On Lions, Impala, and Bigraphs: Modelling Interactions in Physical/Virtual Spaces

STEVE BENFORD, University of Nottingham
MUFFY CALDER, University of Glasgow
TOM RODDEN, University of Nottingham
MICHELE SEVEGNANI, University of Glasgow

While HCI has a long tradition of formally modelling task-based interactions with graphical user interfaces, there has been less progress in modelling emerging ubiquitous computing systems due in large part to their highly contextual nature and dependence on unreliable sensing systems. We present an exploration of modelling an example ubiquitous system, the Savannah game, using the mathematical formalism of bigraphs, which are based on a universal process algebra that encapsulates both dynamic and spatial behaviour of autonomous agents that interact and move among each other, or within each other. We establish a modelling approach based on four perspectives on ubiquitous systems—Computational, Physical, Human, and Technology—and explore how these interact with one another. We show how our model explains observed inconsistencies in user trials of Savannah, and then, how formal analysis reveals an incompleteness in design and guides extensions of the model and/or possible system re-design to resolve this.

1. INTRODUCTION

It is a long established idea in Computer Science that formal modelling can improve the design of computing systems in areas such as reliability, security, and usability. However, formal modelling is also notoriously difficult, especially when computing systems are interactive, as models must account for unpredictable human behaviour as well as complex computation.

The premise of our article is that formal modelling becomes even more challenging—but also necessary—with the advent of ubiquitous computing in which interaction "breaks away from the desktop" to become embedded in the everyday world. Ubiquitous computing systems are now coming into widespread use and are controlling many different aspects of our lives, from our homes to safety-critical infrastructures such as transportation, making it important to be able to model formally this emerging class of system. However, this may prove to be especially challenging. Not only are ubiquitous systems technologically complex and they involve humans, but they are also contextual, meaning that they sense and adapt to aspects of the surrounding environment. This brings into play a further set of unpredictable factors that need to be modelled, from environmental conditions such as weather, to the presence of crowds, to the uncertainties inherent in sensing technologies such as the Global Positioning System (GPS).

In this article, we reveal how modelling ubiquitous computing systems requires a socio-technical approach in which key aspects of human behaviour are modelled alongside system behaviour. We explore a new approach to modelling formally ubiquitous systems in which user interaction is closely coupled with the ability to sense human physical activity within a given context. Our ultimate aim is to enable the modelling of a wide variety of ubiquitous systems from today's location-based services to future experiences that incorporate environmental and physiological sensing and that adapt to complex user behaviours. Our first step towards this has been to develop a formal mathematical model of a specific example of a ubiquitous system called the Savannah game [Benford et al. 2005], studies of which have previously been reported in the ubiquitous computing literature. In turn, this has involved extending the formalism of bigraphical reactive systems (BRSs) [Milner 2009; Sevegnani and Calder 2015], to be able to model and reason about contextual interactions.

Although model-based approaches to interaction have a long tradition within HCI, there has been noticeably less emphasis on formal approaches over the last few years. Early work on HCI was strongly influenced by work in Psychology on mental models [Gentner and Stevens 1983]. A broad range of notations and formalisms were exploited to explore how to most effectively model a user's mental model of a system [Norman 1983]. Designing usable systems centred around the expression and use of these models. A key feature of these early approaches to design was to align the user's mental models with the models and representations used in the constructed computer system. The dominant framing of interaction was as a dialogue between the user and the computer system and this was often viewed from the perspective of the individual and driven by a broad specification/interpretation cycle [Moran 1981; Nielsen 1986].

The emphasis on dialogue between user and machine resulted in a broad set of plan-based theories and representations. Norman's seminal work on interaction framed interaction in terms of a broad goal-plan-execute-evaluate cycle [Norman 1988]. Consequently, formalisms were used to model a range of applications that closely coupled user interaction with display effects, including seminal work on display-oriented text editors [Sufrin 1982]. The emphasis on dialogue also allowed computer systems to be modelled as black box, with an emphasis on the input and output [Dix and Runciman 1985]. Considerable emphasis emerged on dialogue-based approaches, often exploiting state transition networks to represent items such as menu selection. Researchers exploited formalisms such as CSP [Hoare 1978] to represent the dialogue between a user and interface [Alexander 1987]. These formalisms allowed an exploration of potentially problematic issues such as undoing actions in human computer interfaces [Abowd and Dix 1992], and design and analysis of interfaces often involved the use of finite-state machines and event algebras [Gow and Thimbleby 2005].

As interactive systems have developed, the focus on interaction has moved beyond an emphasis of the dialogue between a user and a graphical interface to consider a wider set of issues, including debates around the nature of plans and goals in interaction design [Suchman 1987]. As HCI has matured, and the focus on interaction has moved, the disciplines involved have broadened beyond Computing and Psychology and the phenomena under exploration has widened [Grudin 1990]. A number of researchers have responded by changing the focus of formalism, reframing interaction around distributed cognition [Hutchins 1995] to consider modelling the resources involved [Wright et al. 1996, 2000] and researchers have used formal systems such as PVS [Owre et al. 1996] to model distributed cognition systems [Masci et al. 2015]. However, formal approaches have tended to move away from the mainstream of HCI, finding a strong role in niches such as safety critical interface systems, e.g., the number entry systems found on medical devices [Thimbleby 2015]. Readers are referred to Bolton et al. [2013], which provides a broad overview of the formal approaches to modelling interaction with automation.

In this article, we wish to rebalance this shift away from formal modelling by exploring the role that formal modelling might have with the advent of ubiquitous computing and the growing use of sensor-driven interaction. Rather than focus on the detailed interactive dialogue with a device, we are interested in how multiple user interactions in a shared ubiquitous computing environment might best be understood and reasoned about. This has required us to move beyond a focus on an interface, or even a device, to consider interaction in terms of the *spaces* that people inhabit, their actions in those spaces, and their movements between spaces.

Savannah is a collaborative, location-based educational game in which groups of school children learn about animal behaviour by role-playing being "lions," hunting together on a virtual savannah that appears to be overlaid onto their school playing field. The children, tracked by GPS, move around the playing field in order to steer their lions across the virtual savannah. They must discover prey, and then, gather together in groups in order to launch successful attacks. Although superficially simple, Savannah is an ideal candidate for modelling because it has been fully implemented and also because a usability study revealed how players experienced severe difficulties as a result of unexpected interactions between the various technologies involved and human behaviour within this particular context. The key question that motivates our work is whether formal modelling could account for (and might potentially have predicted) these problems.

BRSs are rewriting systems based on a universal process algebra that deliberately encapsulates both dynamic and spatial behaviour. A BRS consists of a set of bigraphs that describe spatial and interaction or communication relationships alongside a set of bigraphical reaction rules that define how bigraphs can evolve over time. We have previously extended the basic formalism to BRSs with sharing [Sevegnani and Calder 2015; Calder and Sevegnani 2014], to accommodate spatial locations that can overlap. The appeal of this formalism is that it allows us to express directly how the spatial arrangement of entities might drive computational effects within the system, a useful starting point for the formal modelling of ubiquitous systems such as the Savannah game.

In the following, we present a detailed case study of using expanded BRS to model the Savannah game addressing its computational rules, the mapping of these onto the physical environment, significant player behaviours, and finally, key characteristics of GPS within its underlying technical infrastructure. We then use our model to explain the interactional inconsistencies that were observed when Savannah was deployed in practice. Developing this model has involved an intra-disciplinary dialogue between two sub-fields of Computer Science: Formal Computing and Human Computer Interaction

(HCI) that are normally quite separate in terms of their concerns and approaches. Through developing our model and reflecting on this dialogue, we are able to make the following contributions.

—We demonstrate that it is possible formally to model complex and contextual ubiquitous computing systems and that the resulting models are both concise and yet have the power to predict and explain interaction phenomena that arise due to complex interactions between people and technologies within a given context.
—We show that BRSs is a powerful formalism for modelling ubiquitous systems due to its dual treatment of space and relationships, while also extending BRS with new features such as modelling probabilistic events.
—We establish an overarching framework for modelling ubiquitous systems based on the systematic exploration of multiple perspectives (e.g., Computational World, Human, Physical World, and Technology) where the detailed modelling of each perspective is underpinned by concepts and theories drawn from other fields and disciplines.
—Finally, we reflect on the nature of intra-disciplinary collaboration between Formal Computing and HCI and especially how this is enabled by graphical representations of algebraic formalisms.

Our article is structured as follows. Beginning with our background and approach, we briefly review the Savannah game in Section 2 before introducing bigraphs in Section 3. We then present in Section 4 our overall modelling approach, introducing four key perspectives that we consider when modelling Savannah. Next, we gradually develop our formal model of the Savannah game by stepping through each of these perspectives in turn. In Section 5, we model the significant computational aspects of the game. In Section 6, we model those aspects that define how the game connects to the physical world in which it is played, focussing on the human players in the playing field and their connections to the virtual lions. In Section 7, we model the key behaviours of these players as they set about forming groups and coordinating attacks. Finally, in Section 8, we model the impact of the underlying technical infrastructure, especially the characteristics of GPS. Having established our model of the Savannah game, we then put it to work. In Section 9, we employ our model to explain key issues that arose from the prior user study of Savannah; this serves to verify the model and also to demonstrate its power to explain complex interactional phenomena. In Section 10, we show how formal reasoning about properties that express relationships between perspectives reveals an incompleteness in the design of the Savannah game. In Section 11, we introduce a probabilistic extension to our model to take account of the stochastic nature of GPS sensing and "drift." In Section 12, we consider the wider implications of our work, reflecting on the potential use of BRS to model complex ubiquitous systems and considering how our intra-disciplinary approach to modelling might be further expanded to a wider range of systems and contexts. Finally, in Section 13, we finish with some reflections on the nature of formal modelling within HCI.

Please note that we aim to address both HCI and Formal Computing audiences and so throughout we focus on the *contribution* of formality to the challenges of designing and implementing ubiquitous systems rather than on the details of the formalism itself. We, therefore, keep the formal machinery to a minimum, developing our models through examples and using the intuitive graphical representation of bigraphs, instead of the algebraic form. Two appendices contain summaries of all the rules of the model, in graphical and algebraic form, respectively.

2. OVERVIEW OF SAVANNAH GAME

Savannah was a prototype game developed by Futurelab, HP, the BBC, and the Universities of Nottingham and Bristol in order to explore how location-based mobile services

Fig. 1. Lionesses hunt in a virtual savannah overlaid on their playing field.

might enable new approaches to learning. The game was fundamentally spatial in nature, creating a virtual savannah out of digital sounds and images from the BBC Natural History Unit archive, and then, using GPS to overlay this onto an otherwise empty school playing field. Groups of six children at a time were then sent out into this virtual savannah to role play being a pride of lions, progressing through three distinct levels of a game in which they learned about different aspects of lion behaviour.

Each player was given a handheld computer that used the GPS to track their position on the playing field and WiFi to communicate this to the central game server that implemented the game rules (see Figure 1). The server also triggered the display of images, sounds, and messages and maintained a record of all movements and actions in the game. The virtual savannah was structured as discrete, bounded and irregularly shaped locales, as shown in the left part of Figure 2. As players traversed the playing field their GPS position was sent back to the game server, which then mapped this onto the coordinate system of the virtual savannah in order to determine when they entered particular locales. Entering a locale triggered the display of its associated images and sounds (the latter heard through a single-earpiece). Some locales contained potential prey that could be attacked in which case the players would see images and hear sounds representing the prey along with a button that they could use to launch an attack on this prey as shown in Figure 3. Players were told that they needed to work out how many lions were needed to attack each prey, gather together the requisitite number of players, and then, press the attack button at the same time. In practice, the system implemented a timeout of 10s after the first player pressed their attack button during which others could subsequently join in the attack by also pressing their buttons.

Players were initially asked to explore their terrain by systematically walking the playing field, experiencing the sounds and images of different regions of the savannah via their mobile devices and virtually marking key locations in order to stake out their territory. They were then allowed to hunt for prey by scouring the savannah for animals to attack. Players were instructed to choose carefully which animals to attack and in particular to consider how many lions would be needed to take part in a successful attack on each type of prey. Beyond this, the game rules simulated basic aspects of
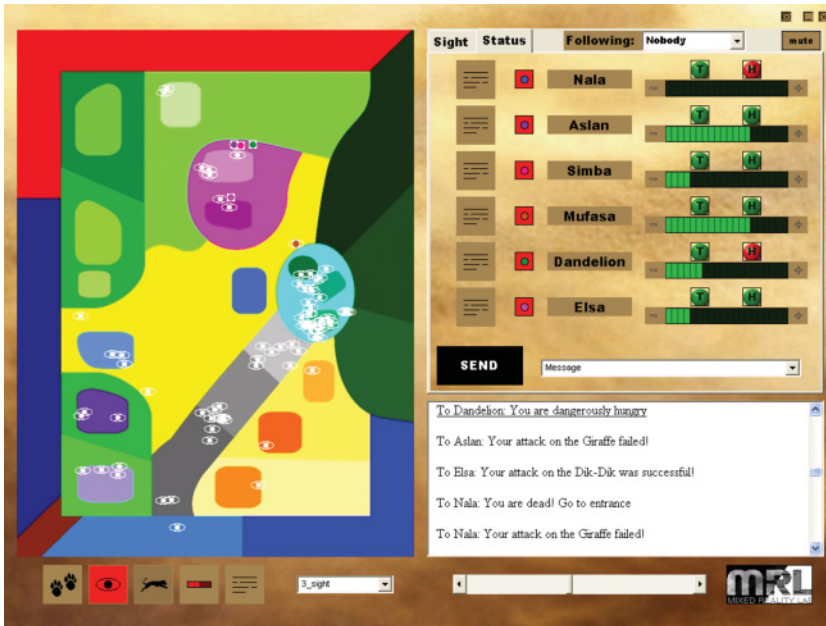
Fig. 2.   The replay interface.

Fig. 3.   Deciding when to attack potential prey.

health, including hunger, thirst, and potential damage from large and angry prey. The overall aim of the game was for the pride to survive and for the children to learn about natural history.

2.1. Understanding the Savannah Experience

A user study of Savannah involved the game being played by six groups of school children over three days of trials. The ages of the children ranged from 9 to 12 years

old and each group contained three girls and three boys. In order to be able to analyse player behaviour, we captured two video views of the action in the field for each session, each following a nominated child throughout the level. However, because the children tended to work together in gender groups, adopting a strategy of following just one girl and one boy in each group enabled us to capture much of the action in the field using only the two cameras. The two nominated children also wore radio microphones enabling us to record their conversations with nearby colleagues.

However, video and audio recordings alone proved insufficient to fully understand the experience of Savannah. Whereas video and audio informed us of players' physical actions, movements, gestures, and conversations, they failed to convey a system-level view of events, which is necessary to fully understand what took place. Specifically, we could not directly observe when players entered different locales, what each player was seeing or hearing on their handheld computer, or what buttons they were pressing. We, therefore, developed an interface to replay system recordings of interactions with the game. Our analysis proceeded by replaying these system recordings synchronised with the two video recordings, swapping between them in order to construct a detailed picture of players' physical activities in tandem with system interactions and responses.

