether property (16) holds for agent *seeker* when the platform is such that property (15) is satisfied (*Agt* is the set of all agents in the platform). Property (15) expresses that there is at least one agent in the neighbourhood of the Seeker that is able to find Nemo. Recall that the communication among agents is something that can be accessed by other agents, therefore we establish that an agent has an ability if it has said it has.

- There is at least one agent j in the neighbourhood of the *seeker* that has communicated possessing the ability to $find_Nemo$ to some agent k , or was told by k that k has the ability to $find_Nemo$:

$$\begin{aligned} \exists j \in NGH(seeker), M_{tell}^{\uparrow k} A \text{ find_Nemo} \in SPEC(j) \\ \text{or } M_{tell}^{\downarrow k} A \text{ find_Nemo} \in SPEC(j). \end{aligned} \quad (15)$$

- It is always true that if we advertise the incentive that the ability to find Nemo can lead to a reward (arbitrarily long), then eventually Nemo will be found:

$$\Box(M_{adv}^{\uparrow seeker} (A \text{ find_Nemo} \rightarrow \Diamond A \text{ earn}) \rightarrow \Diamond B \text{ find_Nemo}). \quad (16)$$

Let j be the agent that is in $NGH(seeker)$ s.t. $M_{tell}^{\uparrow k} A \text{ find_Nemo} \in SPEC(j)$, namely the agent that makes (15) true. Following from (7) and (8), we have that j (in $SPEC(j)$) it will eventually obtain the advert from the *seeker*:

$$\Diamond M_{adv}^{\downarrow seeker} (A \text{ find_Nemo} \rightarrow \Diamond A \text{ earn}).$$

Since (15) is true, and due to (9), (4) and (8), the *seeker* will eventually be contacted by j , namely $SPEC(seeker)$ contains

$$\Diamond M_{tell}^{\downarrow j} A \text{ find_Nemo}.$$

When $M_{tell}^{\downarrow j} A \text{ find_Nemo} \in SPEC(seeker)$ and, due to (11) and (12), we have that j (in $SPEC(j)$) will eventually obtain the delegation from the *seeker*:

$$\Diamond M_{do}^{\downarrow seeker} I \text{ find_Nemo}.$$

Due to (13), (14) and (10) the *seeker* will eventually be told (by j) that $I \text{ find_Nemo}$:

$$\Diamond M_{tell}^{\downarrow j} B \text{ find_Nemo}.$$

Lastly, due to (11) and (12), we obtain that $\Diamond B \text{ find_Nemo} \in SPEC(seeker)$ and so, eventually, the *seeker* will believe that Nemo has been found.

5.2 Software analysis

As a second example we consider a larger task, namely checking the correctness of a substantial piece of software. The task is broken down into small “human intelligence” tasks of similar complexity each of which is then crowdsourced. Each software fragment is considered (in)correct if two out of three crowd members agree on its classification. We need to verify not only that the crowd will check the whole of the software, but also that the software can be kept confidential, namely no individual crowd member can gain access to the full software, nor deduce the whole of the software from the fragments he or she is checking.

Assume that the software S is fragmented into n chunks $\sigma_1, \sigma_2, \dots, \sigma_n$. To be able to participate in the process an agent must have the required software skill level. This is tested as a condition of entry into the crowd. The crowd that tests the code contains all the agents that have been so vetted. The structure of this crowd, encapsulated in Testers, is given in Fig. 4. In the content of Testers, there are as many context-agents as there are subsystems. Individual crowding agents can choose which context they want to join, but they may not be allowed to join some or all of them, based on other contents they are members of. The s_k agents send their content lists to Testers. Then, any content-respective software fragment is only sent to members of the content.

Fig. 4 A crowd of specification testers

The relations among the software fragments are given in the specification of Testers, in form of $(B\sigma_x \wedge \dots \wedge B\sigma_z) \rightarrow B \text{ whole}$, for some $\{\sigma_x, \dots, \sigma_z\} \subset \{\sigma_1, \sigma_2, \dots, \sigma_n\}$. In addition, formulas (17–20) are included as part of the specification for Testers.

