

A roadmap to pervasive systems verification

SAVAS KONUR¹ and MICHAEL FISHER²

¹*Department of Computer Science, University of Sheffield Regent Court, 211 Portobello, Sheffield, S1 4DP, United Kingdom;*
e-mail: S.Konur@sheffield.ac.uk;

²*Department of Computer Science, University of Liverpool Ashton Building, Ashton Street, Liverpool, L69 3BX, United Kingdom;*
e-mail: MFisher@liverpool.ac.uk

Abstract

The complexity of *pervasive systems* arises from the many different aspects that such systems possess. A typical pervasive system may be autonomous, distributed, concurrent and context based, and may involve humans and robotic devices working together. If we wish to formally verify the behaviour of such systems, the formal methods for pervasive systems will surely also be complex. In this paper, we move towards being able to formally verify pervasive systems and outline our approach wherein we distinguish four distinct dimensions within pervasive system behaviour and utilize different, but appropriate, formal techniques for verifying each one.

1 Introduction

We are here concerned with the formal modelling and verification of *pervasive systems*, covering a very broad class of systems that are typically mobile, dynamic, autonomous, distributed, concurrent and context based, and may involve humans and robotic devices working together (Dobson *et al.*, 2010).

Let us consider a very typical pervasive system example that we have used in several earlier papers, for example (Arapinis *et al.*, 2009):

As we walk into a shopping centre, our intelligent clothing interacts wirelessly with shops in the area and then with our mobile phone to let us know that our shoes are wearing out and that the best deals nearby are at shops X, Y and Z.

Our computer, which holds a shopping list, also interacts with our phone to suggest the optimum route to include shoes in our shopping, and with the shops to assess their stock.

At the same time, our computer interacts with the shopping centre's network and finds that one of our friends is also shopping, so a message is sent to the friend's mobile/computer to coordinate shopping plans and schedule a meeting for coffee at a close location in 15 minutes.

This might not yet exist, but such scenarios are neither very complex nor very far from being reality! There are currently several deployed pervasive systems partially employing similar scenarios, for example, Scatterbox (a message forwarding system) (Knox *et al.*, 2008), MATCH (a health-care system for elderly and disabled people at home) (Turner, 2012), CityEvents (a mobile multimedia system for displaying geolocated cultural events) (Rosa *et al.*, 2012), etc.

There are, of course, many definitions of a pervasive system (Want *et al.*, 2002; Greenfield, 2006); even more if we include *ubiquitous systems* (Weiser, 1993) or *adaptive/autonomic systems* (Dobson *et al.*, 2006). In this paper, rather than choosing one of these (often ambiguous) definitions, we will instead describe the

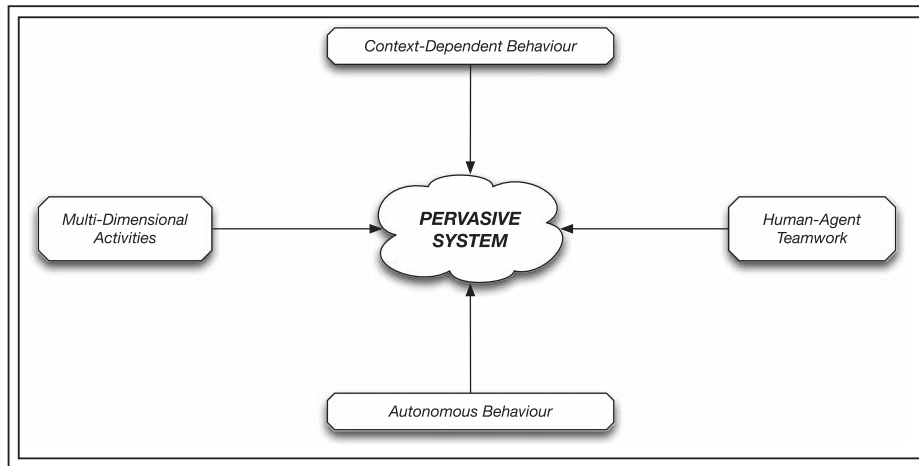


Figure 1 Four orthogonal aspects of pervasive systems

systems we are concerned with in terms of the general classes of behaviour we *expect to see* within them. Thus, the ‘pervasive systems’ we address here typically contain the following aspects:

1. multiple interacting entities, with multi-dimensional properties;
2. truly autonomous decision making;
3. the implicit involvement of humans *within* the system; and
4. context-dependent behaviour.

While we will explain these aspects in more detail in Section 2, it is important to note that many systems might have very simple/limited behaviours in one or more of these areas. For example, a pervasive system that has no interaction with humans clearly has a trivial behaviour in (3) above. It is also important to note that the four elements above are very general and while, for our analysis purposes, we consider them to be distinct, in practice there will be unavoidable overlap.

So, in describing and analyzing such pervasive systems, we choose to treat the above four aspects orthogonally as in Figure 1. As we see later, we verify properties of pervasive systems by verifying each aspect independently. However, in the next section, we will attempt to justify the above categorization, appealing to typical examples of pervasive systems.

Example. At this point, an obvious question is whether any real pervasive system can appropriately be described in terms of our four dimensions. At present we can only answer that all the example systems we have encountered can be viewed in this way. Here, we highlight a large pervasive system that can be potentially be analyzed as we suggest.

The K4Care platform (Isern *et al.*, 2011) is aimed at organizing and providing health-care services in ‘home care’ domains. Specifically, it helps medical practitioners (e.g. doctors, physicians, specialists and nurses) provide administrative and medical services, simplifying the coordination between participants. The K4Care platform also features tools and methods helping the practitioners prepare personalized treatments for patients based on health-care standards.

We briefly identify our four views on the K4Care example:

1. The K4Care model is implemented on the K4Care platform, which is itself distributed and multi-layer. For example, the platform involves a *knowledge layer*, a *data abstraction layer* and an *agent-based layer*. That all these aspects must interact and that real-time response is clearly of importance, together points towards quite complex multi-dimensional descriptions.
2. The K4Care system is agent based and involves a quite sophisticated range of reasoning techniques, for example, about treatments, effects, diseases, etc. There appears to be a certain amount of autonomous decision making in the system, again conforming to one of our key dimensions.

3. The provision of such health services involves not only interaction with a range of medical practitioners (e.g. doctors, nurses) but also with ‘home care patients’. Such patients are typically elderly, with illness and/or physical impairment. Modelling the activities of the K4Care system clearly involves human–system interactions.
4. Finally, the system utilizes *context*, primarily in the form of ontologies, both at the level of reasoning about treatment/procedures, and at the level of reasoning about activities in different situations.

Clearly, semi-automated home care support is an area where safety, security, reliability and efficiency are all vital, and so formal methods, particularly formal verification, can surely be used.

2 Characteristics of pervasive systems

We now consider the four aspects from Figure 1 in more detail.

2.1 Multi-dimensional activity

Pervasive systems comprise many different facets and so we will need to describe/specify not just their basic dynamic behaviour, but also how all these facets/dimensions evolve and interact. Even in the relatively trivial ‘shopping’ example above, we can see that the system involved has many different facets.