Our study revealed how players experienced major interactional difficulties in playing the game. A first pass through our data identified many moments where players appeared to be experiencing serious difficulties and these were then transcribed and analysed in fine detail in order to reveal the underlying causes. A detailed analysis is covered in Benford et al. [2005], while we present and model one illustrative example in Section 9.

For the time being, it is sufficient to note that the problems encountered were significant and frequent and that that they appear to have arisen from a complex interplay of several different factors spanning the design of the game's rules, especially its use of discrete locales; the mapping of these onto the physical environment of the school playing field; human behaviour concerned with forming temporary groups and a tendency to stop at the boundaries of locales on encountering prey; and a degree of uncertainty and instability in the underlying technical infrastructure of GPS and WiFi.

These problems transpired to be somewhat general difficulties arising from the overall design of Savannah and especially from its mapping of virtual onto physical space and assumptions made about player behaviours (rather than arising from faults in this particular implementation). Moreover, using GPS to map discretely bounded locales onto physical environments in this way is a common approach within location-driven ubiquitous systems leading us to anticipate that similar problems might arise in many other applications. Such problems and their underlying causes, therefore, need to be accounted for in any formal model of the system. In turn, a sufficiently rich model that provides such an account might help designers explain and possibly even predict such problems.

The inherently spatial nature of Savannah, in which system events are triggered by changes in the spatial arrangements of players in relation to each other, virtual locales and physical space, lies at the heart of modelling the game and the problems that arose. This appeared to be a good match with the bigraph approach that directly provides formalisms for modelling spatial structures alongside other relationships.

3. INTRODUCTION TO BIGRAPHS

BRSs [Milner 2009] are a universal formalism for modelling interacting systems that evolve in time and space. The formalism is based on a graphical model that emphasises both locality and connectivity. A bigraph consists of a *place graph* representing the location of entities in terms of containment and a *link graph* representing the interactions between entities. Bigraphs as originally defined by Milner do not permit shared or

(a) Graphical notation   (b) Place graph   (c) Link graph

Fig. 4.   Example bigraph *B*.

overlapping locations, and so Calder and Sevegnani have extended the original theory
to BRSs with sharing [Sevegnani and Calder 2015], and implemented a computational
rewriting framework (BigraphER).[1] In the remainder of this article by BRS, we denote
BRS with sharing; in Section 3.1, we give an example that illustrates the advantages
of the extension.

The particular appeal of BRS for us is that the formalism targets the modelling
of locative and ubiquitous systems that are inherently spatial. The approach gives a
particular emphasis to the spatial aspects allowing the location of entities to be used
to drive the dynamic of the system through a series of rewriting rules. An equally
important feature of bigraphs for us is that they strive to bridge between formal math-
ematical modelling and systems design by supporting equivalent diagrammatic and
algebraic representations. This allows systems designers to express graphically the
spatial arrangement of the systems under consideration and to use these graphical
forms as the principle modelling representation. An immediate appeal of this approach
for us is that it provides an accessible entry point for the HCI community to engage
with formal modelling and draws upon the tradition within this community on the use
of diagrams and sketches to design systems. We are interested in the extent to which
this approach might provide a *lingua franca* that will allow greater communication
between user experience designers and formal computer scientists.

In this section, we give an informal overview of BRS, with some examples. The
overview contains only sufficient detail for this article; a concise semantics is defined
in Sevegnani and Calder [2015]. Our introduction presents bigraphs using its diagram-
matic approach.

A BRS consists of a set of bigraphs that describe spatial and communication rela-
tionships, and a set of *reaction rules* that define how bigraphs can evolve over time.

A bigraph is a graphical structure defined as in Figure 4(a). Ovals, circles, and
rectangles encode *entities*, which can be real or virtual. They are assigned a type called
control indicated here by Room, Display, User, and Phone. In general, bigraphs permit
any kind of shape (sometimes coloured) for entities. In this article, we use an intuitive
visual representation (shapes and colours) for entities in the model.[2]

The spatial placement of the entities in a bigraph is defined by the containment rela-
tion and is expressed in the place graph as in Figure 4(b). Interactions and (non-spatial)
relationships between entities are represented by green edges called *links*. These rela-
tionships define the link graph of a bigraph. An example is in Figure 4(c). Links may
be only partially specified, in which case they connect a *name*. Names are links (or
potential links) to other bigraphs representing the external environment or context.
By convention, names are drawn above the bigraph. For example, in Figures 4(a), 5,
and 6, *n* is used to name the (potential) link. The number of links of an entity, also
called *arity*, depends on its control, i.e. entities with the same control have the same

---

[1]http://www.dcs.gla.ac.uk/∼michele/bigrapher.html.
[2]For example, lion-like shapes will be used to indicate entities of control lion (see Figure 9).

Fig. 5.   Example bigraphical reaction rule `leave`.

Fig. 6.   Example transition $B \longrightarrow_{\texttt{leave}} B'$.

number of links (a control is similar to a type). A dashed rectangle denotes a *region* of adjacent parts of the system.

The example bigraph B in Figure 4(a) models a WiFi enabled Display that is situated in the Room identified by name n. A User and her Phone are also in the same Room. The Phone is linked to its owner (i.e., the User) and it is wirelessly connected to the Display via the other link.

Grey squares (for example in Figure 5) indicate *sites*, which encode parts of the model that have been abstracted away. In other words, an entity containing a site can contain zero or more entities of any control. This is a particularly powerful feature in the definition of *reaction rules* that rewrite bigraphs to bigraphs and give dynamics to bigraphs. A reaction rule consists of a pair of bigraphs: the left-hand side specifies the parts of a bigraph to be changed, while the right-hand side specifies how those parts are changed. We use ➤ to indicate the definition of the reaction rules. Example reaction rule `leave` is drawn in Figure 5. It specifies the movement of a User and her Phone from inside the Room (left-hand side) to outside the Room (right-hand side). Note that the wireless link between the Display and the Phone is interrupted in the right-hand side. Other devices may be connected to the Display via name w.

The *reaction relation* on bigraphs induced by the reaction, i.e., the relation resulting from the (iterative) application of the reaction rules, is written with an arrow thus: $\longrightarrow$ or $\longrightarrow_{\texttt{rule}}$ to identify the reaction rule being applied to generate the transition. The transition in Figure 6 shows how bigraph B on the left-hand side evolves to become bigraph B' on the right-hand side. This is the result of applying reaction rule `leave` defined in Figure 5 to bigraph B in Figure 4(a).

3.1. Shared Locations

The models we are going to introduce in the remainder of this article are defined using our generalisation of the standard definition of bigraphs that allows for overlapping (or shared) spatial locations. We now show, through a simple example, the main features of this extension and its advantages over non-sharing bigraphs.

The example bigraph defined in Figure 7(a) contains two entities of control Room that are identified by names n and n'. Each one contains a WiFi enabled Display. In this richer model, we can model explicitly the signal range of each Display with the two azure ovals of control Sig. Note that the signals
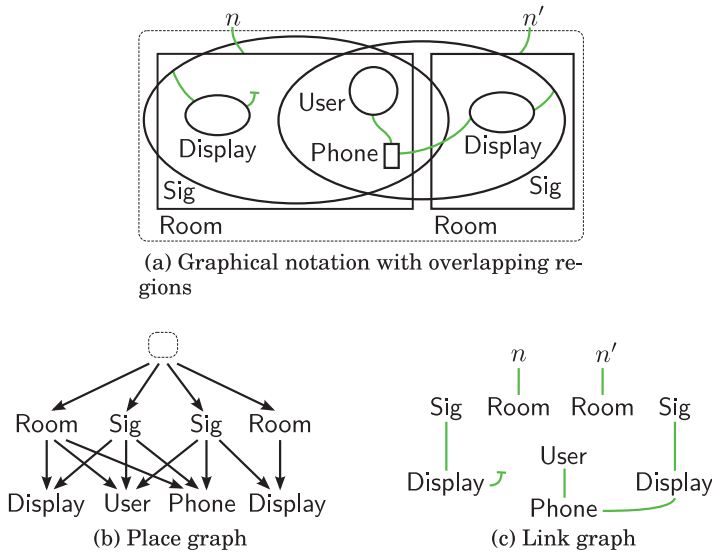
(a) Graphical notation with overlapping regions

(b) Place graph

(c) Link graph

Fig. 7.   Example bigraph with sharing.

—are not contained within the boundaries of the rooms and
—are overlapping.

The User and her Phone are in room n. However, they are in a location in which both WiFi signals can be sensed. This is represented by the fact that the Phone and the User are in the intersection of the two Sig entities. The Phone is connected to the Display in Room n' because it can sense its signal despite being in Room n. This difficulty of wireless video connection from neighbours' houses has been widely reported with the press, see, for example, http://gizmodo.com/007455/the-new-way-to-accidentally-spy-on-your-neighbors.

This system is difficult to model with standard bigraphs; in particular, it is difficult to represent a user being simultaneously within the wireless range of two distinct displays. As indicated in Figure 7(a), bigraphs with sharing allow for an immediate and intuitive definition of complex (e.g., overlapping) spatial configurations. In standard bigraphs, Figure 7(a) is simply not allowed (See Sevegnani and Calder [2015] for more detailed discussion of the benefits of bigraphs with sharing). Note that in bigraphs with sharing, spatial locations are represented by a Directed Acyclic Graph (DAG),[3] as indicated in the place graph in Figure 7(b). However, the definition of the link graph is unchanged, as indicated in Figure 7(c).

Having provided an informal overview of bigraphs, including how we might express dynamic effects using reaction rules, we now turn our attention to modelling the Savannah game using BRS. We begin with our overall approach to modelling.

4. MODELLING APPROACH

In developing a formal bigraphical model of the Savannah game, we were mindful of the complexities of what needed to be modelled. Although Savannah is a relatively simple computer game in terms of its rules, content, number of levels and so forth, the previous user study had revealed how significant difficulties arose from the interplay between the game design, its situation within the physical world, the behaviours of its

---

[3]A tree structure is sufficient in standard bigraphs. See Figure 4(b).

human players, and the ways in which these were mediated by the underlying technical infrastructure, most notably GPS. We, therefore, set about modelling Savannah in terms of four distinct but interlocking *perspectives*, each of which addresses a different facet of the overall system in depth.

This idea of considering different perspectives or viewpoints on system design is well-established within software engineering where it provides a way of breaking down the engineering of complex socio-technical systems into manageable portions that reflect different sensitivities and underlying technical concerns.

In our case, the perspective-driven approach evolved as we iteratively developed our model in response to the challenge of structuring what transpired to be a complex formal model into tractable parts that built upon one another.

While many perspectives might potentially be relevant to the modelling of a ubiquitous system, our reflections on the user study identified four that appeared to be critical to accounting for the design and experience of Savannah.

—The *Computational World* perspective in which we model the design of the application content and software, in our case the significant aspects of the structure and rules of the Savannah game such as the division of the virtual savannah into discrete locales and the rules that govern how lions group together and synchronise their actions in order to successfully attack prey.

—The *Physical World* perspective in which we extend our model to address key aspects of the physical environment within which the application was deployed. In our case, this means modelling the presence of human players on an open school playing field, the connections between individual players and specific lions in the game, and the ways in which players' movements control the locations of their respective lions within the virtual savannah.

—The *Human* perspective in which we then model further aspects of players' behaviours that significantly affect the experience. The use study of Savannah revealed that these were the ways in which players summoned their colleagues to form physically co-located groups before launching an attack as well as a natural tendency to stop when first encountering prey which would often place players on the edge of locales.

—The *Technology* perspective in which we model the impact of the underlying technologies specifically the impact of GPS uncertainty and variability.

In addition to choosing the perspectives that need to be modelled in a given case, it is also necessary to consider which aspects of these perspectives should be modelled. For example, it is not feasible to try to construct a complete model of human behaviour covering all aspects of perception, action, motivation, and so forth in detail. Rather, as the above points suggest, we need to focus on those aspects that are believed (or have been shown to be) most directly relevant to the challenge at hand. Moreover, as we shall see in the following sections, it can be useful to turn to established concepts or models within each perspective, in order to identify relevant abstractions that are grounded in the theory or empirical evidence from other disciplines. Some of the key aspects of the four perspectives for the Savannah game are summarised in Figure 8.

The following sections now step through each of these four perspectives in turn in order to gradually build up a formal bigraphical model of the Savannah game.

5. COMPUTATIONAL WORLD PERSPECTIVE

5.1. Bigraphical Model

We begin with the Computational World perspective that defines the virtual content and rules for the Savannah game. Our first step is to specify the static aspects of the

Fig. 8.   Four perspectives.

Table I. Controls for the Computational World

Control	Arity	Description	Graphical Notation
locale	1	Locale	
localeattack	1	Locale with ongoing attack	
lion	1	Lion	
lionattack	2	Lion initiating an attack	
liongroup	2	Lion in a group	
impala	1	Impala	
impalaseen	2	Impala seen by a lion	
impalaheld	2	Impala held by a group of lions	

bigraphical model. The locale map in the computational world is represented by entities of control locale and localeattack. They are indicated by circles and ovals, with solid and dashed border, respectively. Entities of control localeattack are used to represent locales in which an attack is taking place. The other entities in the Computational World are lions and impala. They are always contained within a locale. Lions and impala, in various states, are indicated by controls lion, lionattack, liongroup and impala, impalaseen, impalaheld, respectively. The controls are summarised in Table I.

A bigraphical representation of an example game configuration in the Computational World perspective is given in Figure 9. The locale on the left-hand side contains three

Fig. 9.　Example game configuration in the computational world.


(a)　　　　　　　　　　　　　　　　(b)

Fig. 10.　Reaction rules for adding a lion (a) and an impala (b) to the Computational World.


Fig. 11.　A lion initiates an attack on an impala in its locale (⟶), idle when the attack timer expires (⟵).

idle lions while the locale on the right-hand side contains one lion attacking an impala.

We now explain the dynamic aspects of the bigraphical model by introducing eight simple reaction rules.

The first two rules describe a lion and an impala appearing in the Computational World. They are specified by the diagrams given in Figure 10(a) and 10(b), respectively. A site in a locale indicates that other lions or impala may be present inside the locale.

Reaction rule ⟶ in Figure 11 models a lion initiating an attack when an impala is present in its locale. The left-hand side indicates an idle lion and impala (outline shapes). The right-hand side of the rule specifies that the lion is committed to killing this particular impala, thus the two entities are linked together. Different controls (i.e., lionattack and impalaseen) need to be used because arity 2 is required to accommodate the new link. They are indicated by solid yellow shapes. Additionally, the locale is assigned control localeattack, which is indicated by a dashed oval. This is required to model that only one attack can take place in a locale at a given moment. Note that once this reaction rule is applied on a locale, it cannot be applied again on the same locale because the left-hand side would not be of control locale but localeattack instead. Also in this case, the sites in the locales allow other impala or lions to be in the locale.

Reaction rule ⟵ in Figure 11 models the attack timer expiring. If no other lions join the attacking lion, then the attack is aborted and all the entities go back to the initial configuration with idle lion, impala, and locale.