- Forward any incentives to the testers:

$$M_{adv}^{\downarrow s_i}(A \text{ test} \rightarrow \Diamond A \text{ earn}) \rightarrow M_{tell}^{\uparrow CN} M_{adv}^{\downarrow s_i}(A \text{ test} \rightarrow \Diamond A \text{ earn}). \quad (17)$$

- S is tested when each fragment is reported as tested:

$$B \text{ tested}(S) \leftrightarrow (M_{tell}^{\downarrow s_x} B \text{ tested}(\sigma_x) \wedge \dots \wedge M_{tell}^{\downarrow s_x} B \text{ tested}(\sigma_z)). \quad (18)$$

- If an agent is safe for a fragment x , it can be added to context s_x :

$$M_{tell}^{\downarrow s_x} M_{tell}^{\downarrow i} I \text{ cx}(s_x) \wedge B \text{ safe}(i, s_x) \rightarrow M_{tell}^{\uparrow s_x} B \text{ ok}(i). \quad (19)$$

- An agent is safe for a fragment x , if it cannot deduce the whole software system from the fragment x and the fragments already sent to it:⁶

$$\neg\left(\left(\bigwedge_{B \text{ safe}(i, s_z), z \neq x} B\sigma_z\right) \wedge B\sigma_x \rightarrow B \text{ whole}\right) \rightarrow B \text{ safe}(i, s_x). \quad (20)$$

The formulas (21–26) are part of the specification for s_x .

- An incentive is sent:

$$M_{adv}^{\uparrow Testers} (A \text{ test} \rightarrow \Diamond A \text{ earn}). \quad (21)$$

- If an agent is interested in joining the context s_x , and the software fragment is not yet tested, then Testers is told about this (and implicitly asked to approve the safety of the agent):

$$(M_{tell}^{\downarrow i} G \text{ cx}(s_x) \wedge \neg B \text{ approve}(\sigma_x) \wedge \neg B \text{ reject}(\sigma_x)) \rightarrow M_{tell}^{\uparrow Testers} M_{tell}^{\downarrow i} I \text{ cx}(s_x). \quad (22)$$

- If an agent is vetted, it is entrusted with a fragment for testing:

$$M_{tell}^{\downarrow s_j} B \text{ ok}(i) \rightarrow M_{do}^{\uparrow i} I \text{ test}(\sigma_x). \quad (23)$$

- If two out of three agents (i , j and k are different) approve the fragment, then the fragment is considered as being approved:

$$(M_{tell}^{\downarrow i} B \text{ approve}(\sigma_x) \wedge M_{tell}^{\downarrow j} B \text{ approve}(\sigma_x) \wedge (M_{tell}^{\downarrow k} B \text{ approve}(\sigma_x) \vee M_{tell}^{\downarrow k} B \text{ reject}(\sigma_x))) \rightarrow B \text{ approve}(\sigma_x). \quad (24)$$

- If two out of three agents (again, i , j and k are different) reject the fragment, then the fragment itself is rejected:

$$(M_{tell}^{\downarrow i} B \text{ reject}(\sigma_x) \wedge M_{tell}^{\downarrow j} B \text{ reject}(\sigma_x) \wedge (M_{tell}^{\downarrow k} B \text{ approve}(\sigma_x) \vee M_{tell}^{\downarrow k} B \text{ reject}(\sigma_x))) \rightarrow B \text{ reject}(\sigma_x). \quad (25)$$

- When the whole code fragment is approved or rejected, Testers is informed:

$$(B \text{ approve}(\sigma_x) \vee B \text{ reject}(\sigma_x)) \rightarrow M_{tell}^{\uparrow Testers} B \text{ tested}(\sigma_x). \quad (26)$$

Finally, every crowd-member is assumed to have the formulas (27–32) in its specification:

⁶ If $B \text{ whole}$ somehow becomes true independent of the specification provided by the testers then this formula will prevent any agent being deduced safe for any fragment. This represents a “fail safe” situation—no more fragments will be assigned and risk compromising the confidentiality of the software but obviously the task can not be completed in this circumstance.