- It clearly has *dynamic* behaviour, together with (changing) knowledge about people and their plans.
- There are *real-time* aspects, not only in scheduling an appropriate meeting time but also in responding to shopper movements in a timely manner.
- As at least some part of the pervasive system has to try to model the ‘real world’, or at least an abstraction of the environment, and since the sensors that are typically used for this can never be perfect, then *uncertainty* must, necessarily, be involved.
- There is *collaboration* and *cooperation*, both at the level of agents ‘representing’ shoppers and at the level of shops and agents cooperating to provide appropriate information and capabilities.
- Many elements are mobile, and there are many programs working at once but in different areas. Thus, *mobility*, *distribution* and *concurrency* are essential.
- Clearly, the software representing a shopper must also carry out *autonomous decision making* (see Section 2.2), for example, to instigate a meeting for coffee.

And so on. As we can see, even from this simple example, a pervasive system can have many different facets. As we wish to specify and verify such systems, then we must ensure our specification (Section 3.1) and verification (Section 4.1) techniques take these dimensions into account.

We will later see how a typical pervasive message forwarding system, such as Scatterbox (Knox *et al.*, 2008), requires a wide range of additional dimensions in order to explain/describe its behaviours:

Uncertain behaviour/accuracy of sensors	<i>probabilistic (behaviours)</i>
Location and tracking	<i>spatial, dynamic</i>
System <i>beliefs</i> about location and activity	<i>knowledge</i>
Importance of messages	<i>priority</i>
Varying levels of user interruptibility	<i>priority</i>
Communication connection failures	<i>probabilistic</i>
Requirements of system efficiency	<i>real time</i>
Security of sensitive data	<i>knowledge, temporal</i>
Modelling intelligent choices	<i>motivations</i>
Roles and responsibilities	<i>multi-agent</i>

So, in general, to be able to clearly specify any non-trivial pervasive system we will require descriptions that combine, and explain interactions between, many such dimensions.

2.2 Autonomous decision making

Pervasive systems contain components that can make autonomous choices, that is, are often separated from direct human control. Thus, with this *autonomy* the computational components can decide, for

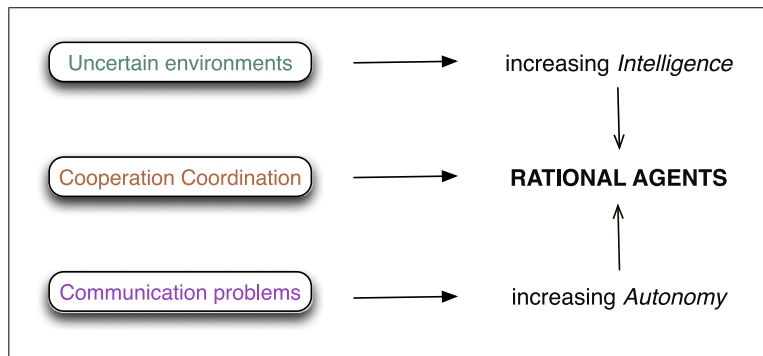


Figure 2 Motivation for use of rational agents

themselves, what to do and when to do it. The specific form of autonomy we consider is that captured by the concept of a *rational agent*. An agent (Wooldridge, 2002) essentially captures the idea of an autonomous entity, being able to make its own choices and carry out its own actions. Beyond this, *rational agents* are typically employed to represent more complex autonomous decision making where communication and explanation are vital (see Figure 2).

These agents are *rational* in the sense that they take their decisions in a *rational* and *explainable* way (rather than, e.g., purely randomly). The central aspects of the rational agent metaphor is that such agents are autonomous, but can react to changes in circumstance and can choose what to do based on their own agenda. In assessing such systems it may not be sufficient to consider *what* the agent will do, but we must often also consider *why* it chooses to do it. The predominant view of rational agents is that provided by the BDI (beliefs–desires–intentions) model (Rao & Georgeff, 1991, 1995) in which we describe the goals the agent has, the ‘beliefs’ it has and the choices it makes. The rational agent’s *intentions* are then determined by its *beliefs* and *goals*.

2.3 Humans ‘in the loop’

Pervasive systems typically involve *humans*, or at least external autonomous entities, and a variety of *tools* or *artefacts* within the system. Often, humans are actually *embedded* within pervasive systems, with computation and communication happening all around them. In many cases, the overall behaviour of the whole pervasive system depends on appropriate behaviour of the humans involved; in this case we can see the human/system relationship as a very weak form of cooperation.

In some areas, computational/robotic elements play an even more collaborative role in the human–system scenario. This is particularly the case if actual physical robots are involved. Although many practical activities suggest human–robot cooperation, space exploration is an area at the forefront of this where *cooperation* and *teamwork* is clearly required. As an example, NASA is interested in *Robot–Astronaut* teams to be deployed on Mars (Sierhuis *et al.*, 2003; see Figure 3(a)).

While not directly involving such space applications, pervasive systems not only can involve multiple robotic elements but also can require human–system interaction of a similar form. Consequently, the problems raised by human–robot interactions can often re-appear in pervasive scenarios. On the domestic front, the close interaction of multiple humans and robots in pervasive scenarios is becoming more prevalent, particularly in health-care situations. Thus, the Robotics industry is actively developing ‘*Robots for the Home*’, particularly as human assistants (see Figure 3(b)). Clearly, this leads to safety concerns:

- Are the robots *safe*?
- How can we tell a robot what we *want*?
- *Will* a robot work out what we want it to do?
- And then, are we *sure* it will decide to do what we want?




Figure 3 Human–robot teamwork (a) in planetary exploration (photo courtesy of NASA (Astronaut-Robot-Team-Concept)) and (b) in the home (RoboSafe, 2013)

This is a natural area where formal verification should, and can, make an impact (Fisher *et al.*, 2013). Indeed, as we will see later, specification and verification techniques developed for human–robot interaction have an important role to play in assessing (high-level and abstracted) human involvement in pervasive systems.

2.4 Context dependency

Pervasive systems are also *context dependent*—their behaviour can change *because* of movement between contexts. This can take (at least) two forms.

2.4.1 Simple context view

In the first form, ‘context modelling’ is essentially concerned with recording someone’s location or activity (Chen *et al.*, 2003). For example, detecting a Bluetooth signal of one’s cellular phone in a seminar room can be used to derive his/her context (in the form of location or activity) in conjunction with other information available, for example, calendar.

This context modelling is a popular technique studied by software engineers who wish to differentiate subtle behavioural differences in systems (Henricksen & Indulska, 2004; Strang & Popien, 2004). These behaviours are often in response to, or in anticipation of, human activity and some notion of ‘context’ of this form is at the heart of very many pervasive systems.

2.4.2 Complex context view

Another, more general, view is that ‘context’ should cover *much* more than location or activity. For example, in the METATEM agent programming language (Fisher & Hepple, 2009) the context in which an agent is acting is a first-class construct. Indeed, contexts are also agents and an agent is considered to reside in (or be contained by) zero or more contexts. Context in METATEM is conceived as the more abstract and flexible notion of *anything* that has an influence over an agent in a system. Specifically, contexts do *not* just concern location, but can be

- roles or societal norms;
- teams or organizational structures;
- Activities;
- styles or preferences;
- Locations;
- *etc.*

This gives a very powerful mechanism for programming and describing quite complex context-dependent behaviours and so complex organizations and systems (Hepple *et al.*, 2008). While there remains debate about exactly what an ‘agent’ is, this view of hierarchical and active agents perhaps needs further comment. While we can see an agent as capturing basic autonomous behaviour, there are many