The reaction rule in Figure 12 models the next phase of the game in which an idle lion joins the attack initiated by a lionattack in its locale. On the right-hand side, both the two lions have control liongroup (in blue) and the impala has control impalaheld (in red). Moreover, both lions are linked to the impala.

A similar reaction rule is defined in Figure 13. In this case, an idle lion joins the attack of a group of two lions in its locale. On the right-hand side, a group is created by linking the three lions (liongroup) to the impala (impalaheld).

The reaction rule in Figure 14 models the attack timer expiring before a third lion could join the group to complete the attack. On the left-hand side, a locale (localeattack) contains a group of two lions (liongroup) holding an impala (impalaheld). On the

Fig. 12.   A second lion joins the ongoing attack and both lions are now in a group.


Fig. 13.   A third lion joins the (attacking) group and the three lions are now in a group.


Fig. 14.   The timer expires before a third lion could join the group.

right-hand side, all the entities revert to their idle state similarly to reaction rule
◄— in Figure 11.

Finally, the reaction rule in Figure 15 models killing an impala (impalaheld) by a
group of three lions (liongroup). On the right-hand side, the three lions and the locale
return to their original, idle, state with control lion and locale, respectively, while the
impala disappears.

6. PHYSICAL WORLD PERSPECTIVE

We now turn to the Physical World perspective to model how the computational content
from the previous section is overlaid onto a physical school playing field. This includes
modelling how individual children become paired with virtual lions so that several
lions may then coordinate an attack on virtual prey.

6.1. Bigraphical Model

Four new controls are required to model the entities populating the Physical World.
Controls child, childattack, and childgroup model human players while control field rep-
resents the physical field of the Savannah game. The graphical notation for each control
is given in Table II. Since there is only one playing field in the physical world, we allow
only one entity of control field in the model.

Our bigraphical model allows for a clean separation between the Physical World and
the Computational World. This is shown by the fact that entities for human players
are in the region containing field while entities for locales, lions, and impala are in a
distinct region.

An example game configuration showing the Computational World and the Physical
World is given in Figure 16. The relation between a player and its lion in the computa-
tional world is represented by the links between child and lion entities. Note that there
is no such link for impala because they are purely computational entities, i.e., they
do not have a counterpart in the Physical World. Also observe that there are no links
between the players and the field (there are no interactions between them) and at this
point, there is no notion of locales in the Physical World: players roam in the field and

Fig. 15. A group of three lions kills an impala.

Table II. Controls for the Physical World

Control	Arity	Description	Graphical Notation
field	0	Field	◼
child	1	Child	
childattack	1	Child initiating an attack	
childgroup	1	Child in a group	

Fig. 16. Example game configuration in the physical world (top region) and computational world (bottom region).

a model of sensors is required in order to identify their position precisely with respect to the locale map. We will model this aspect of the Savannah game in Section 8.

The dynamic aspects of the model are described by seven reaction rules. Our first reaction rule in Figure 17 models a pair child/lion entering the game. The left-hand side consists of two regions: one for the Physical World, containing field and the other for the Computational World, containing locale. On the right-hand side, an idle child and an idle lion are in the field region and inside the locale, respectively. Moreover, child and lion are linked together. Observe that this reaction rule suffices to enforce a one-to-one correspondence between human players and lions.

The remaining reaction rules are obtained by introducing human players in the reaction rules for the Computational World described in the previous section. The reaction rules in Figure 18 correspond to the reaction rules defined in Figure 11. In the reaction rule (⟶), an idle player can initiate an attack only when its lion is also idle and it is in an idle locale with an idle impala. On the right-hand side, the rule specifies that whenever a player initiates an attack, its lion becomes of control lionattack, it is linked to an impalaseen and the locale becomes localeattack. The opposite reaction rule (⟵) models the attack timer expiring.

Fig. 17.   A pair child/lion enters the game.

Fig. 18.   `attack`: a child/lion pair initiates an attack on an impala in its locale (➡), `tmieout1`: idle when the attack timer expires (◀—).

Fig. 19.   `join1`: a child/lion pair joins the attack initiated by a lion in its locale.

The reaction rule in Figure 19 models a child/lion pair joining an attack initiated by a lionattack in its locale. It corresponds to the reaction rule defined in Figure 12. Note that in Figure 19, the players are in different regions. This is not strictly necessary since in our model all players are always in the region containing field. However, we can adopt this formulation because it does not *force* the player to be in distinct regions: one region can be matched by two distinct regions in the matching for bigraphs. Additionally, it will allow for a straightforward extension of the rules in the next section to include a model of the human perspective.

The reaction rule in Figure 20 models the attack timer expiring before a third lion manages to join the ongoing attack. As in the corresponding reaction rule defined in Figure 14, all entities on the right-hand side return to the idle state.

The reaction rule in Figure 21 models an idle child/lion pair joining a group of two lions to kill an impala. On the left-hand side, the child/lion pair is idle, while it is childgroup/liongroup on the right-hand side. Note that the players linked to the two lions already in the group, i.e., lions with names a' and a'' do not need to be included in the definition of the reaction rule. This is because their controls are unaffected by the application of the reaction rule. Moreover, they must both be of control childgroup since they were introduced by a previous application of the reaction rule in Figure 19.

Finally, the reaction rule in Figure 22 models a group of three child/lion pairs killing an impala.

A common feature of the reaction rules we described in this section is that the controls for a human player and the corresponding lion always follow three patterns: child/lion, childattack/lionattack, and childgroup/liongroup. In the graphical notation, this is shown by always assigning the same colour to a player and its lion. Another important feature is that the link connecting a player and its lion is created by an application of the reaction rule in Figure 17 when a new child/lion pair enters the game, and then, never broken or modified by the application of any other reaction rule. Finally, note that the definition of the reaction rules enforces the separation between the Physical World and the Computational World. This can be proved formally by showing that no

Fig. 20. `timeout2`: a group of two child/lion pairs becomes idle when the attack timer expires.

Fig. 21. `join2`: a child/lion pair joins a group in its locale.

Fig. 22. A group of children/lions kill their impala.

reaction rule can move players inside a locale and lions or impala to the field region (and we will consider such proofs in Section 10).

7. HUMAN PERSPECTIVE

The next element of the Savannah game to be modelled is the way in which players gather together into groups in order to launch attacks on prey. The user study revealed how players tended to exhibit a specific behaviour in this regard, one in which a player discovering prey would summon other players over to them by shouting; thus leading to a group of players forming a tight knit circular formation, looking inwards to compare their screens so that they could carefully coordinate the pressing of their "attack buttons." As we shall discuss later, we might potentially model many aspects of human behaviour: for example, intentions, likely patterns of behaviour, or cognition, motor skills and task models as people interact with a device.

The essential behaviour to capture here is that of gathering into a closely proximate group so as to launch an attack; for this we turn to the field of Proxemics, drawing inspiration from Hall [1966]. Proxemics has its roots in cultural anthropology where it emerged as a theory to explain the influence of personal distance and territory on human communication and has grown to be a rich and complex field in its own right. Proxemics-based approaches have grown in popularity in ubiquitous computing with a number of researchers exploring their use to understand and drive interaction in ubiquitous computing [Greenberg 2011].

For the purposes of this article, a simple interpretation of proxemics will suffice: the idea of an aura representing the *personal space* layer of proxemics (note, there are several such layers in the core proxemic model). Specifically, this means one aura per player. Further, for the purpose of modelling Savannah and based on the findings from the user study, we take the view that when several players' auras (i.e., personal spaces) overlap, the players believe that they are in a group and may be ready to make an attack.

Fig. 23.   A human player with his/her auras.

Fig. 24.   Example configuration with Physical World, Computational World and Human perspective.

Fig. 25.   A pair child/lion enters the game.

7.1. Bigraphical Model

We introduce a new control aura. Entities of this kind are associated with each player,
forming concentric bubbles around her/him as shown in Figure 23. Each aura repre-
sents a different degree of proximity (i.e., distance from the corresponding child), which
in turn determines what kind of interaction is possible between the child and other
entities within the same bubble. Entities of control aura are drawn as circles in shades
of yellow and can be overlapping. A link is used to associate each child to her/his aura.

An example configuration showing both the Physical World augmented with the
Human perspective and the Computational World is given in Figure 24.

Each entity denoting human players is contained by and linked to an entity of
control aura. Note that auras may be overlapping and the children are placed in their
intersection. This models the fact that two or more children are interacting. The link
between a player and her/his aura allows us to distinguish the player's own aura from
the other auras. Note that lions in different locales may be linked to players being in
overlapping auras. This models the (possible) discrepancy between the player's belief
of game configuration in the physical world and the actual configuration of the game
in the Computational World.

The dynamic aspects of the model for the Human perspective are specified by three
reaction rules.

The first reaction rule is given in Figure 25. It is a simple extension of the reaction
rule in Figure 17: when a child/lion pair enters the game, an aura is also associated to
the child in the Physical World. This reaction rule specifies that entities of control aura

Fig. 26. `aura-join`: overlapping auras are created ($\longrightarrow$), `aura-disjoin`: uncoupled ($\longleftarrow$).

may only appear in the region containing field and that children are always contained by an aura.

The other two reaction rules are new to the Human perspective and define how overlapping auras are created and uncoupled. The graphical representation is given in Figure 26.

The reaction rule from left to right ($\longrightarrow$) models two children moving closer in order to make social interactions possible. This is encoded by the two auras becoming overlapped on the right-hand side and the two children being placed in their intersection. The opposite reaction ($\longleftarrow$) models two children leaving a region where interactions were possible. The sites inside each aura and the regions surrounding each child are needed to allow the two reaction rules to be applied when an arbitrary number of children share their aura and not just when exactly two children interact. These reaction rules assume that auras are symmetric, i.e., whenever a child moves into another child's aura, then both children are in the intersection of the two auras. Note that the rules are defined only when the two children are idle. This is based on the assumption that children stand still during the phase of the game in which three lions form a group to kill an impala. The interaction specified by these two reaction rules is always binary. Therefore, when more than two players are involved, a sequence of reaction rule applications is required in order to obtain the correct configuration. For instance, if we want to form a group of three players A, B, and C, then the first application of ($\longrightarrow$) will create the intersection between players A and B, the second application will create the intersection between B and C, and finally, the third application will create the intersection between A, B, and C. Note that these applications can be performed in any order and the final configuration is always the same.

8. TECHNOLOGY PERSPECTIVE

We now turn to the last of our four perspectives, the Technology perspective. This is concerned with significant characteristics of the underlying technical infrastructure of devices, communications, and sensing that impact on the experience and so need to be captured formally in the model. While there might potentially be many relevant technologies and characteristics to consider, the Savannah user study highlighted one in particular: GPS mediates the connection between the players and their associated lions. GPS inaccuracies and variability were highlighted as being contributing factors to players' problems.

8.1. Bigraphical Model

Two new families of controls are introduced to model the entities characterising the Technology perspective: GPS positions and areas of the Savannah field in the Physical World. The former have control $[xy]$ where $x \in [-90, 90]$ and $y \in [-180, 180]$ indicate the latitude and the longitude,[4] respectively. The latter have control area_A, where A

---

[4]In the current model, altitude values (i.e., the Z-axis) are ignored.

Table III. Controls for the Technology Perspective

Control	Arity	Description	Graphical Notation
area_A	1	Area containing all coordinates in A	
$[xy]$	1	GPS position with coordinates	

Fig. 27. Complete example configuration with Physical World, Computational World, Human, and Technology perspective.

specifies the set of Cartesian coordinates belonging to an area.[5] The intended interpretation is that an entity of control $[xy]$ can be contained by an entity of control area_A only if $[xy] \in A$. Positions and areas are represented graphically as orange bullets and hexagons, respectively. The mapping from areas to locales is encoded by linking each area to the corresponding locale in the virtual space. Thus, the arity of both controls area_A and locale is 1. The mapping and the definitions of sets $A, A', \dots$ are parameters of the model that depend on how an instance of the Savannah game is set up and on the physical characteristic of the Savannah field. A summary is given in Table III.

A complete example configuration is given in Figure 27. It corresponds to the configuration defined in Figure 24 with the addition of the Technology perspective on the right-hand side.

Every child/lion/aura and every impala is linked to a GPS position. Also every locale is linked to an area. In our model, the convention is that region 0 is for the physical world with the Human perspective, region 1 models the Computational World and region 2 hosts the entities of the Technology perspective. Observe that at this point, there are no outer names in the model: all the links are closed because no other entities need to be added.

The dynamics of the bigraphical model for the Technology perspective is defined by nine reaction rules.

The first two reaction rules model a child/lion pair and an impala entering the game at GPS position $[xy]$.

Their graphical representation is given in Figure 28. Both reaction rules are parameterised on the entry position $[xy]$ which is within area A. Our model only represents the GPS positions associated to the child/lion pairs and imapala present in the game. Note that when two entities are at the same GPS position, they are associated to two distinct entities with the same control $[xy]$. Note also that if these rules are only

---

[5]All sets $A, A', \dots$ in the model are assumed disjoint, because locales are disjoint.

(a) enter-l

(b) enter-i
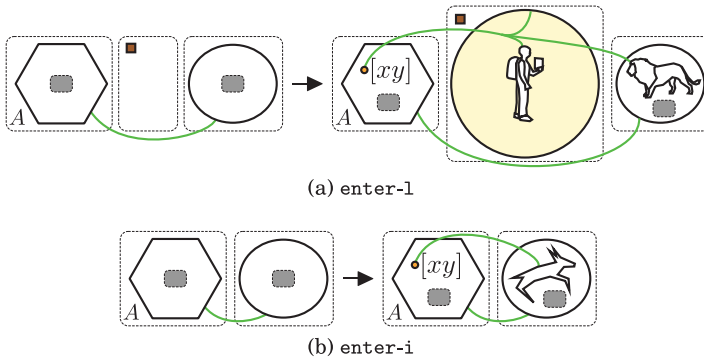
Fig. 28. Reaction rules for adding a child/lion pair (a) and an impala (b) to the game at GPS position $[xy] \in A$.

Fig. 29. kill: a group of three children/lions kill their impala.

Fig. 30. mov1: a lion moves within the same locale: $[xy]$, $[x'y'] \in A$.

applicable when the game is initialised, then we could use rule priorities to enforce this restriction.

The third reaction rule is given in Figure 29. It models a group of three childgroup/liongroup killing an impalaheld. The only difference with the reaction rule we defined in Figure 22 is that GPS position $[xy]$ linked to the impalaheld on the left-hand side is now removed from the region modelling the Technology perspective on the right-hand side. In this way, no idle (i.e., without links) GPS positions are ever introduced by the reaction rules.