- It is an agent's goal to earn, and it is able to test software:

$$G \text{ earn} \wedge A \text{ test}. \quad (27)$$

- If it is able to test software and it actively pursues the goal to test software then eventually the software will be either approved or rejected:

$$A \text{ test}(\sigma_x) \wedge I \text{ test}(\sigma_x) \rightarrow \Diamond(B \text{ approve}(\sigma_x) \vee B \text{ reject}(\sigma_x)). \quad (28)$$

- If an incentive is received, then pursue this to seize the opportunity:

$$M_{tell}^{\downarrow Testers} M_{adv}^{\downarrow s_x} (A \text{ test} \rightarrow \Diamond A \text{ earn}) \rightarrow M_{tell}^{\uparrow s_x} G \text{ cx}(s_x). \quad (29)$$

- If delegated a code fragment to test, then adopt the intention to do so:

$$M_{do}^{\downarrow s_x} I \text{ test}(\sigma_x) \rightarrow I \text{ test}(\sigma_x). \quad (30)$$

- Any code fragment can be either approved or rejected, but not both:

$$(B \text{ approve}(\sigma_x) \wedge \neg B \text{ reject}(\sigma_x)) \vee (\neg B \text{ approve}(\sigma_x) \wedge B \text{ reject}(\sigma_x)). \quad (31)$$

- When the code is tested, send the results to the relevant context:

$$\begin{aligned} B \text{ approve}(\sigma_x) &\rightarrow M_{tell}^{\uparrow s_x} B \text{ approve}(\sigma_x). \\ B \text{ reject}(\sigma_x) &\rightarrow M_{tell}^{\uparrow s_x} B \text{ reject}(\sigma_x). \end{aligned} \quad (32)$$

Establishing that (33) holds for Testers should then be straightforward.

$$\Diamond B \text{ tested}(S) \wedge \Box (M_{tell}^{\uparrow s_x} B \text{ ok}(i) \rightarrow B \text{ safe}(i, s_x)). \quad (33)$$

5.3 Utilising the formal basis

Above we have given two examples showing both how our formal syntax can be used to specify digital crowd scenarios, and then how formal verification processes can use these specifications to establish properties. The important aspects concerning this process are:

1. the general properties of the crowd, and environment, are formalised using our language;
2. the specific properties of individual agents within the digital crowd are also formalised using the same language;
3. we also formalise any assumptions we make about agent (hence, human) or environmental behaviour;
4. for an individual agent, formal verification of its properties can be carried out by invoking existing agent model-checking tools such as [8], which prove that all possible executions of the agent conform to the logical requirements;

5. once proved for individual agents, reasoning about combined properties can take place, either using manual proof (as above) or by some form of automated process (note there are many automated provers for modal and temporal logics); and
6. once complete, we have verified properties of the digital crowd scenario under the assumptions specified and so can be sure what behaviour will occur (if the assumptions are satisfied).

Once this process is complete, it is natural to then revisit the assumptions, weaken them and see if the verification can still be carried through. If it cannot, then we can see where the required behaviour can fail and so can either take this back to the crowd application designer, or weaken the properties being verified. And so on. The process continues in this way until sufficient verified behaviour has been extracted. Note that we clearly cannot completely specify all crowd behaviour, but can explore classes of behaviour and prove what will occur under these assumptions.

6 Summary

Over the last few years, the business practice of crowdsourcing has begun to capture the attention of computer scientists. All crowdsourcing processes involve the participation of a digital crowd, a large number of agents that access a single Internet platform or web service. Although a digital crowd is a collection of agents, it is not a structure traditionally studied within the area of multi-agent systems. Logic-based formal methods are available for analysing the behavior of, and the dynamics within, a multi-agent system before the actual system is constructed. In particular, the formal verification of systems is the process of establishing that a designed system has the intended properties. Crowdsourcing can be made more reliable and effective by applying such logic-based formal methods by, for example, determining important properties of a digital crowd, under given assumptions about its members, before the crowd is assembled.

Our aim is to enable automated formal verification, in particular model-checking, to be utilised in digital crowd systems, and crowdsourcing in general. To this end we extend the paradigm of an agent, in particular a *BDI* agent, as used in traditional multi-agent systems. Our extended agent encompasses both communication behaviour and further individual agents and sub-structures. We accomplish this by abstracting from the member agents' mental states and directly specifying the communication exchanged by the agents.