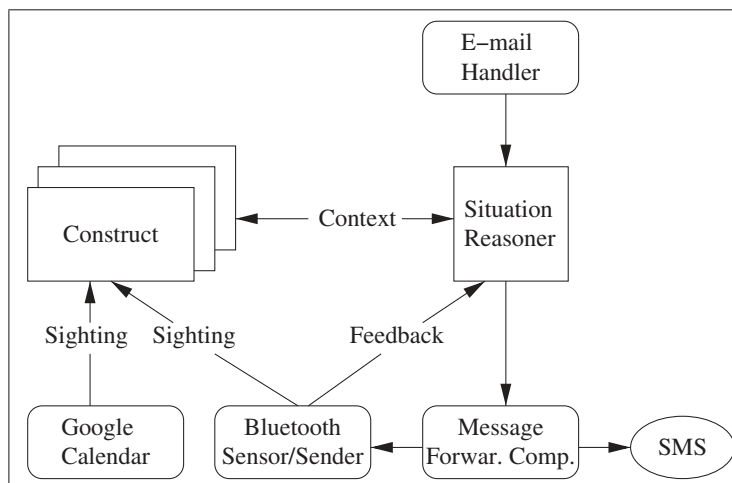


Figure 4 Scatterbox message forwarding system (Konur *et al.*, 2014)

different levels of abstraction over which this can be taken. Consider, for example, a flock of birds¹. Seen at a high level of abstraction, the ‘flock’ can be viewed as an autonomous entity with its emergent behaviour being opaque to an external viewer. So, we could model the flock as an agent. But then, of course, if we move into the flock we can discern local interactions between individual birds. Correspondingly, once we look inside the agent’s content we can see (and interact with) representations of individual birds.

This duality, being able to treat the larger entity as an agent while at the same time being able to move inside to interact with further agents in its content, is very powerful. It clearly has ancestry on object-based systems, but extends this to agents with active, autonomous behaviour. The context/content view is, thus, very useful for modelling entities that are autonomous and active at many different levels of abstraction and often do not correspond to physical entities, for example, roles and norms, teams and other organizational structures, styles and preferences, etc.

It is important, in our later sections, that we are able to verify both varieties of the notion of ‘context’ described above. The first is simpler, but is used in many deployed pervasive systems. The second approach is more powerful, and something like this is likely to become used in future pervasive system development. Below we give an example for each approach.

Example (Knox *et al.*, 2008). The Scatterbox system, as well as exhibiting multi-dimensional behaviour as described in Section 2.1, also involves context-dependent behaviours. As described in Knox *et al.* (2008), Scatterbox is a message forwarding system, primarily designed as a testbed for context-aware computing in pervasive scenarios. Each Scatterbox user’s context is defined by tracking his/her location and by monitoring his/her daily schedule. These data are analyzed and ‘situations’ are identified that indicate the user’s level of interruptibility. Once a message arrives, Scatterbox forwards it to the subscribed user provided the user’s available context suggests they are in a situation where they may be interrupted. A schematic diagram of Scatterbox is given in Figure 4.

Example (Hepple, 2010). Consider an abstract ‘cooking’ example. Here, a *cooking agent* simply wants to ‘cook’. But, to begin with, it does not know *how* to cook and does not know about any contexts. Specifically, it initially has one goal (i.e. to ‘cook’) but has no plans telling it how to go about this. So, this *cooking agent* now moves into a *chef role* context, which is itself another agent (see Figure 5(a)). This context agent provides the cooking agent with appropriate information, plans and capabilities associated with the ‘role’ of being a ‘chef’ (see Figure 5(b)). Once the agent has received and assimilated these aspects then it can legitimately take on the role of a ‘chef’ and can indeed satisfy its goal to be able

¹ Our thanks to an anonymous reviewer for highlighting this.

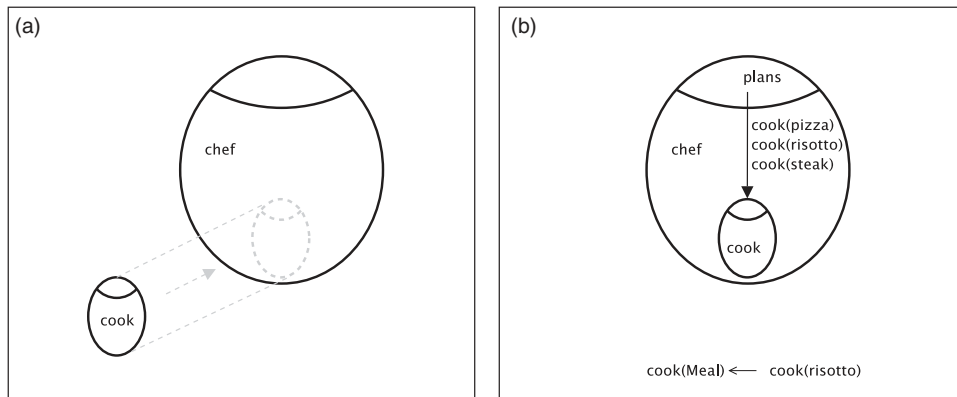


Figure 5 Cooking scenario (Dennis *et al.*, 2008)

to 'cook'. So, just by moving into a new context, the cookery agent has obtained new computational capabilities. Moving into further contexts can bestow additional capabilities, or even constraints, upon the agent. In this way, contexts are active, providing more than just a representation of 'location'.

3 Specifying pervasive systems

Following our theme, if we are to program a pervasive system or specify system properties, then we must address each of our four dimensions in turn (see Figure 6). In this section, we will both discuss how the properties of multidimensionality are specified and how certain aspects of systems are programmed.

3.1 Multi-dimensional activity

In the *specification* of different aspects and dimensions, we use *formal logics* (Konur, 2013). Logics provide a well-understood and unambiguous formalism for describing the *behaviours*, *requirements* and *properties* of systems. In addition, as we are concerned with the verification of pervasive systems, the well-researched proof, complexity and satisfiability in checking results can provide us with many practical tools such as automated deduction or model checking systems.

An obvious issue with formally specifying pervasive systems is that formal logics have limitations. For example, there might be system properties that cannot be represented formally in a particular logical system. In order to maximize the expressivity of logics and hence to be able to verify more complex properties, we combine different aspects of different logics.

Beyond classical Boolean logic there are many logics devised for a huge array of different types of system. Choosing the appropriate logic provides a level of abstraction close to the key concepts of the software, for example:

- dynamic communicating systems → *temporal logics*;
- systems managing information → *logics of knowledge*;
- situated systems → *logics of belief, contextual logics*;
- timed systems → *real-time temporal logics*;
- uncertain systems → *probabilistic logics*;
- cooperative systems → *cooperation/coalition logics*.

As we cannot easily describe all aspects of a pervasive system in one of these logics, we will often need to *combine* formalisms. There are several standard ways to combine logics, with the three standard ones being (Gabbay *et al.*, 2003): *temporalization* (Finger & Gabbay, 1996); *fusion* (or *independent join*); and *product* (or *full join*). Now, imagine we have two logics to combine, logic A (typically, a temporal logic) and logic B.

We assume the logics A and B are graphically represented as in Figure 7.