Finally, the lion movements caused by GPS updates are modelled by six conditional reaction rules of two kinds: movement within the *same* locale and movement to a *different* locale. A reaction rule for the first kind of movement is given in Figure 30. The other reaction rule of this kind models an idle lion moving within a locale in which an attack is taking place. It is defined by substituting the entities of control locale with entities of control localeattack. The fact that the movement takes place within a locale is specified in the side condition where the new GPS position $[x'y']$ is within area$_A$, the same area that also contains the old position $[xy]$ on the left-hand side. The four reaction rules for the second kind of movement take the form of the reaction rule in Figure 31. In this case, the side conditions are $[xy] \in A$ and $[x'y'] \in A'$. The other three reaction rules of this kind are obtained by replacing locale with localeattack. In more detail, the other three possible combinations are: locale–localeattack, localeattack–locale, and localeattack–localeattack. A distinctive feature of

Fig. 31.   mov2: a lion moves to a different locale: $[xy] \in A$ and $[x'y'] \in A'$.
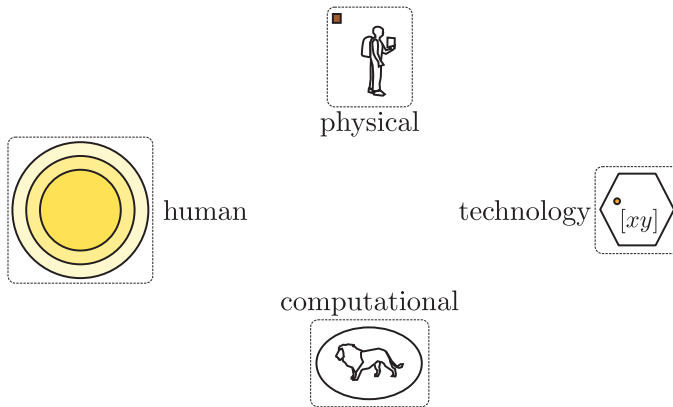
Fig. 32.   Overview of bigraphs for the four perspectives.

the rules for lion movements is that they involve only entities in the Computational World that have a technology perspective. In particular, this means that our model does not enforce any correspondence between GPS updates and aura reconfigurations in the human perspective, nor does it model impala movements (which cannot be sensed).

We have now traversed each of our four perspectives, progressively extending our formal model to capture the key element of the Savannah game and user study as reported in previous work. A complete summary of the controls and the reaction rules of the model is given in Appendix A. Figure 32, a version of Figure 8 specialised to bigraphs, summarises how we model the different perspectives.

9. THE "THREE GIRLS, A BOY, AND AN IMPALA PROBLEM"

Having demonstrated the effectiveness of bigraphs for developing a basic model of the Savannah game, we now explore how we might put our model to use. This involves using it to explain key interactional difficulties that were observed in the Savannah user study and to reveal how these emerged from complex dependencies between our four perspectives. Given space limitations, we focus on one particularly challenging example of struggling to play Savannah that was first reported in Benford et al. [2005]. This single example that we call "the three girls, a boy, and an impala problem," in which three lionesses struggle to attack an impala, demonstrates various interactional challenges. Applying our model to this example serves to both verify the model and demonstrate its power to explain complex interactional phenomena.

Our particular focus is a moment when three players attempt to launch an attack on an impala (Figure 33). The impala is first discovered by the lioness Elsa upon entering the purple virtual locale labelled ADT8. Figure 33 shows stills from our video recording alongside the positions of the players as recorded by our game server at the start (left)
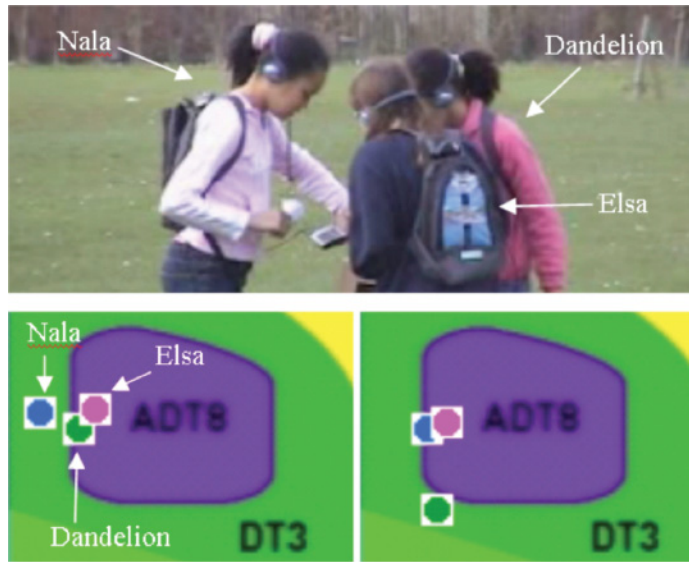
Fig. 33. Three lionesses struggle to attack an impala.

and towards the end (right) of this passage of play. Elsa immediately gathers two other players, Dandelion and Nala, around her to form a tight huddle facing inwards while all three align their handheld computers just inches apart in an attempt to all see the prey before launching an attack. However, this proves to be difficult. The transcript of their conversation makes it clear that at first Nala cannot see the impala while Elsa and Dandelion can, but that later on Dandelion cannot see the prey while the other two now can, even though none of them have shifted their physical positions in the meantime. The three girls become increasingly frustrated throughout the sequence and in the end accuse each other of prematurely launching the resulting attack (which does eventually take place) before all three of them are able to join in. In fact, the attack occurs because a boy who is passing through locale ADT8, but who never joins in their huddle, launches the attack on his own—an action that the three girls never spot.

The underlying causes of these difficulties become somewhat clearer when one examines the system logs of interactions alongside the video. Figure 33 clearly shows that Elsa remains inside the locale containing the impala (labeled ADT8) throughout the sequence, while first Nala, and then, Dandelion find themselves outside at different times. It also shows that this problem occurs in part because Elsa, who first notices the impala and summons the others to gather around her, is positioned close to the edge of the locale. In turn, this exacerbates the effect of a small degree of drift in their GPS positions, which itself may arise from a self-correction that is applied by the GPS unit to its dead-reckoning algorithm whenever a moving player comes to a halt. While it might just be bad luck that Elsa is positioned near to a locale edge in this case, analysis of the video recordings of other sequences reveals this to be relatively commonplace, in large part due to the natural human behaviour of immediately halting when first encountering a prey animal, which tends to be on first entering a locale. However, in this case, it is the boy who passes through the locale without stopping who actually launches the attack.

In summary, their considerable interactional difficulties arise from various interleaved causes that span

Fig. 34. Game configuration according to the girls' perception: Elsa thinks she is initiating the kill. The girls are unaware of the presence of the boy in locale A. The real configuration is S_6 given in Table V.

Fig. 35. Girls' perception of the game configuration: Nala thinks she is joining Elsa in killing the impala. They see that Dandelion is not in their locale but they hope her lion will show up in locale A before the timer of their killing expires. The real configuration is S_7 given in Table V.

—*computational design* in which the virtual Savannah is divided into discretely bounded locales whose boundaries are invisible to players;
—*human behaviour*, including tendencies to form close knit huddles before launching an attack, for the first player to discover prey to stop at the edge of its locale, and for players in a huddle to ignore others (the boy) who are not in the huddle but may in fact be in the (invisible) locale;
—*technology behaviour*, especially GPS inaccuracy including a tendency for GPS to self-correct its dead reckoning mechanism and/or to drift a little short after players come to a halt following walking.

The challenge we face is unpacking the complex ways in which these factors combine to contribute to the difficulties experienced by the players so as to understand how we might either refine the system to prevent these difficulties or mitigate them when they arise. Our starting point is in understanding this episode in terms of a sequence of bigraphs.

9.1. Bigraph Evolution of the "Three Girls, a Boy, and an Impala Problem"

We now give the sequence of bigraphs, $S_0 \ldots S_9$ in Tables IV and V, which illustrates clearly the sequence of events described above, namely how and why the boy was able to become part of the hunting group and make the kill, but the girls were unaware of his participation. For simplicity, we give only the entities relevant to the trace: we omit other locales, areas, children, lions, and impala. Also, we refer to "girls" and "boy," rather than child/lion pairs; recall that locales are discrete, i.e., they are non-overlapping. The sequence is as follows:

$$S_0 \underset{\texttt{mov2}}{\Rightarrow} S_1 \underset{\texttt{mov2}}{\Rightarrow} S_2 \underset{\texttt{mov2}}{\Rightarrow} S_3 \underset{\texttt{mov2}}{\Rightarrow} S_4$$
$$\underset{\texttt{attack}}{\Rightarrow} S_5 \underset{\texttt{join1}}{\Rightarrow} S_6 \underset{\texttt{join2}}{\Rightarrow} S_7 \underset{\texttt{kill}}{\Rightarrow} S_8 \underset{\texttt{mov2}}{\Rightarrow} S_9.$$

Note that at this stage, after the successful attack, and the departure of the boy, the girls have a faulty perception of what has happened. Figures 34 and 35 give an

Table IV. Sequence $S_0 \xrightarrow{}_{\texttt{mov2}} S_1 \xrightarrow{}_{\texttt{mov2}} S_2 \xrightarrow{}_{\texttt{mov2}} S_3 \xrightarrow{}_{\texttt{mov2}} S_4 \xrightarrow{}_{\texttt{attack}} \cdots$ (Continues in Table V)

State	Graphical Notation	Description
S_0 $\downarrow$ mov2		Initially, there are three girls in one shared aura: Elsa (top player in the aura), Dandelion (lower player right aura) and Nala (lower player left in aura). Elsa enters a locale with an impala (i.e. $A =$ ADT8), the other two girls are in a different local, A'.
S_1 $\downarrow$ mov2		Dandelion (lower player right) enters the same locale as Elsa and the impala.
S_2 $\downarrow$ mov2		Nala (lower player left) enters the same locale as the other two girls and the impala. At this point, any of the girls could initiate an attack.
S_3 $\downarrow$ mov2		Dandelion drifts out the locale before an attack is initiated by any of the girls.
S_4 $\downarrow$ attack		The boy enters the locale, but not the shared aura.

Table V. Sequence $\cdots \xrightarrow{\text{attack}} S_5 \xrightarrow{\text{join1}} S_6 \xrightarrow{\text{join2}} S_7 \xrightarrow{\text{kill}} S_8 \xrightarrow{\text{mov2}} S_9$ (Starts in Table IV)

| State	Graphical Notation	Description
S_5 $\downarrow$ join1		The boy has seen the impala and initiates a kill.
S_6 $\downarrow$ join2		Elsa joins the group. She mistakenly *thinks* she initiated the kill, i.e. she thinks the situation is as given in Figure 34.
S_7 $\downarrow$ kill		Nala joins the group, mistakenly thinking she was joining Elsa's group, i.e. she thinks the situation is as given in Figure 35.
S_8 $\downarrow$ mov2		Now there are three in the attack group and so the impala is killed (and disappears).
S_9		The boy leaves the locale.

Fig. 36.   Relationships between the four perspectives (a) and missing relationship (b).

account of what they (mistakenly) perceived as two game configurations they experienced during the sequence.

9.2. Reflections on the "Problem" and the Formal Model of Four Perspectives

The observations of user difficulties led the game designers to propose several potential solutions, from implementing fuzzy locales with a sense of hysteresis to the use of 3D spatialised media that would lead players towards the centre of a locale. The merits or otherwise of these ideas are not our focus here. Instead, the key point for this article is to recognize the multi-layered social-technical causes of interactional difficulties and that a formal model of the system needs to account for the subtle complexities of how humans and systems interact with one another. As we stated earlier, the key question for us is how could we have modelled formally the original design of the Savannah game in a way that might have helped us more quickly understand or even predict the problems that ultimately arose, and what does this tell us about formal approaches to modelling such systems?

We have already seen how our approach based on modelling the four perspectives makes clear and precise the interactions between the perspectives. Specifically, the model reveals that the game was designed such that movement within a locale, as sensed by a new GPS position, only involves entities in the Computational World that have a Technology perspective. This means that a GPS update may invoke a lion movement, but there is no enforced correspondence between GPS updates in the Technology perspective and aura reconfigurations in the Human perspective. In other words, the relations between the perspectives are *not complete* and this can be a source of confusion for the players.

9.3. A Missing Pairwise Relationship?

Recall that the Computational World perspective involves entities such as *lions*, *impalas*, and *locales*, the Physical World perspective involves *children*, the Technology perspective involves areas and *GPS* coordinates in *areas*, and the Human perspective involves *auras*. Figure 36(a) indicates the pairwise relationships between the bigraph models of the four different perspectives. Specifically, a solid line between two perspectives indicates that a change to entities in one perspective induces a change to entities in the other. (Note, this Figure is not a bigraph, but simply an enhancement of Figure 32 indicating relationships between the perspectives.) It is easy to observe there is a missing pairwise relationship: between Human and Technology, as indicated by the dotted line in Figure 36(b).

In the next section, we consider how reasoning about properties of the model can reveal this incompleteness.

10. DOING MORE WITH THE MODEL

We have shown in Section 9 how the formal model can explain behaviours after the fact, but can the model also help us to predict behaviours and confirm (or otherwise) design decisions and assumptions?

We answer this question in two stages. First, we extend bigraphs with *bigraphical patterns* that enable us to reason formally about some of the invariant properties that we would expect to hold for a well designed game and also to analyse whether certain game states are reachable at all or within a given number of steps. Second, we consider how to extend the model by introducing further reaction rules that would remove its inconsistencies and that in turn suggest possibilities for re-designing the game itself.

10.1. Bigraphical Patterns

We begin by considering invariant properties. Specifically, we consider relationships between perspectives, for example, properties such as a child can never be inside a locale or that a child is always linked to a lion. To define properties, we introduce the concept of bigraphical patterns.

Intuitively, properties about relationships will involve checking whether one or more patterns occur in a bigraph. So, we consider properties that consist of standard Boolean operators[6] and *bigraphical patterns* as atomic propositions. Bigraphical patterns are represented graphically in a notation similar to the one used for bigraphs and are indicated by boldface type (e.g., **Lion**, **Field**). They are bigraphs in which

—entities can assume *disjoint controls* and
—link identifiers can be matched.[7]

A pattern is true, for a given bigraph, if an instance of the pattern, with any of the controls, occurs in that bigraph and it exposes the named link. We use solid green shapes to indicate entities that match any control (i.e., colour) and a solid purple border to indicate a locale of any control (i.e., locale or localeattack). The patterns we use in our analysis are given in Tables VI and VII. For example, the green lion in formula **Lion**$_a$ indicates an entity of control lion or lionattack or liongroup, linked to a. If name identifiers are omitted (see, for instance, the patterns in Table VII), then this is not a problem: any open link can be matched, regardless of name.

In order to express properties of bigraphical reaction rules, we introduce the notation **P**^{lhs} and **P**^{rhs} to indicate that pattern **P** refers to the left-hand side and the right-hand side of a reaction rule, respectively. This lightweight notation is sufficient for our analysis, based on bigraph matching as defined in Sevegnani and Calder [2015], though we note that more extensive logical properties of bigraphs can be expressed in the full-blown spatial logic **BiLog** [Conforti et al. 2005].

10.2. Analysing Relationships

We now consider relationships between perspectives, starting with the Computation and Physical World perspectives.
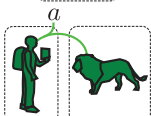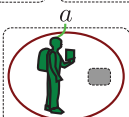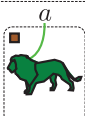
10.2.1. Computational/Physical World. Consider the following properties.