We proposed an abstract agent specification language, adding the representation of messages to a temporal *BDI* language. Although most agent programming languages allow messages to be passed among agents, these messages are primarily signals (to *stop*, *go*, etc) and are not part of the reasoning within the agent. The content of agent communication typically is not considered a separate informational state of the agent. We represent the messages in the same fashion as the informational, motivational and dynamic mental attitudes of the *BDI* agent are represented. While nothing is assumed about the mental states of a particular agent, the information it exchanges with other agents can be sufficient to reason about a structure in which that agent belongs, and can represent loose social structures such as digital crowds. To exhibit this, we develop

two examples to illustrate how our formal specification language can be used as the basis for formally describing properties of a digital crowd.

The possibility of applying formal methods to digital crowds opens many interesting avenues for future research. Model-checking requires the presence of an executable model of the system under investigation. There are executable specification languages for agent systems based on temporal logic [13,42] and we are interested in adapting these so that our specifications can be easily converted into appropriate executable models. We anticipate that we will need to develop abstraction mechanisms e.g., a suitable crowd size for a model may be the number of distinct crowd members referred to in the property description. Extending the problem with probabilistic aspects would also be of interest. This would enable reasoning about the probable success of some task given to a digital crowd based on (stochastic) assumptions about the probable abilities (and reliability) of members of the crowd. Probabilistic Model-checking tools (e.g., PRISM [26]) could be adapted to study digital crowds.

Finally, while we do not need to axiomatize our specification language to use it in model-checking it would be interesting to develop an axiomatization for the M operators and formally study its properties. We currently make no use of plans and planning in our examples, though planning is an important part of multi-agent systems and it would be interesting to integrate the problem of planning with, and for, digital crowds.

Acknowledgments This work was partially funded by the EPSRC within the “Verifying Interoperability Requirements in Pervasive Systems” (EP/F033567) and the “Reconfigurable Autonomy” (EP/J011770) projects.

Open Access This article is distributed under the terms of the Creative Commons Attribution License which permits any use, distribution, and reproduction in any medium, provided the original author(s) and the source are credited.

References

1. Barbuceanu, M., Fox, M.S.: COOL: A Language for Describing Coordination in Multi Agent Systems. In: ICMAS, pp. 17–24. The MIT Press, San Francisco (1995)
2. Belnap, N., Perloff, M.: Seeing to it that: a Canonical form for agentives. *Theoria* **54**(3), 175–199 (1988)
3. Bordini, R.H., Hübner, J.F., Vieira, R.: Jason and the Golden Fleece of agent-oriented programming. In: Bordini, R.H., Dastani, M., Dix, J., El Fallah Seghrouchni, A. (eds.) *Multi-Agent Programming: Languages, Platforms and Applications*, pp. 3–37. Springer (2005)
4. Clarke, E.M., Grumberg, O., Peled, D.: *Model Checking*. MIT Press, Cambridge (1999)
5. de Boer, F., van Eijk, R., van der Hoek, J.-J., Meyer, W.: A fully abstract model for the exchange of information in multi-agent systems. *Theor. Comput. Sci.* **290**(3), 1753–1773 (2003)
6. de Boer, F.S., Hindriks, K.V., van der Hoek, W., Meyer, J.-J.Ch.: A verification framework for agent programming with declarative goals. *J. Appl. Logic* **5**(2), 277–302 (2007)
7. Dennis, L.A., Farwer, B.: Gwendolen: A BDI Language for Verifiable Agents. In: *Proceedings of the AISB Workshop on Logic and the Simulation of Interaction and Reasoning*. AISB (2008)
8. Dennis, L.A., Fisher, M., Webster, M., Bordini, R.H.: Model checking agent programming languages. *Autom. Softw. Eng.* **19**(1), 5–63 (2012)
9. Dix, J., Fisher, M.: *Multiagent Systems. Specification and verification of multi-agent systems*, vol. 14. MIT Press, Cambridge (2013)
10. Elgesem, D.: The modal logic of agency. *Nordic J. Philos. Logic* **2**(2), 1–46 (1997)