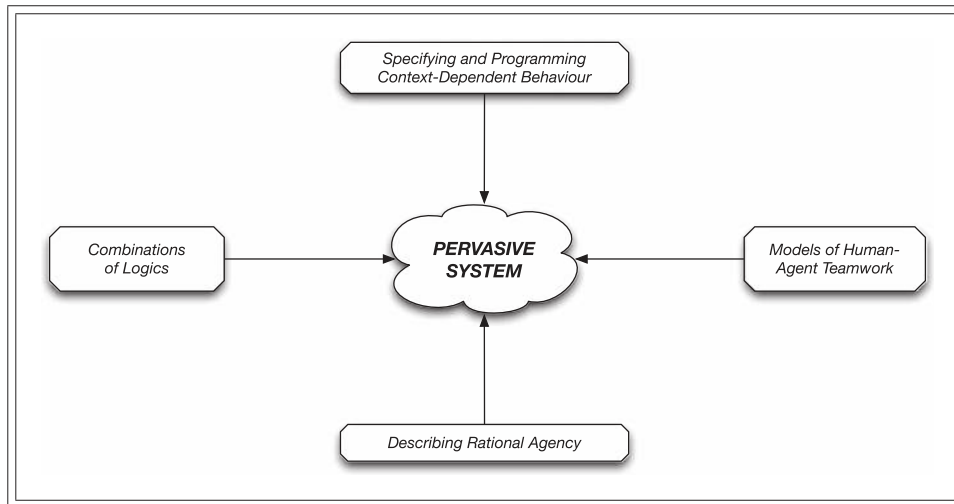


Figure 6 Four orthogonal aspects for specifying pervasive systems

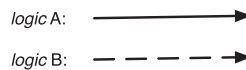


Figure 7 Notation (Konur *et al.*, 2013)

The *temporalization* permits only a restricted way that two component languages can interact with each other. It is denoted by $A(B)$, where a pure subformula of logic B can be treated as an atom within logic A. This combination is *not* symmetric—A is the main logic, but at each world/state described by logic A we might have a formula of B describing an embedded ‘B-world’ (see Figure 8).

As Figure 8 suggests, after we take some A-step(s) (straight arrow), we can take a B-step (dotted arrow). From this point on, we are in the B-world, and we cannot take an A-step anymore.

The *fusion* $A \oplus B$ is more symmetric than temporalization in that, at any state/world, we can either take an ‘A-step’ or a ‘B-step’ (see Figure 9). It is important to note that the two logical dimensions are essentially independent (hence ‘independent join’).

The *product* combination, $A \otimes B$, is similar to the fusion, but with a much tighter integration of the logics (see Figure 10). Operators of the constituent logics tend to be *commutative*. This provides a stronger integration of the component logics but, correspondingly, often results in logics of higher complexity than fusions or temporalizations.

Example. Once we can combine logics, for example, using fusions, then we can combine several logics together in order to provide a quite powerful descriptive framework. As an example of why combinations of logics may be appropriate for formally specifying multi-dimensional behaviour in a pervasive system, consider the following statement from an assistive health-care scenario.

If a patient needs medicine, then the patient believes that there is a probability of 65% that, within 10 minutes, a nurse robot will want to assist the patient.

In formalizing this we might write:

$$needs_medicine(patient) \rightarrow B_{patient}^{\geq 0.65} \diamond^{\leq 10} G_{nurse} assist(patient)$$

where

$B_{patient}^{\geq 0.65}$ means *patient believes with 95% probability*: probabilistic logics;

$\diamond^{\leq 10}$ means *within 10 minutes*: temporal logics;

G_{nurse} means the nurse robot *has a goal*, and so *wants* to help the patient: BDI logics.

This means that if the agent intends to eat and if it believes it has money, but does not have food, then it will intend to go to the shop. Once the agent achieves its desire (i.e. goal), it will buy some food and then intends to go home. When it is at home, the agent will eat and then update its beliefs to record that it believes it has eaten.

Modern agent programming languages in this BDI paradigm were pioneered by the PRS in the early 1990s (Georgeff & Lansky, 1987; Rao & Georgeff, 1992), which sought to abstract *decision making* away from the more procedural aspects where control systems were used. This BDI architecture has been shown to generalize widely to any situation in which autonomous decision making is required to control complex software processes and a family of programming languages has been developed to support the paradigm, for example, AGENTSPEAK (Rao, 1996; Bordini *et al.*, 2005), 2APL (Dastani *et al.*, 2005), GOAL (Hindriks *et al.*, 2001) and Gwendolen (Dennis & Farwer, 2008).

3.3 Modelling teamwork: humans ‘in the loop’

In a pervasive system, we can often see humans and autonomous components as working together as a *team*. But how might we try to describe behaviours within such human–computer teams? To simplify, we choose to model the *high-level* cooperation and interaction, behaviours of the team participants, abstracting away from low-level details such as movement. Specifically, we choose to carry out the modelling of these high-level behaviours using the *Brahms* framework (van Hoof, 2000; Sierhuis, 2001; Sierhuis & Clancey, 2002). This is a framework designed to model both human and robot activity, again using *rational agents*. As it was developed to represent activities in real-world contexts, it also allows the representation of artefacts, data and concepts in the form of classes and objects. Both agents and objects can be located in a model of the world (the geography model) giving agents the ability to detect objects and other agents in the world and have beliefs about the objects. Agents can move from one location in the world to another by executing *move* actions, so simulating physical movement (for a more detailed description, see Sierhuis (2001, 2006).

The Brahms language was originally conceived of as a language for modelling contextually situated activity behaviour of groups of people. This has now evolved into a language for modelling both people *and* robots/agents. As such it is ideal for describing human–agent/robot teamwork and indeed has been used for this both in mundane scenarios and in modelling the ‘Robot–Astronaut teams on Mars’ mentioned earlier. It has been particularly successful as it allows us to describe, concisely and precisely, the overall team behaviours and so assess how well such teams work together.

Example (Stocker *et al.*, 2012). We now consider a domestic home care application as an example to illustrate the Brahms specification language. In this scenario, the helper robot reminds the assisted person Bob to take his medication. The scenario is explained as follows:

The Robot has a workframe to deliver medicine to Bob; activated at pre-allocated times. The Robot places the drugs on Bob’s tray and then monitors them hourly to check if they have been taken. A detectable takenMedicationC aborts the workframe if the drugs have been taken and then updates the Robot’s beliefs. If the drugs have not been taken the workframe reminds Bob to take his medication. The Robot counts the number of times it reminds Bob, and after 2 reminders it notifies the House.

A fragment of the Brahms specification of this scenario is given below (Stocker *et al.*, 2012):

```
workframe wf_checkMedicationC
{
  repeat: true;
  priority: 3;
  detectables:
    detectable takenMedicationC
    {
      when (whenever)
        detect ((Bob.hasMedicationC = false), dc:100)
        then abort;
    }
}
```

```

when (knownval (current.perceivedtime > 14) and
      knownval (Bob.hasMedicationC = true) and
      knownval (current.checkMedicationC = true))
do
{
  checkMedication ();
  remindMedicationC ();
  conclude ((current.checkMedicationC = false));
  conclude ((current.missedMedicationC = current.
            missedMedicationC+1));
}
}

```

A *workframe* corresponds to a plan in standard rational agent programming languages. This particular one has a certain priority level, events that can be detected, a set of triggers (captured within the *when* clause) and a list of outcomes (after they *do*) such as calling further workframes or updating beliefs.

3.4 Context dependency

3.4.1 Simple context view

In the shallow view of ‘context’ we can simply describe/specify contextual information in terms of changes in location. For example, we might describe a simple model of user movement by a straightforward state-machine (see Figure 11). This then allows us to use a simple formal language, such as *temporal logic* to describe changes in context, such as ‘ \diamond SHOPPING’, meaning ‘eventually the user will be in a SHOPPING context’.