(1) *"If there exists a link between a player and a lion, then that link persists in all subsequent states."*
(2) *"A player cannot be in a locale."*
(3) *"A lion / impala cannot be in a field."*

---

[6]The Boolean operators we use are: negation ($\neg$), disjunction ($\vee$), conjunction ($\wedge$) and implication ($\Rightarrow$).
[7]Note that in bigraphical matching an open link can be matched to any link, even to a closed one.

Table VI. Patterns for Properties Computational/Physical World

Formula	Description	Graphical Notation
$\mathbf{Child}_a$	Child with name a	
$\mathbf{Lion}_a$	Lion with name a	
$\mathbf{Impala}_b$	Impala with name b	
$\mathbf{Child\text{-}Lion}_a$	Child-lion pair linked over name a	
$\mathbf{Child\text{-}Loc}_a$	Child with name a within a locale	
$\mathbf{Lion\text{-}Field}_a$	Lion with name a in a field	
$\mathbf{Impala\text{-}Field}_b$	Impala with name b in a field	

We expect all these properties to hold for *any* configuration of the game (starting from a valid initial state with no players or lions). Our proof strategy consists of showing that the properties are preserved by all reaction rules, and then, by induction, they hold for all configurations reach by application of the rules. In the following, we show how we use bigraphical patterns to reason about the reaction rules.

(1) *"If there exists a link between a player and a lion, then that link persists in all subsequent states."*
This requires checking the following formula holds for all reaction rules:
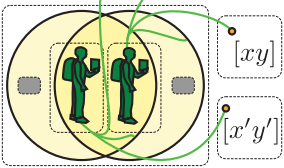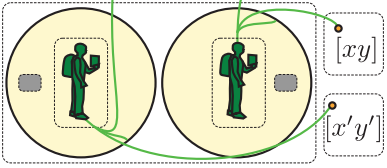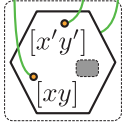
$$\left(\mathbf{Child}_a^{lhs} \Rightarrow \mathbf{Child}_a^{rhs}\right) \wedge \left(\mathbf{Lion}_a^{lhs} \Rightarrow \mathbf{Lion}_a^{rhs}\right)$$

with $a \in \{a, a', a''\}$ and patterns $\mathbf{Child}_a$ and $\mathbf{Lion}_a$ defined in Table VI. Informally, this specifies that whenever a child or a lion with an open link a occurs on the left-hand side of a reaction rule, then the same child or lion also occurs on the right-hand side and its link is still open. It is easy to see, by inspection, that the formula holds for all reaction rules.
We note that our original, intuitive formulation was

$$\left(\mathbf{Child\text{--}Lion}_a^{lhs} \Rightarrow \mathbf{Child\text{--}Lion}_a^{rhs}\right)$$

Table VII. Patterns for Properties Technology/Human

Formula	Description	Graphical Notation
Over-Aura	Two children with overlapping auras	
Disjoint-Aura	Two children with disjoint auras	
Same-GPS	Two positions in the same area	
Different-GPS	Two positions in two distinct areas	

However, this is not sufficient because some rules specify only a sub-part of the system, i.e., only a child or a lion but not both (see, for example, reaction rule `attack`). Moreover, this formula is not necessary because it is subsumed by the use of patterns **Lion**$_a$ and **Child**$_a$.

(2) *"A player cannot be in a locale."*
This requires checking the following formula holds for all reaction rules:

$$\neg\,\textbf{Child-Loc}_a^{rhs}$$

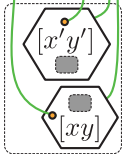with $a \in \{a, a', a''\}$ and pattern **Child-Loc**$_a$ defined in Table VI. Informally, the formula says simply that pattern **Child-Loc**$_a$ may never occur on the right-hand side of any reaction rule. (Since the formula holds in the initial state, it is sufficient to check only the right-hand sides of the reaction rules). It is easy to see that the formula holds for all reaction rules.

(3) *"A lion/impala cannot be in a field."*
This requires checking the following formula holds for all reaction rules:

$$\neg\big(\textbf{Impala-Field}_a^{rhs} \vee \textbf{Lion-Field}_a^{rhs}\big)$$

Similar to the case above, since the formula holds in the initial state, it is sufficient to check only the right-hand sides. Again, it is easy to see that the formula holds for all reaction rules.

We now turn our attention to the relationship between Technology and Human perspectives.

10.2.2. Technology/Human. Consider the following properties.

(1) *"If two children at positions [xy] and [x′y′] have overlapping auras, then positions [xy] and [x′y′] are in the same area."*
(2) *"If positions [xy] and [x′y′] are in the same area, then the two children at positions [xy] and [x′y′] have overlapping auras."*
(3) *"If two children at positions [xy] and [x′y′] have disjoint auras, then positions [xy] and [x′y′] are in different areas."*
(4) *"If positions [xy] and [x′y′] are in different areas, then the two children at positions [xy] and [x′y′] have disjoint auras."*

We observe that unlike the previous pair of perspectives, no reaction rule involves entities belonging to both perspectives. This suggests that there are no constraints or synchronisations on updates to the two perspectives, which leads us to hypothesise that the properties do not hold. In the following, we show this to be the case by giving the corresponding patterns and a counter-example. In each case, the counter-example is a sequence of states and reactions leading to a state in which the given property does not hold.

(1) *"If two children at positions [xy] and [x′y′] have overlapping auras, then positions [xy] and [x′y′] are in the same area."*
This requires checking the following formula holds for all configurations and pairs of GPS positions:

$$\textbf{Over-Aura} \Rightarrow \textbf{Same-GPS} \qquad (1)$$

where the patterns are defined in Table VII. Note that in both patterns, open links are not associated with names. This is because the property concerns bigraphs representing configurations, or states, of the system (not reaction rules) and they do not have open links.
The following sequence is a counter-example:

$$G_0 \underset{\text{enter-l}}{\longrightarrow} G_1 \underset{\text{enter-l}}{\longrightarrow} G_2 \underset{\text{aura-join}}{\longrightarrow} G_3 \qquad (2)$$

where G_0 is an initial state of a game with no children, the first child enters the game at position $[x, y] \in A$ in state G_1 and the second child enters at position $[x′y′] \in A′$ in state G_2. The formula is not true in G_3 because the two children have overlapping auras but the corresponding positions are in different areas. We note that the property *can* be true, for example it holds for state S_2 in the examples presented in Section 9.
(2) *"If positions [xy] and [x′y′] are in the same area, then the two children at positions [xy] and [x′y′] have overlapping auras."*
This requires checking the converse of the previous formula:

$$\textbf{Same-GPS} \Rightarrow \textbf{Over-Aura}$$

In this case, a counter-example is

$$G_0 \underset{\text{enter-l}}{\longrightarrow} G_1 \underset{\text{enter-l}}{\longrightarrow} G_2 \qquad (3)$$

where the first child enters at a position $[x, y] \in A$ in state G_1 and the second child enters at position $[x′y′] \in A$ in state G_2. The formula does not hold in state G_2 because the two GPS positions are in the same area but the two auras are disjoint. Note that the formula is also not true in all states S_i with $4 \leq i \leq 8$ in Section 9.
(3) *"If two children at positions [xy] and [x′y′] have disjoint auras, then positions [xy] and [x′y′] are in different areas."*

This requires checking the following formula holds for all configurations and pairs of GPS positions:

$$\textbf{Disjoint-Aura} \Rightarrow \textbf{Different-GPS}.$$

The counter-example is the same sequence used for the previous property. The formula does not hold in state G_2 because the two children have disjoint auras but their positions are in the same area. Note also that the formula is not true in all states S_i with $4 \leq i \leq 8$ in Section 9.

(4) *"If positions* [xy] *and* [$x'y'$] *are in different areas, then the two children at positions* [xy] *and* [$x'y'$] *have disjoint auras."*
This requires checking the following formula holds for all configurations:

$$\textbf{Different-GPS} \Rightarrow \textbf{Disjoint-Aura}$$

The counter-example is the same sequence used for the first property above. The formula is false in G_3 because the two positions are in distinct areas but the corresponding children have overlapping auras.

10.3. What does Analysis Reveal?

What can we conclude from this analysis? We have shown that the Computational/ Physical World properties (1)–(3) *are* invariants, whereas Technology/Human properties (1)–(4) *are not* invariants. We note this does not mean that the latter do not hold for some configurations: we have seen that they can be true for some sequences. But if we accept that the properties in Section 10.2.2 are desirable (and we do), then this analysis indicates a design flaw in the game, and a plausible explanation of the major interactional difficulties experienced by players. It also indicates that a possible re-design would involve the addition of reaction rules that synchronise Technology and Human (perspective) updates.

We note that whereas it is relatively straightforward to verify invariants, we require to unfold computation (i.e., rewrite) to disprove them.

Finally, we note that we have concentrated on properties concerning interaction and human experience, rather than on the functional behaviour of the system that is traditionally a focus for analysis of formal models. An indication of the latter is analysis that includes reasoning about game sequences to determine:

(1) Given an example configuration (e.g., four players, three locales, two impala), can any player initiate an attack and that attack lead to a successful kill? What assumptions would we have to make about player movements?
(2) Can a player initiate an attack within n steps, or put another way, what is the minimum length of path to a kill, given a specific initial configuration?

These are temporal properties and proof (e.g., by model checking) could involve exploring large state spaces, especially when including arbitrary movements with large numbers of players and locales. However, abstractions are possible. For example, we could reduce the number of locales: only locales containing impala are explicitly represented while all the other locales are merged into a logical locale called *inactive*, in which killings cannot take place. In Section 10.5, we present two results concerning state spaces.

10.4. Re-Design the Model to Introduce Technology/Human Relationship

So far, we have shown how our model might help predict various inconsistencies in the game design that might in turn lead to interactional difficulties. We now show that it is possible to extend to model with further reaction rules in order to remove these inconsistencies, which in turn, might inspire ways of re-designing the game itself.
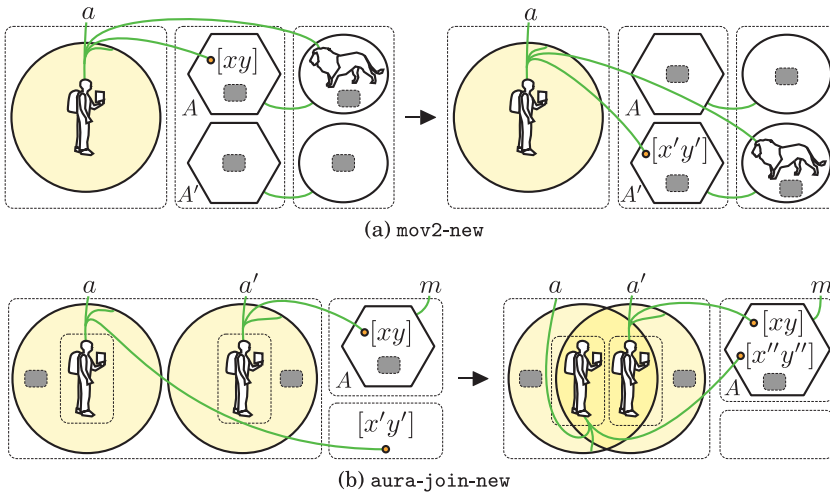
(a) `mov2-new`

(b) `aura-join-new`

Fig. 37. New reaction rules for Technology/Human relationship.

Recall Technology/Human property (1): *"If two children at positions $[xy]$ and $[x'y']$ have overlapping auras, then positions $[xy]$ and $[x'y']$ are in the same area."*

In order to make this property invariant, the reaction rules of the BRS have to be modified in such a way that any rule application can *never* lead to a state in which the following formula[8] holds

$$\textbf{Over-Aura} \wedge \neg\textbf{Same-GPS}. \tag{4}$$

By inspection of the current model, we observe that this formula may hold only in the states obtained as follows:

(1) by application of reaction rule `mov2` to a state in which property (1) holds (i.e., overlapping auras and corresponding GPS positions in the same area),
(2) by application of reaction rule `aura-join` to a state with disjoint auras and corresponding GPS positions in different areas.

Intuitively, `mov2` may break "correct" configurations, while `aura-join` may build "incorrect" configurations. These two kinds of transition are allowed because no constraint on aura configurations is specified by the left-hand side of `mov2` (in the first case) and because of the lack of constraints on GPS positions on the right-hand side of `aura-join` (in the second case).

This analysis suggests the changes required for the BRS to satisfy the property: reaction rules `mov2` and `aura-join` have to be modified in order to avoid the generation of the two kinds of transitions described by the two cases above.

Diagrams for the two new rules are given in Figure 37. Reaction rule `mov2-new` replaces `mov2`. It specifies that only children with disjoint auras are allowed to move to a different area. Reaction rule `aura-join-new` replaces `aura-join`. It specifies that the event of joining auras is associated to a GPS update (i.e., $[x'y'] \rightarrow [x''y'']$). Moreover, on the right-hand side, new position $[x''y'']$ is in the same area of the GPS position linked to the other child. All the other reaction rules are left unchanged.

The counter example given in Equation (2) cannot be generated by the new model. However, this is not sufficient to prove the new BRS satisfies the property. Instead,

---

[8]Note this formula is the negation of the formula given in Equation (1).

Table VIII. Patterns for Properties Technology/Human in the New Model

Formula	Description	Graphical Notation
$\textbf{Over-Aura}_{a,a'}$	Children a and a' with overlapping auras	
$\textbf{Same-GPS}_{a,a'}$	Positions a and a' in the same area	
$\textbf{Aura}_a$	Child a with disjoint aura	
$\textbf{GPS}_{A,a}$	Position a in area A	

we define a logical formula and prove it holds for all reaction rules[9] as we did for the verification of Computational/Physical World properties (1)–(3). This is possible because the new model includes reaction rules involving entities belonging to the Technology and the Human perspectives.

Before defining the formula, we introduce four new bigraphical patterns as in Table VIII. They are similar to the patterns given in Table VII. The main difference here is that patterns specify name identifiers. This is necessary because the new patterns are used to define a formula over reaction rules instead of states.

The formula we need to check is defined as follows:

$$\left(\left(\textbf{GPS}_{A,a}^{lhs} \wedge \textbf{GPS}_{A',a}^{rhs} \right) \Rightarrow \textbf{Aura}_a^{lhs} \right) \wedge \left(\textbf{Over-Aura}_{a,a'}^{rhs} \Rightarrow \textbf{Same-GPS}_{a,a'}^{rhs} \right).$$

It is routine to verify it holds for all reaction rules and for all the pairs of GPS positions (in the reaction rules). Therefore, Technology/Human property (1) is invariant in the new model. In more detail, the left and right conjuncts handle reaction rules `mov2-new` and `aura-join-new`, respectively. The formula holds for all the other reaction rules since the antecedents of both implications are always false.

Observe that the formula does not hold in the original BRS because it is false for rules `mov2` and `aura-join`. Furthermore, new reaction rules `mov2-new` and `aura-join-new` are not sufficient to make Technology/Human properties (2) and (3) invariants as the counter example defined by Equation (3) can still be generated. On the other hand, it can be proved that property (4) is invariant in the new model.