11. Estellés-Arolas, E., González-Ladrón-De-Guevara, F.: Towards an integrated crowdsourcing definition. *J. Inf. Sci.* **38**(2), 189–200 (2012)
12. Estellés Arolas, E., González-Ladrón-De-Guevara, F.: Clasificación de iniciativas de crowdsourcing basada en tareas. *El Profesional de la Información* **21**(3), 283–291 (2012)
13. Fisher, M.: A survey of concurrent METATEM: the language and its applications. In: Gabbay, DovM, Ohlbach, HansJürgen (eds.) *Temporal Logic, Lecture Notes in Computer Science*, vol. 827, pp. 480–505. Springer, Berlin (1994)
14. Fisher, M.: *An Introduction to Practical Formal Methods Using Temporal Logic*. Wiley, Hoboken (2011)
15. Fisher, M., Dennis, L., Hepple, A.: *Modular Multi-Agent Design*. Technical Report ULCS-09-002, Department of Computer Science, University of Liverpool (2009)
16. Fisher, M., Kakoudakis, T.: Flexible Agent Grouping In Executable Temporal Logic. In: *Proceedings of the 12th International Symposium on Languages for Intensional Programming (ISLIP)*. World Scientific Press (1999)
17. Fornara, N., Okouya, D., Colombetti, M.: Using OWL 2 DL for Expressing ACL Content and Semantics. In: *Multi-Agent Systems*, volume 7541 of *Lecture Notes in Computer Science*, pp. 97–113. Springer, Heidelberg (2012)
18. Hepple, A., Dennis, L., Fisher, M.: Languages, Methodologies and Development Tools for Multi-Agent Systems. A common basis for agent organisation in BDI languages, pp. 71–88. Springer, Berlin (2008)
19. Hindriks, K.V., de Boer, F.S., van der Hoek, W., Meyer, J-JCh.: Agent programming in 3APL. *Auton. Agents Multi-Agent Syst.* **2**(4), 357–401 (1999)
20. Ho, C.J., Zhang, Y., Wortman Vaughan, J., van der Schaar, M.: Towards Social Norm Design for Crowdsourcing Markets. In: *Proceedings of Human Computation Workshop* (2012)
21. Horling, B., Lesser, V.: A survey of multi-agent organizational paradigms. *Knowl. Eng. Rev.* **19**(4), 281–316 (2004)
22. Howe, J.: The Rise of Crowdsourcing. *Wired*, Issue 14.06, (2006)
23. Jennings, N.R.: On agent-based software engineering. *Artif. Intell.* **117**(2), 277–296 (2000)
24. Jennings, N.R., Sycara, K., Wooldridge, M.: A roadmap of agent research and development. *Auton. Agents Multi-Agent Syst.* **1**(1), 7–38 (1998)
25. Kamar, E., Horvitz, E.: Incentives for Truthful Reporting in Crowdsourcing. In: *Proceedings of the International Conference AAMAS*, pp. 1329–1330 (2012)
26. Kwiatkowska, M., Norman, G., Parker, D.: PRISM: Probabilistic Symbolic Model Checker. In: *Proceedings of the 12th International Conference Modelling Techniques and Tools for Computer Performance Evaluation (TOOLS)*, vol. 2324 LNCS. Springer, Berlin (2002)
27. Lofland, J.: Collective behaviour: the elementary forms. In: Curtis, R.L., Aguirre, B.E. (eds.) *Collective Behaviour and Social Movements*, pp. 70–75. Pearsons, Boston (1993)
28. Mascardi, V., Martelli, M., Sterling, L.: Logic-based specification languages for intelligent software agents. *Theory Pract. Log. Program.* **4**(4), 429–494 (2004)
29. McBurney, P., Parsons, S.: Locutions for argumentation in agent interaction protocols. Agent communication. In: van Eijk, R., Huget, M.-P., Dignum, F. (eds.) *Lecture Notes in Computer Science*, vol. 3396, pp. 209–225. Springer, Berlin (2005)
30. Olson, P.: *We Are Anonymous: Inside the Hacker World of LulzSec, Anonymous, and the Global Cyber Insurgency*. Little, Brown and Company, New York (2012)
31. Papadimitriou, C.H.: *Computational Complexity*. Addison-Wesley, Reading (1994)
32. Pitt, J., Mamdani, A.: Some remarks on the semantics of fipa’s agent communication language. *Auton. Agents Multi-Agent Syst.* **2**(4), 333–356 (1999)
33. Pokahr, A., Braubach, L., Lamersdorf, W.: Jadex: A BDI reasoning engine. In: Bordini, R.H., Dastani, M., Dix, J., El Fallah Seghrouchni, A. (eds.) *Multi-Agent Programming: Languages, Platforms and Applications*, pp. 149–174. Springer, Berlin (2005)
34. Raimondi, F., Lomuscio, A.: Automatic verification of multi-agent systems by model checking via ordered binary decision diagrams. *J. Appl. Logic* **5**(2), 235–251 (2007)
35. Rao, A. S.: AgentSpeak(L): BDI Agents Speak Out in a Logical Computable Language. In: *Agents Breaking Away: Proceedings of 7th European Workshop on Modelling Autonomous Agents in a Multi-Agent World*, vol. 1038 LNCS, pp. 42–55. Springer, New York (1996)
36. Rao, A.S., Georgeff, M.P.: Modelling Agents within a BDI-Architecture. In: *Proceedings of the International Conference Principles of Knowledge Representation and Reasoning, KR*. Morgan Kaufmann (1991)