3.4.2 Complex context view

In the case where we have the much more complex view of context, then (because each context is itself an agent) we must embed this notion within an *agent programming language*. The formal semantics for such a language, a variety of AGENTSPEAK, is provided in Dennis *et al.* (2008). We can see from the simple ‘cookery’ example earlier that it is straightforward to view contexts as agents and so develop contexts that share full agent-level behaviours amongst their contents. For example, when the cook enters the chef context (see Figure 5), the sending of plans is triggered, and the cook agent receives plans for preparing risotto, steak and pizza. This behaviour is written in the agent programming language as the following plan (Dennis *et al.*, 2008):

```

+ content (Agent) [source(self)]
  <- .print ("Sending ", Agent, " plans...");
  .send (Agent, tellHow, "+!cook (risotto)
        : check_constraint (cook,risotto)
        <- make (risotto).");

```

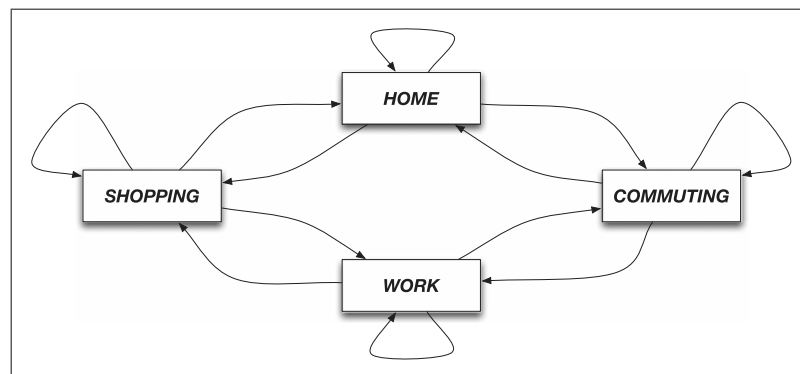


Figure 11 A simple model of user movement

```

.send (Agent, tellHow, "+!cook (steak)
      : check_constraint (cook, steak)
      <- make (steak).");
.send (Agent, tellHow, "+!cook (pizza)
      : check_constraint (cook, pizza)
      <- make (pizza).").

```

The plan is triggered, in the Context agent, by a new Agent entering it. Once this happens, the Context agent prints a statement then sends three messages to the Agent that just entered it. These messages contain the new plans that the Agent can then add to its own capabilities.

4 Verification

As described above, in order to carry out formal verification of pervasive systems, we will actually verify our four distinct dimensions separately (see Figure 12). So, let us examine all these four in more detail. As we will see below, all are based on model checking (Clarke *et al.*, 1999), but in very different ways.

4.1 Verifying combined specifications

Formal description of pervasive and ubiquitous systems can involve many different logical dimensions. This makes the verification task extremely difficult, as the interactions between these dimensions increase complexity significantly. The techniques currently used for formalizing and verifying aspects of pervasive and ubiquitous systems are very limited, often trying to describe system behaviour by considering just one dimension rather than dealing with multiple formalisms at the same time.

One idea to tackle the problem would be to introduce complex logics comprising all the aspects needed and develop new model checking methods and tools. However, devising new model checking methods/tools for such complex logics is both difficult and time consuming.

In Konur *et al.* (2013), we present a verification technique that allows the combination of model checking techniques for different formal frameworks. Namely, instead of introducing new logics for different combined aspects and developing model checkers for these logics, we combine existing logics and devise a *generic* model checking method for these combined frameworks. A combined model checking procedure can then be synthesized, in a modular way, from model checkers for constituent logics. This avoids the re-implementation of model checking procedures.

An overview of our approach is presented in Figure 13. As the figure illustrates, a combined model checker $\mathbf{MC}_{A\oplus B}$ for the combined logic $A\oplus B$ employs the model checkers \mathbf{MC}_A and \mathbf{MC}_B for the constituent logics A and B, respectively. Note that in the generic model checking procedure there are several calls to \mathbf{MC}_A and \mathbf{MC}_B , but they are not shown in the figure to keep the presentation simple.

This technique can be used to verify combined real-time, probabilistic, informational aspects, etc., which are at the core of pervasive and ubiquitous systems. Using this method, we can model check the combinations of the logics discussed in Section 3.1: *temporal logics*, *logics of knowledge*, *logics of beliefs*, *real-time temporal logics* and *probabilistic logics*.

4.2 Verifying autonomous behaviour

An advantage of using the BDI paradigm to program and model autonomous decision making is that several BDI programming languages now have associated formal verification techniques (Bordini *et al.*, 2006; Fisher *et al.*, 2007; Jongmans *et al.*, 2010; Dennis *et al.*, 2012). The property specification languages within such model checkers have to be extended in order to describe the key reasoning components (beliefs, desires, etc.) but as these environments generally do not implement full logics of knowledge or belief, such modalities can be treated in a simple and efficient fashion, normally as propositions within a temporal logic.

The points where an autonomous decision-making layer must interact with lower-level control layers or the real world are more of a challenge and involve the construction of abstract models of these parts of

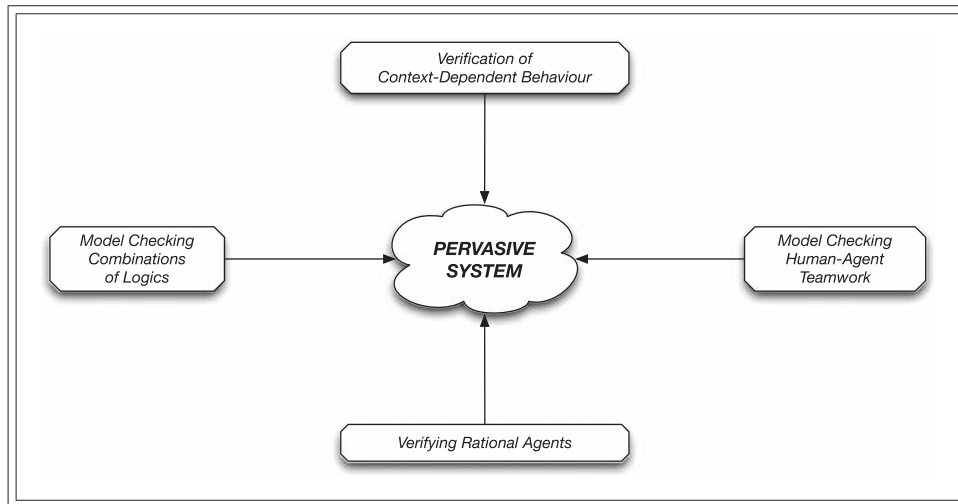


Figure 12 Four orthogonal aspects for *verifying* pervasive systems

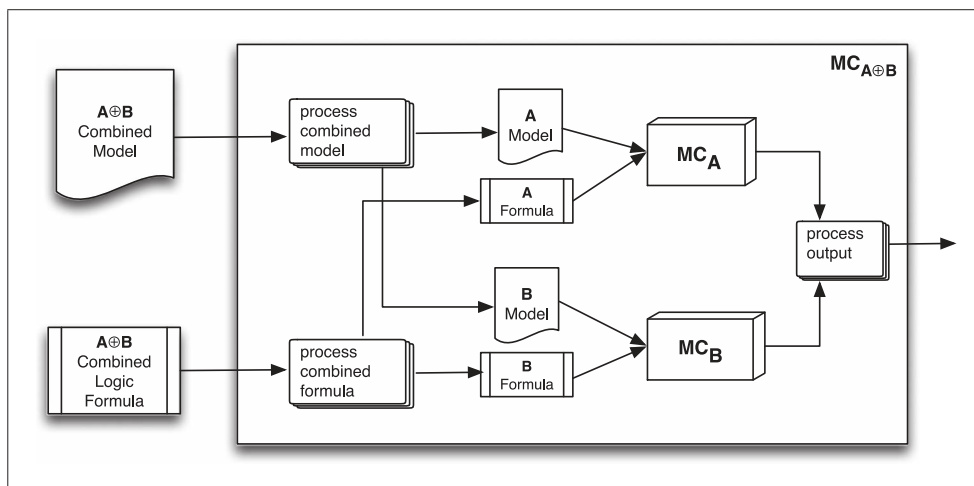


Figure 13 Overview of the combined model checking approach of Konur *et al.* (2013)

the system. There are two approaches to this problem: one is to use advanced mathematical techniques such as hybrid automata to model the real world and the other is to use simple, highly abstract models, and focus on capturing, clearly, the conditions under which the rational layer makes valid decisions.