---

[9]Alternatively, we could unfold computation and use model checking to prove that the formula in Equation (1) holds for all states.

10.4.1. Implications for System (Re)design. We have seen that it is possible to extend the model of the Savannah game to make it formally consistent by introducing further reaction rules that properly connect the Human and Technology perspectives to the Computational and Physical World perspectives. It is important to note that the new model cannot generate the example sequence describing the "three girls, a boy, and an impala problem" we described in Section 9. This means the new BRS is not a faithful representation of the current implementation of the Savannah game. In particular, the new reaction rules introduce a correspondence between proximity and social interaction that does not emerge from the analysis of the user trials (recall that the boy was physically close to the girls but he did not interact with them). Therefore, we argue that the original model is the correct model for the current specification of the Savannah game.

The new model however, could be used to indicate future development directions and test the effects of the introduction of new features or new game rules. The question becomes how might the system designers respond to the new reaction rules? Given that the source of the problems appears to lay in the Human and/or Technology perspectives, we might naturally turn to these for new design solutions.

Considering the Human perspective, we might think that the root of the problem is that the notion of forming social groups by overlapping auras only exists in the players heads, i.e., in the original Savannah game it is a purely social phenomenon that has no first-class representation in the system itself. One general solution might then be to bring auras and groups into the system itself. We might implement an explicit grouping mechanism based on overlapping auras (driven by GPS), and then, reveal to players when they are or are not in a group with others. We might then ensure that all group members see a consistent state by temporarily shifting locale boundaries to cover the whole group (essentially this is the proposal for fuzzy locale boundaries that was discussed in Benford et al. [2005]).

We might also look for solutions in the Technology perspective, perhaps introducing additional sensing modalities to determine relative local proximity between players (various near-field sensing and communication technologies might support this). Yet another approach might be to better connect the Human and Technology perspectives by revealing to the players at the interface how GPS has positioned them with respect to locale boundaries so that they can then adjust their own positions to make the game consistent, essentially pushing the solution into the Human perspective by providing players with the resources to be able to adapt their behaviours to the technology.

Whether or not such ideas lead to a better game experience in practice will depend on a variety of additional factors such as whether they can be wrapped up in engaging game mechanics and narratives. However, our point is that the introduction of new reaction rules not only serves to make the game formally consistent, but also may help inspire possible new design features.

As a final note on this topic, formal modelling of the system may also serve to reveal key underlying assumptions. For example, the earlier rule: *"If there exists a link between a player and a lion, then that link persists in all subsequent states"* exposes an important and hitherto hidden assumption in the design of the game. What if players were to swap devices? That this did not happen in the Savannah study was perhaps due to the presence of teachers and researchers giving out devices to players (including labelling them with the names of lions), and then, expecting the children to report back on what they had done. However, this was a social convention. It would have been quite possible for players to swap devices (and perhaps even reasonable to do so if batteries had run down) and in other contexts we might more naturally expect people to temporally swap phones. Formal modelling then can also help expose potentially significant assumptions as well as reveal inconsistencies.

10.5. Experiments With State Space

We conducted two experiments to explore the state space of the model using the BigraphER system.[10] BigraphER is an implementation of BRS and stochastic BRS that supports place graphs with sharing. The tool consists of an OCaml library and a command-line tool that provides efficient manipulation and simulation of BRS and stochastic BRS (see Sevegnani and Calder [2015] for further details).

First, we considered the example trace presented in Section 9. For this, we added all the rules excepting enter-i and enter-l (we assumed a fixed number of players (4) in the initial state), and restricted the GPS updates to the five movements observed in the trace. From the initial state, BigraphER generated[11] the full transition system (107 states and 252 transitions) in 1.3s; the "problem trace" is just one (portion of a) path through that state space. Second, we relaxed the restriction on updates to allow all possible updates within the same area (with reaction rules $mov1(a, p, p')$[12]) and between any two areas (with reaction rules $mov2(a, a', p, p')$[13]). These rules contain free variables and BigraphER automatically computes all the possible reaction rules generated from them. For example, $mov1(0, p0, p0')$ generates 25 reaction rules, one for each combination of values of $p0$ and $p0'$ (5 possible values each). As expected there is state-space explosion, and BigraphER generated the full transition system (20,666 states and 486,586 transitions) in 4.8h. These two experiments indicate it is possible to generate all the sequences of configurations with BigraphER in a reasonable amount of time. However, some care needs to be taken when specifying the exact reaction rules employed; to avoid state-space explosion yet still remain faithful to the problem, some mitigations are possible. For example, we can

(1) terminate each computation path after the impala has been killed;
(2) disallow GPS updates within the same area—they only add interleaving and do not enable any further reaction rules;
(3) employ counter abstractions. As discussed earlier we can model only two locales: one indicating the locale in which the impala is roaming, the other representing all the other locales, and instead of modelling each player individually, we can record how many players are in each locale.

11. DEALING WITH AN UNCERTAIN WORLD

Ubiquitous computing systems are suffused with uncertainty. This uncertainty ranges from the nature of the sensors that underpin these systems to the human behaviour that drives interaction with them. Consequently, we need to reflect the various forms of uncertainty involved in ubiquitous systems in our models. In order to do this, we now turn to one final extension to our model, the modelling of probabilistic events that represent possible sources of randomness in the Savannah game. We illustrate this by considering how we might deal with the uncertainty inherent in GPS sensing and how we might reflect the uncertain nature of user-initiated events.

11.1. Probabilistic Sensing and GPS Positions

So far, we have assumed that a sensed GPS position uniquely identifies a position in the physical world and that position is within the Savannah field. However, this assumption is often too strong as GPS sensing is intrinsically probabilistic due to multiple factors

---

[10]The source code for these experiments is available at http://www.dcs.gla.ac.uk/~michele/savannah.html.

[11]On a machine running FreeBSD 10.1 amd64 with four i5-3570 cores at 3.40GHz and 10597MB of memory.

[12]Notation $mov1(\_, \_, \_)$ indicates the expression in the BigraphER language defining bigraphical reaction rule mov1 defined in Figure 30. Argument a indicates an area, while arguments p and p′ are the GPS positions on the left-hand side and the right-hand side of the reaction rule, respectively.

[13]Arguments a and a′ indicate areas, while p and p′ are GPS positions as in $mov1(\_, \_, \_)$.

Fig. 38. Sensed position $[c_0 c_1]$ in example physical Savannah field.

that affect accuracy and precision. These include the positions and availability of GPS satellites, atmospheric noise, multi-path reflections from nearby objects, and also the use of software techniques such as dead reckoning to try and compensate for these factors under assumed circumstances (when a player or vehicle is in constant motion) but that may exacerbate them in others (when they suddenly stop). For example the *GPS drift problem* occurs when the sensed coordinates "drift" even though the GPS receiver is at a standstill. We, therefore, now define a richer model for GPS sensed positions as part of the Technology perspective in which a probability distribution over the possible positions in the physical world induces a probabilistic mapping of entities in the model to Cartesian coordinates. In other words, a sensed GPS position is treated as an estimate of an actual physical position. In this richer model, we probabilistically assign children (i.e., the only entities equipped with a GPS sensor) to entities of control area$_A$, i.e., coordinates of form $[xy]$.

Our approach is illustrated by example as follows. Consider a physical Savannah field consisting of areas A_0, A_1, and A_2 with the topology shown in Figure 38 and the sensed GPS position $[c_0 c_1]$.

The probability distribution over all the possible physical positions is given by a three-dimensional mesh and heat map as shown in Figure 39. It can be seen that the physical position with highest probability is within A_2, and for positions in the other two areas, the corresponding probabilities are much lower.

Let us now formalise our approach. A sensed GPS position $[c_0 c_1]$ is interpreted as a random variable X with bivariate Gaussian distribution as follows:

$$X \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma}) \qquad \text{with } \boldsymbol{\mu} = \begin{pmatrix} c_0 \\ c_1 \end{pmatrix} \quad \boldsymbol{\Sigma} = \begin{pmatrix} \sigma^2 & 0 \\ 0 & \sigma^2 \end{pmatrix}.$$

Note that, we assume a simplistic model in which both coordinates have the same variance σ^2 and the covariance between them is equal to 0. The probability p_i of a physical position to be within a given area A_i is given by

$$\mathrm{P}(X \in A_i) = p_i = \iint_{A_i} f_X(x, y) \, \mathrm{d}x \, \mathrm{d}y \tag{5}$$

where f_X is the density of X. The total area of a Savannah field in the physical world is finite but we may obtain readings that outwith this area. Namely, f_X is defined over $\mathbb{R}^2$ and so we also define p_{err}, i.e., the probability that the position is outside the field (not contained by any A_i):

$$p_{err} = 1 - \sum_i p_i. \tag{6}$$

Fig. 39.   Sensed GPS position $[c_0 c_1]$ with corresponding probability density function.


Fig. 40.   Probabilistic interpretation of GPS sensing. The actual physical position (upper part of diagram) is interpreted as several possible bigraphs (lower part of diagram). Only the interpretation within the red box corresponds to the actual position in the physical world.

This is typically close to 0 but it may grow depending on the size and shape of the field and the precision of the GPS sensors.

Each GPS update corresponds to different probabilistic interpretations in the sensor space as shown in Figure 40. This can be encoded at the bigraphical level by adding probabilities to the reaction rules for GPS updates defined in Figures 30 and 31, i.e.,

reaction rules `mov1` and `mov2`. In more detail, a GPS update event is encoded by a family of $n + 1$ (with n is the number of areas in the Savannah field) *probabilistic* reaction rules as follows:

—one reaction rule in the form of `mov1` (See Figure 30) encodes the movement within the same locale,
—$n - 1$ reaction rules in the form of `mov2` (See Figure 31) encode the movement to a different locale,
—one reaction rule in the form of `mov2` encodes a GPS sensing error with the movement to dummy entity $\text{area}_{A_{err}}$.

With this approach all the updates are regarded as probabilistic, with a Gaussian distribution. Note that we do not distinguish between drift, other sources of noise, and actual physical movements that happen to be very small. This is because we have no ability to distinguish the nature of the movement by inspection of a single GPS update, such a distinction would require an inference process over a timed sequence of updates.

11.2. Probabilistic User-Initiated Events

Another source of probabilistic behaviour is the occurrence of user-initiated events such as to initiate an impala kill or join a group. Our model so far assumes that a child always performs an action whenever this is possible, as shown by the definition of reaction rules `attack`, `join1` and `join2` given in Figures 18, 19, and 21, respectively. However, this is not always the case as children may decide to ignore an impala or a group. This can be modelled by assigning probability p to the reaction rules encoding the event (i.e., reaction rules `attack`, `join1`, and `join2`) and probability $1 - p$ to the reaction rules modelling the user ignoring the event. The latter kind of rule is defined as an identity reaction rule in which the left-hand side is the same as the right-hand side. Probabilities for these kinds of user-initiated events could be inferred from logs of user trials. For example, the recent work on inferring *activity patterns* from user logs [Andrei et al. 2014] may indicate a fruitful direction. Here, a finite number of (activity) patterns of usage behaviour (sets of probabilities of transitions between states) are inferred from mass trials (e.g., thousands of users) involving user logs extracted from an instrumented system. Each activity pattern is represented by a discrete time Markov chain, and a distribution over the Markov chains is also inferred. Temporal logic properties are then analysed in different activity patterns to gain insight into how the system has actually been used, for different clusters of behaviour and for different periods of time. The motivation is potential system re-design; we could take a similar approach and test (using a probabilistic temporal logic) whether hypotheses, e.g., about proxemics, are actually observed in practice. It is important to note that these approaches are predicated on modelling observed phenomena, we are not considering intention or higher level semantic goals.

11.3. Technical Note: Probabilistic Bigraphs

To our knowledge, there is no published extension of bigraphs to a probabilistic (DTMC) setting (note the extensions to a stochastic setting [Krivine et al. 2008]). However, we outline such an extension here. Each event is modelled by a family of probabilistic reaction rules, with the sum of their probabilities equal to 1. When n events can occur in one state, the transition probabilities have to be normalised. Different strategies for normalisation can be applied

—when n events are equiprobable, transition probabilities take the form $\frac{p_i}{n}$ where p_i is the probability of the reaction rule generating the transition,
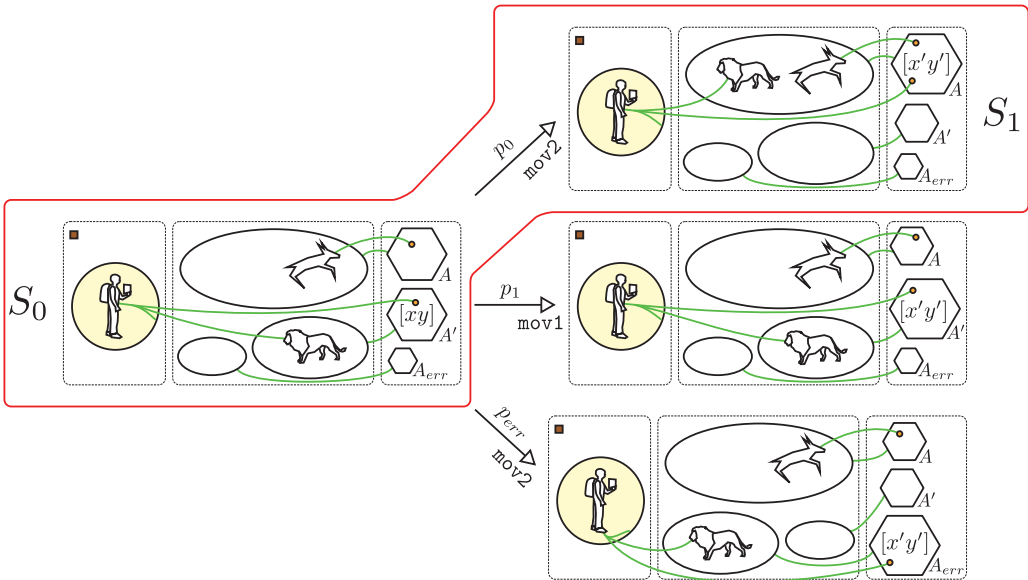
Fig. 41.   Probabilistic transitions from state S_0.

—otherwise, one event has probability q and all the transition probability for the other
  events are in the form $p_i(1 - q)$.

  Some examples of probabilistic transition will be discussed in the next section.

11.4. How to Model Complex Behaviours with Probabilities

The analysis of the user trials suggests that some behaviours emerge because particular
events are more likely to occur than others. One such behaviour, highlighted by the
analysis of Benford et al. [2005], and mentioned in Section 2 is where players tend to
stop (stand still) on the boundaries of locales instead of moving further inside. One
explanation for this behaviour is that as soon as a lion (in the Computational World
perspective) enters a locale populated by impala, a prey appears on the screen of the
corresponding child's device and she/he instinctively stops moving to initiate an attack.
Observe that there are no rules of the game that encourage players to move well into
a locale. Hence, it would be inappropriate to model either moving well into a locale, or
indeed, stopping near the boundary, by introducing ad-hoc reaction rules in the BRS
(ad-hoc in the sense that they do not correspond to rules of the game).