37. Rao, A.S., Georgeff, M.P.: BDI Agents: from Theory to Practice. In: Proceedings of the 1st International Conference Multi-Agent Systems (ICMAS), pp. 312–319, San Francisco, USA (1995)
38. Rheingold, H.: *Smart Mobs: The Next Social Revolution*. Perseus Books, Cambridge (2003)
39. Schenk, E., Guittard, C.: Towards a characterization of crowdsourcing practices. *J. Innov. Econ.* **1**(7), 93–107 (2011)
40. Schild, K.: On the relationship between BDI logics and standard logics of concurrency. *Auton. Agents Multi-Agent Syst.* **3**(3), 259–283 (2000)
41. Searle, J.R.: *Language, Mind, and Knowledge* (Minneapolis Studies in the Philosophy of Science). A Taxonomy of Illocutionary Acts, vol. 7, pp. 344–369. University of Minneapolis Press, Minneapolis (1975)
42. Shoham, Y.: Agent-oriented programming. *Artif. Intell.* **60**(1), 51–92 (1993)
43. Singh, M.P.: Agent Communication Languages: Rethinking the Principles. Communication in Multi-agent Systems. In: Huget, Marc-Philippe (ed.) *Lecture Notes in Computer Science*, vol. 2650, pp. 37–50. Springer, Berlin (2003)
44. Stirling, C.: *Handbook of Logic in Computer Science. Modal and temporal logics*. Oxford University Press, Oxford (1992)
45. Troquard, N.: Reasoning about coalitional agency and ability in the logics of “bringing-it-about”. *Auton. Agents Multi-Agent Syst.* **28**, 1–27 (2013)
46. van Ditmarsch, H., van der Hoek, W., Kooi, B.: *Dynamic Epistemic Logic*, 1st edn. Springer, Berlin (2007)
47. Vieira, R., Moreira, A.F., Wooldridge, M., Bordini, R.: On the formal semantics of speech-act based communication in an agent-oriented programming language. *J. Artif. Intell. Res. (JAIR)* **29**, 221–267 (2007)
48. Winikoff, M.: Implementing Commitment-Based Interactions. In: Proceedings of the 6th International Conference Autonomous Agents and Multiagent Systems (AAMAS), pp. 1–8. ACM (2007)
49. Wobcke, W.: Agents and Multi-Agent Systems Formalisms, Methodologies, and Applications. *Agency and the Logic of Ability*, vol. 1441, pp. 31–45. Springer, Berlin (1998)
50. Wooldridge, M.: Semantic issues in the verification of agent communication languages. *Auton. Agents Multi-Agent Syst.* **3**, 9–31 (2000)
51. Wooldridge, M., Ciancarini, P.: *Agent-Oriented Software Engineering. Agent-Oriented Software Engineering: The State of the Art*, vol. 1957, pp. 1–28. Springer, Berlin (2001)
52. Wooldridge, M., Jennings, N.R.: Intelligent agents: theory and practice. *Knowl. Eng. Rev.* **10**(2), 115–152 (1995)
53. Zhao, Y., Zhu, Q.: Evaluation on crowdsourcing research: current status and future direction. *Inf. Syst. Front.* **8**, 1–18 (2012)