The AIL/AJPF model checking toolkit (Dennis *et al.*, 2012), which we have developed, not only takes the second of these approaches but, through the generic *Agent Infrastructure Layer*, provides a Java framework to enable the implementation of interpreters for various agent programming languages and the direct use of the AJPF agent program model checker. Through this AIL/AJPF approach (see Figure 14) verification for several high-level agent programming languages has already been carried out.

Here, operational semantics can be provided for a range of agent programming languages, including AGENTSPEAK and METATEM, and then these semantics can be used within the generic model checking procedure provided by Java PathFinder (Visser *et al.*, 2000).

4.3 Verifying human–agent teams

As we have seen from Section 3.3, the *Brahms* simulation/modelling language is not only appropriate for describing human–agent teamwork, but has been used in practice for non-trivial human–robot interactions

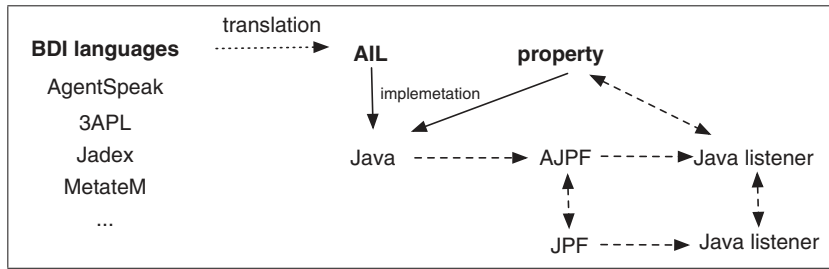


Figure 14 Verification of multiple agent programming languages (Bordini *et al.*, 2008)

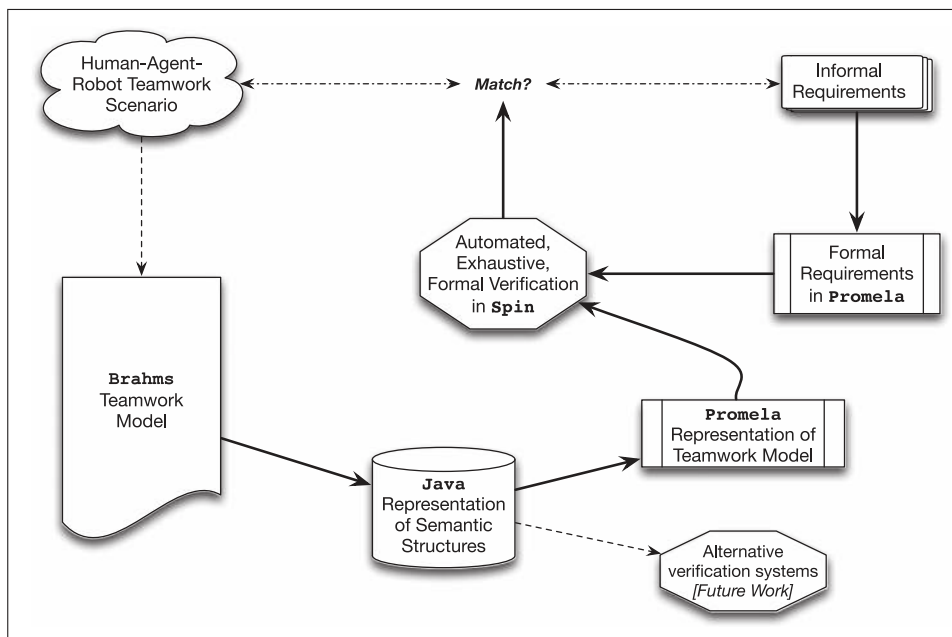


Figure 15 Formal verification process of Brahms (Stocker *et al.*, 2012)

(van Hoof, 2000; Sierhuis, 2001; Sierhuis & Clancey, 2002). Consequently, Brahms provides us with a ready-made framework for describing human-agent teamwork, as well as providing a range of examples in specific areas (Clancey *et al.*, 2003; Sierhuis *et al.*, 2003; Sierhuis, 2006). It seems natural, therefore, to aim towards the formal verification of Brahms models. The formal operational semantics for the Brahms language produced in Stocker *et al.* (2011) now forms the basis of our formal verification of Brahms (see Figure 15). In Stocker *et al.* (2012) a Brahms model of a human-agent team scenario is translated to a Java representation of the corresponding semantics structure, which is then translated into the SPIN model checker’s input language, Promela (Holzmann, 2003). Formal properties are also represented in Promela, and the verification problem is passed on to the SPIN tool. This approach has been extended and refined, particularly improving efficiency by translating to alternative model checkers (Hunter *et al.*, 2013).

4.4 Verifying context-dependent behaviour

4.4.1 Simple context view

In this view, the context information is obtained by getting the location or the activity. For example, in the Scatterbox system, the user context is obtained both by tracking the user’s location (via location sensors) and by monitoring his/her daily schedule. The Scatterbox system analyzes this context data, and determines the user’s context and corresponding level of *interruptibility*. However, the user’s actual context

might be different than the context derived by Scatterbox. Therefore, the exact interruptibility level, determined by the exact user movement, and Scatterbox's *perceived* user interruptibility might also be different. We therefore distinguish between the *exact* interruptibility and the Scatterbox's *belief* about the interruptibility, as these two might have different values. The system's belief about the interruptibility depends on the accuracy of the context acquisition process.

The Scatterbox system also analyzes incoming e-mails to determine the *importance* level. Similar to the interruptibility case, the exact importance level and Scatterbox's perceived importance level might be different. We therefore distinguish between these two variables. For the same reason, we distinguish the exact calendar information and the perceived calendar information.

When we verify the overall operation of the Scatterbox system, we usually refer to the perceived information by Scatterbox. For example, if we want to verify that the Scatterbox accurately finds the interruptibility of the user, we need to check whether the formula $\Box (B_intr = intr)$ is true over the model describing the behaviour of the Scatterbox system (*intr* and *B_intr* denote the exact interruptibility and perceived interruptibility, respectively).

In Konur *et al.* (2014), we formally verified the Scatterbox system using probabilistic model checker PRISM (PRISM, 2013). As the user location is determined via location sensors, a quantitative understanding of the effects of uncertain sensor input and inference was an essential part of this analysis. We also analyzed the dynamic temporal and stochastic behaviour of the system, allowing us to verify formal requirements even in the presence of uncertainty in sensors.

4.4.2 Complex context view

As discussed in Section 3.4, context can be modelled in BDI-style agent programming languages in a variety of ways—it can be subsumed into the belief system, used as part of a modular construction of agent systems or modelled separately as a first-class component of the agent state (Dennis *et al.*, 2008; Fisher & Hepple, 2009). All of these approaches can be incorporated easily into the operational semantics of the BDI languages and thus agent model checkers extended in a straightforward fashion to reason about them.