  We propose instead to *infer* probabilities of events from the user trials, and then,
assign these probabilities to different families of reaction rules (see Section 11.3 for
the concept of families of probabilistic reaction rules). We now illustrate this approach
with an example consisting of

—one child/lion pair,
—one impala,
—two locales (plus dummy locale A_{err} to handle GPS probabilistic sensing), and
—GPS update rules in the form $[xy] \rightarrow [x'y']$ and $[x'y'] \rightarrow [x''y'']$.

This is sufficient to show how probabilistic events can model the "stop at boundaries"
behaviour.

  Consider initial configuration S_0 drawn on the left-hand side of Figure 41. Note
that in this small example, at this stage only one event is possible (GPS update),

Fig. 42. Probabilistic transitions from state S_1.

which is modelled by two reaction rules: `mov1` and `mov2`, resulting in three possible new states (one of which is S_0). The probabilities for the three transitions are computed as described above by using equations (5) and (6) and with GPS position $[x'y']$ and areas A, A' and A_{err}. Note that transition $S_0 \dashrightarrow S_1$ models the behaviour of a lion entering a new locale containing an impala.

Now consider the two events (attack and GPS update) and the four possible transitions from S_1, as shown in Figure 42. Let the probability of attack inferred from the analysis of the user trials be indicated by q. Consequently, $1 - q$ is the total probability of the other three transitions modelling the GPS update event. As for the transitions from S_0, the probabilities associated with the GPS updates (namely p'_0, p'_1 and p'_{err}) are computed by integrating the Gaussian centred in $[x''y'']$ over the three areas. Note that if the inferred value of q is close to 1, then the probability of a GPS update is very low. Nevertheless, we would want to retain the possibility of these rare transitions, to model, for example, GPS drift at standstill.

This probabilistic modelling approach allows the BRS to match closely the behaviours observed on the field, without sacrificing formality and clarity of the model by cluttering it with reaction rules that do not have a counter part in the rules of the Savannah game. We hypothesise that if developers can see clearly the effect of q, they may be motivated to re-design the game such that the value assigned to q is lower. For example, they may introduce a delay in the appearance of prey on the screen upon entry to a locale, which means there would be several GPS updates before the first attack.

Of course, the assignment/inference of appropriate probabilities is crucial here. One likely source of probability distributions is more formal experimental HCI studies that systematically assess user behaviours. Such studies have underpinned the optimisation

of conventional graphical user interfaces over the past decades. We anticipate that similar studies might help inform formal models of ubiquitous systems in the future.

12. DISCUSSION

Having introduced our formal model of the Savannah game and put it to work to explain and help resolve previously observed problems, we now draw out wider reflections on modelling ubiquitous systems in general.

It is evidently possible to model formally relatively simple ubiquitous systems such as the Savannah game. Moreover, it is also evident that such models can help reveal underlying inconsistencies in their designs and so explain problems that are revealed during use. Our experience also suggests that such models might predict these flaws in the early stages of design and also inform re-design to address them. However, our article has also revealed that formally modelling even simple ubiquitous systems is not a trivial exercise.

12.1. The Challenges of Modelling Ubiquitous Computing System

Ubiquitous systems pose wide-ranging challenges to formal modelling that go beyond traditional concerns with modelling computational behaviour. First, like interactive systems in general, human behaviour needs to be taken into account. While there is a longstanding tradition of modelling aspects of human cognitive behaviour as part of the design of graphical user interfaces using approaches such as GOMS (Goals, Operators Methods and Selection Rules) [John and Kieras 1996] and the Keystroke Level Model [Card et al. 1980], ubiquitous systems require us to model further aspects of human behaviour such as the use of space to mediate social interactions as we saw in the Savannah game. Second, ubiquitous systems are embedded within the everyday world and so potentially subject to a wide variety of environmental influences. While many agree that context is richer than merely location [Schmidt et al. 1998], including elements that encapsulate the human sense of "place" as well as those that describe the physical structure of space [Dourish 2006], a generalised definition of context remains elusive, making this key concept particularly difficult to model. In fact, there has been considerable debate about the modelling of context and the reductionist dangers involved in adoption a stance that operationalises context [Dourish 2004].

Finally, ubiquitous systems typically rely on invisible and wireless sensing systems, which are notorious for their variable accuracy and coverage, limitations that have come to be known as "seams" in the HCI community, leading researchers to propose strategies for designing around them or even turning their limitations into features (a strategy known as "seamful design"). Our article has shown that, even in a simple case, this combination of factors makes modelling ubiquitous systems complex and that demands new approaches to modelling that in turn are supported by new underlying formalisms.

12.2. Mapping These Challenges Through Perspectives

In order to wrestle with these challenges, we established an overarching approach of systematically adopting different perspectives on a ubiquitous system. This helped make the overall modelling task tractable by allowing us to focus on a specific challenge at a time and also helped explain design flaws in terms of inconsistencies between perspectives. An important feature of this approach is to identify appropriate abstractions for modelling relevant phenomena within each perspective, which may involve drawing on theories and approaches various disciplines, for example, drawing on proxemics to model the concept of aura as a way of encapsulating a key player behaviour when forming groups.

There are three important ideas to stress here. First, is that models are, by their very nature, abstractions of the world. Second, is that it is important to choose abstractions that abstract the relevant features of a given perspective to the problem at hand. It would not, for example, be feasible to model all aspects of player behaviour—cognitive, physiological, emotional, and so forth—within the human perspective. Third, once we move beyond the computational perspective, we will inevitably need to turn to disciplines beyond Computer Science whose research addresses the Human, Physical and Technical perspectives, drawing on their work to identify appropriate abstractions. With these points in mind, we now revisit each of our four perspectives to consider the kinds of the theory and abstractions that might prove useful in the future as we seek to generalise our approach beyond the initial relatively simple example of the Savannah game.

The Computational World perspective represents the conventional focus for formal modelling within Computer Science and is broadly concerned with mathematically verifying the design of system software. While there are potentially many different aspects to this, the modelling of ubiquitous systems will necessarily focus on software mechanisms that map between the external actions of humans in real contexts, as mediated by sensors, to system actions. The abstractions involved may range from relatively simple mappings such as the use of locales to trigger actions as seen in Savannah (although we have seen that even the simplest treatment of locales can give rise to surprising complexity) through to sophisticated machine learning techniques that infer and learn about behavioural or contextual information from sensor data and whose operation may vary over time as well as location.

The Physical World perspective is concerned with modelling key aspects of the physical environment within which a ubiquitous system is situated. Savannah's choice of an open school playing field made this relatively straightforward in our initial example, but it is easy to envisage how this could become far more complex, requiring us to draw on external theories and models of spatial structure. The physical world is characterized by varying terrain, boundaries, buildings, natural features, roads, landmarks, and so forth, all of which might be relevant to the design of ubiquitous systems. The discipline of Geographical Information Systems [Star and Estes 1990] is replete with concepts and techniques for modelling such factors that might be abstracted into our formal models. In turn, Architecture and Urban Design provides concepts for modelling the salient properties of the built environment (e.g., Space Syntax [Hillier and Hanson 1984]), while Environmental Science is concerned with developing models of weather, climate and other dynamic environmental conditions.

We have already seen how modelling the Human perspective might abstract key concepts from the social sciences. Again, Savannah is a relatively simple case. More complex ubiquitous systems might require us to model other "onion layers" from the proxemic model [Greenberg 2011] such as "intimate" or "public" space, or draw on sociological concepts such as *f-formations* [Marshall et al. 2011] that describe how people arrange themselves into small groups during face to face communication, or even socially inspired concepts from the field of Computer Supported Cooperative Work such as *focus* and *nimbus* that introduced notions of directionality and asymmetry in face to face communication [Greenhalgh and Benford 1995]. Stepping up to a different level of scale, the theory of Space Syntax from urban planning explains how the structure of space (from building to city scale) predicts patterns of human movement. While ubiquitous systems tend to be inherently spatial in nature, they may also exploit other connections to the physical world that require us to further unpack that elusive notion of context. The increasing use of cameras and physiological sensors to sense gestures, facial expressions and bodily responses may require us to draw on the theory from Psychology and Medicine in an attempt to abstract the salient features of emotion and

physiology, while an increasing focus on modelling behaviour may require us to draw on theories that account for motivation and intention.

Finally, we turn to the Technology perspective that models key features of the enabling infrastructure of (usually) wireless sensing and communication technologies. This will require us to draw on research from Engineering that explains the detailed operation of such technologies. These might range mathematical models of signal propagation that predict how the positioning of transmitters and receivers (from cell towers and WiFi access points to the predictable orbits of constellations of GPS satellites) affect coverage and accuracy. Such models may be probabilistic in nature, as indeed may be those that attempt to capture key aspects of human and environmental behaviours.

In short, a wide variety of existing concepts and theories might potentially be brought to bear on the modelling of a given perspective. Which are appropriate will depend on the nature of the ubiquitous system at hand: in what kind of environment is it set? What sorts of behaviours might we anticipate or have we observed? Which sensing technologies are being used? There is also the question of the utility of different theories for the purposes of modelling. How readily do they yield abstractions that can be formally modelled? The appeal of proxemics, for example, is that it appears to be sufficiently rich to capture key elements of human social behaviour and yet sufficiently simple that it can be modelled using a formalism such as bigraphs. Although, it is worth recalling that the simplifications involved in proxemics have being critiqued in the social science literature for the reductionist approach towards culture and social interaction.

12.3. The Nature of Bigraphs

While there are clearly many challenges to be addressed in the formal modelling of ubiquitous systems, our experience with bigraphs suggests that this particular formalism provides a good basis for future explorations. This is because the theory of bigraphs provides explicit support for modelling spatial interaction at its core. This fundamental characteristic of bigraphs makes it possible to model spatial structures, and then, to mediate other relationships through these. This said, we found it necessary to extend bigraphs in several ways to be able to model the Savannah game. While previous work had introduced bigraphs with sharing, this article has introduced bigraphical patterns and support for probabilistic modelling. While these extensions are important, others may be required in the future. For example, the current abstraction of space in bigraphs is based upon a notion of containment, whereas adopting some models of social interaction in space such as f-formations [Marshall et al. 2011] would require us take account of the relative orientations of participants and artefacts.

There is however, a further key aspects of bigraphs that we have found to be important here: the ability to represent them diagrammatically as well as algebraically. It is challenging enough for many Computer Scientists and Software Engineers to engage with the mathematical machinery of Formal Computing, but this is likely to prove even more of a barrier as we move out of the Computational World and into the Physical World, Human, and Technology perspectives where we may need to undertake dialogues with other disciplines altogether. Being able to reason diagrammatically about bigraph models has proved essential to enabling an intra-disciplinary dialogue between Formal Computing and HCI within the discipline of computing. It is likely to become more so as other disciplines come into the picture.

12.3.1. Other Formal Modelling Techniques. To our knowledge, there are few formal modelling techniques, or applications of formal techniques, to the design of mixed reality

systems that address both human and system behaviour. Worthy of mention is the design notation ASUR [Dubois and Gray 2007], which is the basis of the functional aspects of two object based development methods for mixed reality systems [Dupuy-Chessa et al. 2010; Dubois et al. 2014]. ASUR is a graphical notation that includes four types of entities: Adapters that bridge the physical and digital worlds, System depicting the digital entities involved in the system, the User of the system, and Real objects taken from the physical world and involved in the interaction. Each of these entities may interact with the others through interaction channels, denoted by arrows, indicating information exchanges when using the interactive system. The latest evolution of ASUR refines the concept of channel to include a characterisation by medium (e.g., light, physical contact, infrared, and air) and a representation that expresses the coding scheme used on a channel, i.e., the syntax of the data. While our chosen perspectives can be mapped to parts of the ASUR entities, we note the notation is entirely graphical, without an underlying semantics of the interactions or dynamic system evolution. We also note that other people have looked at how to model pervasive systems including [Bruegger et al. 2009] who proposed a layered framework for modelling pervasive applications.

13. REFLECTIONS ON FORMAL MODELLING WITHIN HCI

We finish with some wider reflections on the nature of formal modelling within HCI. There is a longstanding debate within Computer Science surrounding tensions between the rigours of theoretical formal modelling and the pragmatics of applied software engineering [Calder 1998; Rushby 2007; Parnas 2008]. This debate also encompasses HCI [Bellotti et al. 1995], which has seen its share of formal modelling, especially in areas such as cognitive modelling and interface architectures, but whose practice has come to be dominated by pragmatic user-centred approaches that involve iterative prototyping and user-testing, supported by methods such as ethnographic field studies and participatory design. More than 20 years ago, HCI researchers were drawing attention to the gulf between those who were advocating the abstract modelling of interactive systems in order to produce optimal designs and those who were seeking immediately practical methods that were seen to be "good enough" [Shum and Hammond 1994]. While the landscape of approaches has undoubtedly become more complex since then, a recent extensive review and critique of the HCI theory has argued that the field still appears to be struggling with a fundamental tension between theory and practice [Rogers 2012].

Our first broad reflection is to challenge any notion that formal modelling and user-centred design are fundamentally opposed. Indeed, a key contribution of our work here is to show how they might be better connected. The results of user testing of the Savannah prototype have driven directly the modelling process, determining what needed to be modeled to explain the findings and motivating and populating the human, environment, and technology perspectives to complement the more conventional (from a modelling point of view) computational perspective. Our work also suggests how mathematical analysis of formal models might inspire re-design of the system to resolve problems as part of an iterative design cycle. Thus, our work directly reflects Bellotti et al's earlier challenge to reconsider the role of formal modelling in interaction design that while formal modelling has often been seen as following on from the derivation of formal requirements and formal specifications, it might also be incorporated into "less structured" design processes, being introduced selectively and at later stage. Our work provides a concrete and detailed example of how this might be achieved.

Thus, we argue that both formal modelling and user-centred design have their strengths. User-centred approaches are readily comprehensible and implementable by interaction designers, while their situated nature can reveal unanticipated issues and expose underlying assumptions and omissions about how humans actually behave in complex real-world settings. Formal methods on the other hand may speak more readily to software designers and can offer a level of proof of correctness that is important in some application areas, for example, safety-critical systems. The question then becomes how to integrate these approaches into a holistic design process that can leverage the benefits of each in a complementary way. With this in mind, we reflect that the two approaches may naturally speak to different audiences interaction designers and software designers who will need to work together on large projects as part of multidisciplinary teams. We also suggest that they also naturally speak about different kinds of issues. User-centred design naturally emphasises the human and environmental factors that impact on a system, but typically has less to say about the detailed design of the software itself or of underlying infrastructures such as wireless networking and sensing. Conversely, formal modelling naturally captures the functional operation of software and hardware, but may struggle to account for the complexities of human and environmental behaviours.