As described above, the AIL/AJPF verification framework can, in principle, verify both AGENTSPEAK and METATEM programs. The latter already has built in capabilities to reason about languages involving Content and Context.

5 Concluding remarks

There have been some approaches to the formalization of pervasive systems (e.g. Birkedal *et al.*, 2006; Ranganathan & Campbell, 2008; Arapinis *et al.*, 2009; Bakhouya *et al.*, 2012). In addition, significant aspects of the overall verification problem have already been tackled, for example, *probabilistic* verification (Kwiatkowska *et al.*, 2002), analysis of *context* (Boytsov & Zaslavsky, 2011), stochastic resource analysis (Gorrieri *et al.*, 2002) and the verification of *autonomy* (Dennis *et al.*, 2012). However, few of these have considered both specification *and* verification, and none have incorporated all the different aspects we describe here.

In this article, we have identified four orthogonal views and have tackled the verification of each of these independently, and each to a different extent. Below we give a summary of our work carried out so far:

1. *combinations of logics*: theoretical work on algorithms for model checking combined logics (Konur *et al.*, 2013);
2. *autonomous decision making*: characterizing rational agents, and then developing agent model checking (Dennis *et al.*, 2012) for analyzing autonomous decision making (Fisher *et al.*, 2013);
3. *human-agent teamwork*: developing a verification method for the Brahm workflow language (Stocker *et al.*, 2011, 2012) and its application to human-robot teamwork; and
4. *Context-based computation*:
 - a. probabilistic verification of context-based systems (Konur *et al.*, 2014); and
 - b. extending agent model checking to agent languages with context (Dennis *et al.*, 2008; Fisher & Hepple, 2009).

At present the verification of all these aspects is carried out separately; in the future we will work on integrating all these efforts into a holistic verification framework for pervasive systems. This is a difficult problem: combining techniques for populations and individuals, discrete logic and feedback control, abstracted human behaviour and stochastic activity, etc. requires a sophisticated framework and it may well be that the complexity of this approach requires simplification and limitations on potential interactions.

Acknowledgement

The work described in this paper is partially supported in the United Kingdom by the following EPSRC projects: ‘Verifying Interoperability Requirements in Pervasive Systems’ (EP/F033567) and ‘ROADBLOCK’ (EP/I031812).

References

- Arapinis, M., Calder, M., Dennis, L. A., Fisher, M., Gray, P. D., Konur, S., Miller, A., Ritter, E., Ryan, M., Schewe, S., Unsworth, C. & Yasmin, R. 2009. Towards the verification of pervasive systems. *ECEASST Electronic Communications of the EASST* **22**, 1–15.
- Bakhouya, M., Campbell, R., Coronato, A., de Pietro, G. & Ranganathan, A. 2012. Introduction to special section on formal methods in pervasive computing. *ACM Transactions on Autonomous and Adaptive Systems* **6**(1–6), 9.
- Birkedal, L., Bundgaard, M., Damgaard, T. C., Debois, S., Elsborg, E., Glenstrup, A. J., Hildebr, T., Milner, R. & Niss, H. 2006. Bigraphical programming languages for pervasive computing. In *Proceedings of the International Workshop on Combining Theory and Systems Building in Pervasive Computing*, 653–658.
- Bordini, R. H., Dastani, M., Dix, J. & El Fallah Seghrouchni, A. (eds) 2005. *Multi-Agent Programming: Languages, Platforms and Applications*, Springer.
- Bordini, R. H., Dastani, M., Dix, J. & El Fallah-Seghrouchni, A. (eds) 2009. *Multi-Agent Programming: Languages, Tools and Applications*, Springer.
- Bordini, R. H., Dennis, L. A., Farwer, B. & Fisher, M. 2008. Automated verification of multi-agent programs. In *Proceedings of the 23rd IEEE/ACM International Conference on Automated Software Engineering (ASE)*, 69–78.
- Bordini, R. H., Fisher, M., Visser, W. & Wooldridge, M. 2006. Verifying multi-agent programs by model checking. *Journal of Autonomous Agents and Multi-Agent Systems* **12**(2), 239–256.
- Bordini, R. H., Hübner, J. F. & Vieira, R. 2005. Jason and the Golden Fleece of agent-oriented programming, chapter 1. In Bordini, R. H., Dastani, M., Dix, J. & Seghrouchni, E. F (eds), 3–37.
- Boytsov, A. & Zaslavsky, A. 2011. *Formal Verification of the Context Model — Enhanced Context Spaces Theory Approach*. Technical report, Department of Computer Science, Space and Electrical Engineering, SE-971 87, Lulea University of Technology.
- Chen, H., Finin, T. & Joshi, A. 2003. An ontology for context-aware pervasive computing environments. *Knowledge Engineering Review*. Special Issue on Ontologies for Distributed Systems, **18**, 197–207.
- Clancey, W., Sierhuis, M., Kaskiris, C. & van Hoof, R. 2003. Advantages of Brahms for specifying and implementing a multiagent human-robotic exploration system. In *Proceedings of the Sixteenth International Florida Artificial Intelligence Research Society Conference (FLAIRS)*, 7–11. AAAI Press.
- Clarke, E. M., Grumberg, O. & Peled, D. A. 1999. *Model Checking*, MIT Press.
- Dastani, M., van Riemsdijk, M. B. & Meyer, J.-J. C. 2005. Programming multi-agent systems in 3APL, chapter 2. In *Multi-Agent Programming*, Bordini, R. H., Dastani, M., Dix, J. & Seghrouchni, E. F (eds), Springer, 39–67.
- Dennis, L. A. & Farwer, B. 2008. Gwendolen: a BDI language for verifiable agents. In *Proceedings of the AISB’08 Workshop on Logic and the Simulation of Interaction and Reasoning*, Löwe, B. (ed.). AISB.
- Dennis, L. A., Fisher, M. & Hepple, A. 2008. Language constructs for multi-agent programming. In *Proceedings of the 8th Workshop on Computational Logic in Multi-Agent Systems (CLIMA)*, LNAI **5056**, 137–156. Springer.
- Dennis, L. A., Fisher, M., Webster, M. & Bordini, R. H. 2012. Model checking agent programming languages. *Automated Software Engineering* **19**(1), 5–63.
- Dobson, S., Denazis, S., Fernández, A., Gati, D., Gelenbe, E., Massacci, F., Nixon, P., Saffre, F., Schmidt, N. & Zambonelli, F. 2006. A survey of autonomic communications. *ACM Transactions on Autonomous and Adaptive Systems* **1**(2), 223–259.
- Dobson, S., Sterritt, R., Nixon, P. & Hinchey, M. 2010. Fulfilling the vision of autonomic computing. *IEEE Computer* **43**(1), 35–41.
- Finger, M. & Gabbay, D. M. 1996. Combining temporal logic systems. *Notre Dame Journal of Formal Logic* **37**(2), 204–232.
- Fisher, M., Dennis, L. & Webster, M. 2013. Verifying autonomous systems. *Communications of the ACM* **56**(9), 84–93.

- Fisher, M. & Hepple, A. 2009. Executing logical agent specifications. In *Multi-Agent Programming: Languages, Tools and Applications* R. H. Bordini, M. Dastani, J. Dix & A. El Fallah-Seghrouchni (eds), 1–27 Springer.
- Fisher, M., Singh, M., Spears, D. & Wooldridge, M. 2007. Guest editorial: logic-based agent verification. *Journal of Applied Logic* **5**(2), 193–195.
- Gabbay, D., Kurucz, A., Wolter, F. & Zakharyashev, M. 2003. *Many-Dimensional Modal Logics: Theory and Applications*. Studies in Logic and the Foundations of Mathematics, **148** Elsevier Science.
- Georgeff, M. P. & Lansky, A. L. 1987. Reactive reasoning and planning. In *Proceedings of the 6th National Conference on Artificial Intelligence (AAAI)*, 677–682. AAAI Press, .
- Gorrieri, R., Herzog, U. & Hillston, J. 2002. Unified specification and performance evaluation using stochastic process algebras. *Performance Evaluation* **50**(2/3), 79–82.
- Greenfield, A. 2006. *Everyware*, New Riders Publishing.
- Henricksen, K. & Indulska, J. 2004. A software engineering framework for context-aware pervasive computing. In *Proceedings 2nd IEEE Conference on Pervasive Computing and Communications*, 77–86.
- Hepple, A. 2010. *Agents, Context, and Logic*. PhD thesis, Department of Computer Science, University of Liverpool.
- Hepple, A., Dennis, L. A. & Fisher, M. 2008. A common basis for agent organisations in BDI languages. In *Proceedings of the International Workshop on Languages, Methodologies and Development Tools for Multi-Agent Systems (LADS)*, Lecture Notes in Artificial Intelligence **5118**, 171–188. Springer-Verlag.
- Hindriks, K., de Boer, F., van der Hoek, W. & Meyer, J.-J. 2001. Agent programming with declarative goals. In *Intelligent Agents VII*, Lecture Notes in Artificial Intelligence **1986**, 228–243. Springer-Verlag.
- Holzmann, G. 2003. *Spin Model Checker, The: Primer and Reference Manual*, 1st edition. Addison-Wesley Professional.
- Hunter, J., Raimondi, F., Rungta, N. & Stocker, R. 2013. A synergistic and extensible framework for multi-agent system verification. In *Proceedings of International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 869–876. IFAAMAS.
- Isern, D., Moreno, A., Sanchez, D., Hajnal, A., Pedone, G. & Varga, L. 2011. Agent-based execution of personalised home care treatments. *Applied Intelligence* **34**, 155–180.
- Jongmans, S.-S. T. Q., Hindriks, K. V. & van Riemsdijk, M. B. 2010. Model checking agent programs by using the program interpreter. In *Proceedings of the 11th International Workshop on Computational Logic in Multi-Agent Systems (CLIMA)*, LNCS **6245**, 219–237. Springer.
- Knox, S., Shannon, R., Coyle, L., Clear, A., Dobson, S., Quigley, A. & Nixon, P. 2008. Scatterbox: context-aware message management. *Revue d'Intelligence Artificielle* **22**(5), 549–568.
- Konur, S. 2013. A survey on temporal logics for specifying and verifying real-time systems. *Frontiers of Computer Science* **7**(3), 370–403.
- Konur, S., Fisher, M., Dobson, S. & Knox, S. 2014. Formal verification of a pervasive messaging system. *Formal Aspects of Computing* **26**(4), 677–694.
- Konur, S., Fisher, M. & Schewe, S. 2013. Combined model checking for temporal, probabilistic, and real-time logics. *Theoretical Computer Science* **503**, 61–88.
- Kwiatkowska, M., Norman, G. & Parker, D. 2002. PRISM: probabilistic symbolic model checker, Lecture Notes in Computer Science **2324**, 200–204. Springer.
- NASA (Astronaut-Robot-Team-Concept). NASA astronaut robot partner. <http://history.nasa.gov/DPT/DPT.htm>.
- PRISM 2013. Probabilistic symbolic model checker. <http://www.cs.bham.ac.uk/dxp/prism>.
- Ranganathan, A. & Campbell, R. H. 2008. Provably correct pervasive computing environments. In *IEEE International Conference on Pervasive Computing and Communications*, 160–169.
- Rao, A. S. 1996. AgentSpeak(L): BDI agents speak out in a logical computable language. In *Proceedings of the 7th European Workshop on Modelling Autonomous Agents in a Multi-Agent World, (MAAMAW)*, LNAI **1038**, 748–752. Springer-Verlag.
- Rao, A. S. & Georgeff, M. P. 1991. Modeling agents within a BDI-architecture. In *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning (KR)*. Morgan Kaufmann.
- Rao, A. S. & Georgeff, M. P. 1992. An abstract architecture for rational agents. In *Proceedings of the International Conference on Knowledge Representation and Reasoning (KR)*, 439–449.
- Rao, A. S. & Georgeff, M. P. 1995. BDI agents: from theory to practice. In *Proceedings of the 1st International Conference on Multi-Agent Systems (ICMAS)*, 312–319. IEEE Press.
- RoboSafe 2013. Trustworthy robotic assistants project. <http://www.robosafe.org>.
- Rosa, P. M. P., Dias, J. A., Lopes, I. M. C., Rodrigues, J. J. P. C. & Lin, K. 2012. An ubiquitous mobile multimedia system for events agenda. In *WCNC*, 2103–2107.
- Sierhuis, M. 2001. *Modeling and Simulating Work Practice. BRAHMS: A Multiagent Modeling and Simulation Language for Work System Analysis and Design*. PhD thesis, SIKS Dissertation Series No. 2001-10, Social Science and Informatics (SWI), University of Amsterdam.
- Sierhuis, M. 2006. Multiagent modeling and simulation in human-robot mission operations. <http://ic.arc.nasa.gov/ic/publications>).

- Sierhuis, M., Bradshaw, J. M., Acquisti, A., Hoof, R. V., Jeffers, R. & Uszok, A. 2003. Human-agent teamwork and adjustable autonomy in practice. In *Proceedings of the 7th International Symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS)*.
- Sierhuis, M. & Clancey, W. J. 2002. Modeling and simulating work practice: a human-centered method for work systems design. *IEEE Intelligent Systems* **17**(5), 32–41.
- Stocker, R., Dennis, L., Dixon, C. & Fisher, M. 2012. Verifying Brahms human-robot teamwork models. In *Proceedings of the 13th European Conference on Logics in Artificial Intelligence, JELIA'12*, 385–397. Springer-Verlag.
- Stocker, R., Sierhuis, M., Dennis, L., Dixon, C. & Fisher, M. 2011. A formal semantics for Brahms. In *Proceedings of the 12th International Conference on Computational Logic in Multi-Agent Systems, CLIMA'11*, 259–274. Springer-Verlag.
- Strang, T. & Popien, C. L. 2004. A context modeling survey. In *Workshop on Advanced Context Modelling, Reasoning and Management, UbiComp 2004 – The Sixth International Conference on Ubiquitous Computing*.
- Turner, K. J. 2012. *Advances in Home Care Technologies: Results of the MATCH Project*, IOS Press.
- van Hoof, R. 2000. Brahms website. <http://www.agentisolutions.com>.
- Visser, W., Havelund, K., Brat, G. & Park, S. 2000. Model checking programs. In *Proceedings of the 15th International Conference on Automated Software Engineering (ASE)*, 3–12. IEEE Computer Society.
- Want, R., Pering, T., Borriello, G. & Farkas, K. 2002. Disappearing hardware. *Pervasive Computing* **1**(1), 36–47.
- Weiser, M. 1993. Some computer science issues in ubiquitous computing. *Communications of the ACM* **36**(7), 74–84.
- Wooldridge, M. 2002. *An Introduction to Multiagent Systems*, John Wiley & Sons.