What is required then are new ways of bringing the different approaches together so that interaction and system designers can work with both formal methods and user-centred approaches when designing complex ubiquitous systems. Our article offers a contribution in the overall approach to modelling based upon four key perspectives. We propose that two of these: the Computational World and Technology perspectives, may be broadly characterised as being concerned with the design of the system itself. They require design teams to engage with the fine details of software and hardware design and so are perhaps where formal modelling most comfortably sits. The other two: the Human and Physical World perspectives, are broadly concerned with the situated use of the system in the "real world." They require design teams to reason about the behaviours of the non-digital elements of the overall system, i.e., the behaviours of people. User-centred approaches are perhaps most at home in these perspectives. Our overall approach encourages teams to explore systematically and, ultimately, integrate all four perspectives into a single formal model as part of a holistic view of the system, its users and their environment.

Our second broad reflection is on the nature of the diagrammatic and algebraic notations: the diagrammatic form of bigraphs can play an especially important role in bridging between our four perspectives, and hence, between formal modelling and user-centred design. It was certainly the case in developing the model reported here that diagrammatic representations of bigraphs proved useful for our own interdisciplinary team to develop and communicate ideas. All team members were able to engage with the diagrams, while only one took responsibility for the algebraic manipulations. We, therefore, suggest that a key motivation for including diagrammatic representations alongside algebraic ones is to provide a *lingua franca*, or a boundary object [Star and Griesemer 1989] around which different disciplines can establish common ground. This mirrors previous studies of formal modelling into HCI design that emphasized the need to enrich terse analytic representations [Bellotti et al. 1995].

It is, however, an open question as to whether the diagrammatic representation can significantly replace the algebraic one (now, or in the future) to become the primary means through which the formalism is manipulated. This would require addressing a number of further challenges, not least the challenge of scale. While there have been numerous attempts to develop visual programming and specification languages of various kinds (e.g., LabView, Petri nets, SDL, Statecharts), the scalability of diagrammatic notations is always a major challenge [Green and Petre 1996].

Scalability challenges for bigraphs are focussed on the *number* of reactions rules in a model, and the (graphical) complexity or *detail* of the left-hand sides and right-hand sides of the rules, i.e., the numbers of regions and the detail of entities and links within each region. We note that in the model presented here, the number of rules does not depend on the number of players, nor on the number of prey: there are only 12 (parameterised) reaction rules, expressed in 3 pages, including controls. Regions and sites are key visual abstractions in the graphical form, and in our approach, there is, more or less, a one to one correspondence between perspectives and regions (note, in the Savannah example, the human perspective does not introduce a new region but a new control into the region that models the physical world). The bigraphical concepts of containment and patterns, which we employed here, and the use of priorities for rules, as employed in Calder and Sevegnani [2014], may offer some traction for introducing more scalable structure in the graphical form, for example, with respect to techniques such as zoomable interfaces [Bederson and Hollan 1994].

While the graphical form is a good communication mechanism, the algebraic form is especially useful for type checking and debugging. For example, implementing our model in BigraphER helped us to uncover several basic programming errors. The compiler can return, automatically, a graphical representation of any algebraic form, though a hand-crafted representation is usually more visually appealing. We can also use BigraphER for real-time, online verification/monitoring: generating sequences of bigraphs, according to the rewrite rules that correspond to events. Trials on synthetic and actual event data for a network management application [Calder et al. 2014] indicated the slowest update in a bigraph sequence was just under 0.10s.

While it may be that future research can establish more scalable visual representations of bigraphs, for the time being, we see the immediate value of the diagrammatic form of bigraphs supporting inter-disciplinary discussions of the fine details of a model rather than serving as a general replacement for the algebraic expression and manipulation.

We finish by briefly noting two wider long-term challenges for the kind of formal modelling that we are proposing here. The first concerns the potential of formal modelling to help inspire new designs rather than explaining existing problems. Our work here shows how analysis of a formal model may help suggest refinements to a given system, but it is less clear as to whether it can inspire new ones. Of course, the same charge has been levelled at the HCI theory in general, for examples, in recent calls to develop more "strong concepts" as forms of intermediate design knowledge that bridge between specific instances and the generalized theory and that can help generate designs [Höök and Löwgren 2012]. Other practice-based methods also suffer from the same limitation: user-testing, for example, may reveal problems with a system, but it is less clear how it inspires radical new designs. Thus, we return to the argument that formal modelling needs to be integrated with the panoply of HCI's methods including user-testing, ethnographic studies, and inspirational design approaches, rather than being seen as an entirely different paradigm that is somehow in opposition to them. Our second challenge concerns that of scale; how can our approach scale to real-world systems that are far more complex that Savannah? This indeed is a challenge for future research. However, we note there are examples elsewhere of formal techniques being applied to the design of large scale and/or critical industrial computing systems [Newcombe et al. 2015; Ball et al. 2004]. Conversely, the charge of scalability might be levelled at many other techniques within HCI: how do ethnographic studies or user-tests scale? In short, dealing with large-scale real-world systems is certainly a challenge for formal modelling, as it is for HCI in general. We hope that our work here has helped addressed some of the "real-world" challenges; the question of scale is one for future research.

14. CONCLUSIONS

This article had its genesis in an endeavour to explore how HCI might engage with Formal Computing to respond to and help tackle the significant challenges raised by the emergence of ubiquitous computing. The foundational challenge has been to consider how we can formally reason about a world in which computation becomes embedded into physical environments and in which human social interactions are increasingly mediated by sensors and mobile devices. The formal modelling of Savannah represents a foothill project within this overall challenge. By bringing together formal Computer Scientists with those working in Human-Computer Interaction to iteratively develop a formal model of Savannah using bigraphs, we have been able to

(1) demonstrate that it is possible to formally model the complexities of a real-world ubiquitous system, albeit a relatively simple one;
(2) show that bigraphs, with appropriate extensions, offers a tractable and potentially powerful formalism for modelling ubiquitous systems, at least in part because its formalisms treat space—an inherent aspect of many ubiquitous systems—as a first-class entity in the modelling process;
(3) propose key extensions to the basic bigraphs approach to extend its power to model ubiquitous systems including introducing bigraphical patterns to help analyse invariant properties in a design, and showing how bigraphs can model probabilistic events including some behaviours of sensing systems and of people.
(4) illustrate how such a formal model can be used to account for interactional problems that were observed in a previous user study of Savannah;
(5) demonstrate how the model might be used to predict that such problems might arise, based on a formal verification of the incompleteness of the overall system design (including the humans, physical environments, and sensing technologies that are also part of the overall system).
(6) show how the formal model of Savannah might be made verifiably complete through the addition of new bigraphical reaction rules, which in turn, suggest potential directions for the re-design of the system.
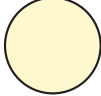
Looking beyond these specific contributions, we have also established an overall approach to modelling based on explicitly considering different perspectives on a ubiquitous system—Computational, Physical, Human, and Technology—and exploring how these interact with one another. Modelling each perspective involves turning to theories and concepts from other disciplines, for example, proxemics from cultural anthropology, in order to generate appropriate abstractions.

And yet the Savannah game remains a relatively simple example of a ubiquitous computing system. Extensive future work is required to explore richer examples, develop more sophisticated models, show how these may be used in practice, generate new abstractions from appropriate theories, and maybe further extend the formalisms and diagrammatic expression of the bigraphs approach. So far, modelling Savannah with bigraphs has led us into a stimulating and we believe unusual debate between what are all too often disconnected fields in our discipline. It is one that we hope others will join as we move forward.

APPENDIX

A. CONTROLS AND BIGRAPHICAL REACTION RULES

Table IX. Controls for the Savannah Game

Control	Arity	Description	Graphical Notation
locale	1	Locale	
localeattack	1	Locale with ongoing attack	
aura	1	Aura	
area$_A$	1	Area containing all coordinates in A	
lion	1	Lion	
lionattack	2	Lion initiating an attack	
liongroup	2	Lion in a group	
impala	1	Impala	
impalaseen	2	Impala seen by a lion	
impalaheld	2	Impala held by a group of lions	
field	0	Field	
$[xy]$	1	GPS position with coordinates x, y	
child	1	Child	
childattack	1	Child initiating an attack	
childgroup	1	Child in a group	

(a) `attack`: a child/lion pair initiates an attack on an impala in its locale (⟶), `timeout1`: idle when the attack timer expires (⟵).

(b) `join1`: a child/lion pair joins the attack initiated by a lion in its locale.

(c) `timeout2`: a group of two child/lion pairs becomes idle when the attack timer expires.

(d) `join2`: a child/lion pair joins a group in its locale.

(e) `kill`: a group of children/lions kill their impala.
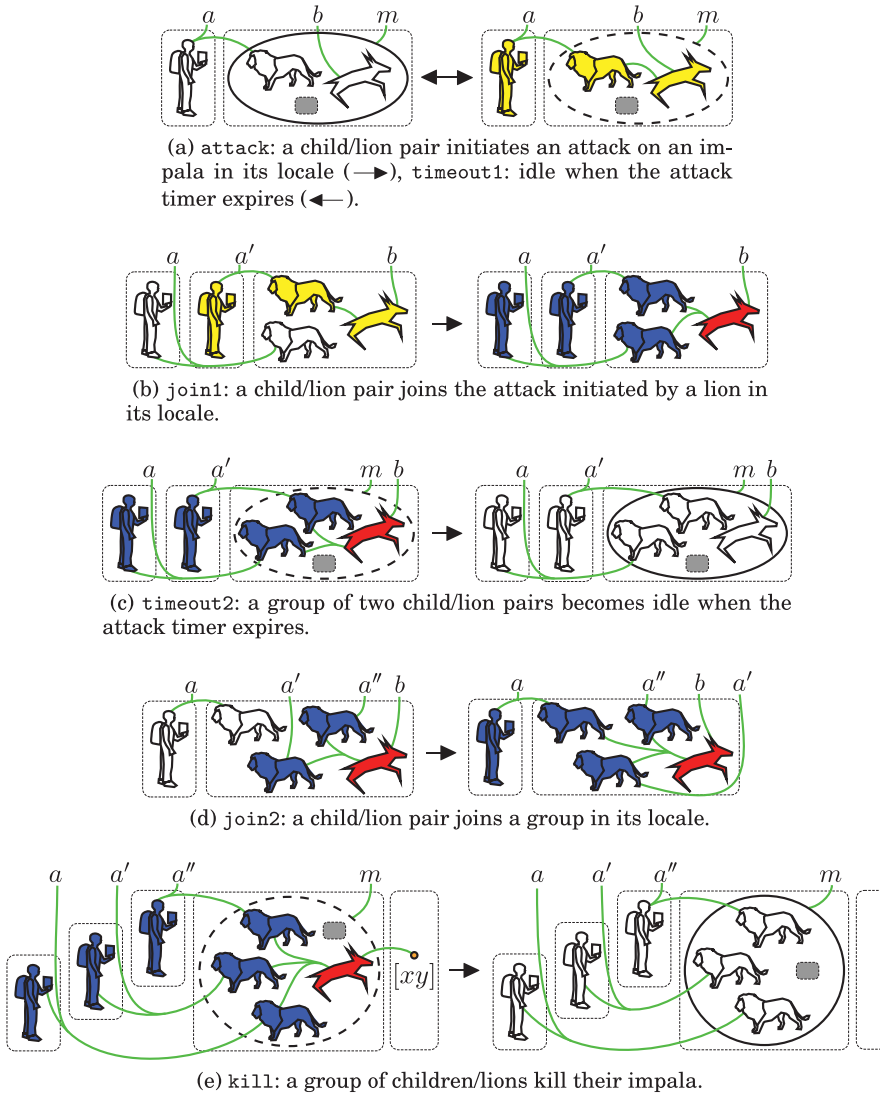
Fig. 43.   Reaction rules for the Savannah game.

(a) enter-l: a child/lion pair enters the game.

(b) enter-i: an impala enters the game.

(c) mov1: a lion moves within the same locale: $[xy], [x'y'] \in A$.

(d) mov2: a lion moves to a different locale: $[xy] \in A$ and $[x'y'] \in A'$.

(e) aura-join: overlapping auras are created ($\longrightarrow$), aura-disjoin: uncoupled ($\longleftarrow$).
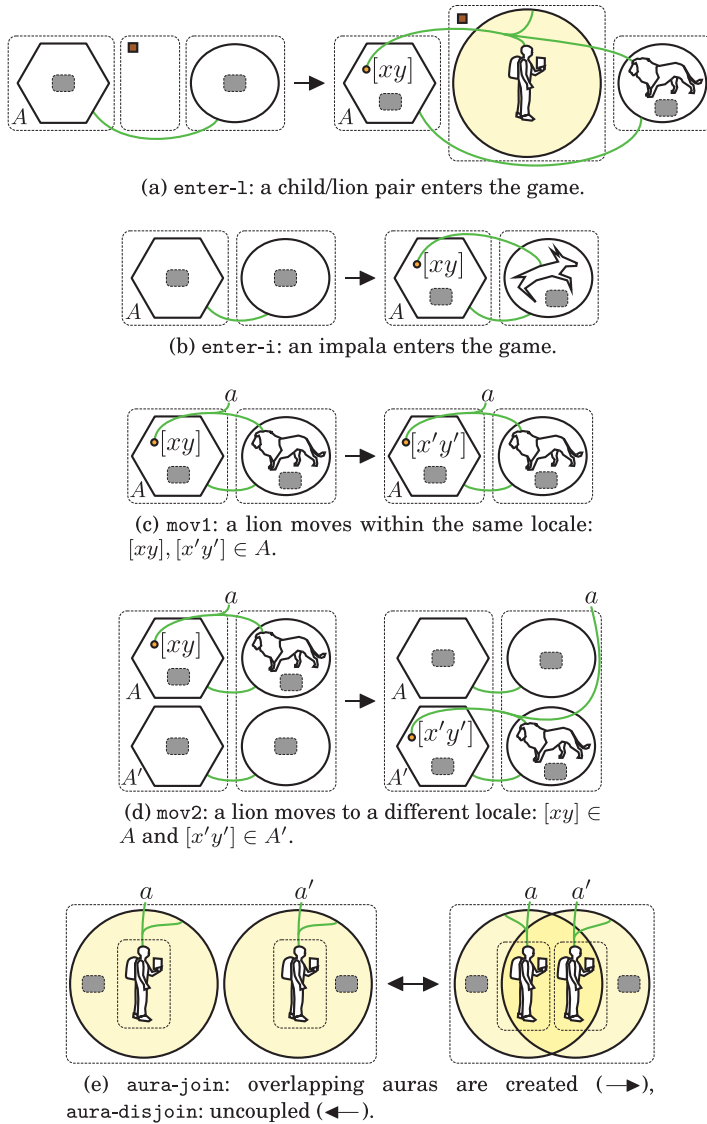
Fig. 44. Reaction rules for GPS events. Similar rules with entities of control localeattack in place of control locale are not shown.

B. ALGEBRAIC DEFINITIONS

Refer to Sevegnani and Calder [2015] for a complete account on the algebraic form for bigraphs with sharing. All the controls listed in Table IX except for locale, localeattack, aura, and area$_A$ are *atomic*.[14] The algebraic form of the reaction rules given in Figure 43 is as follows: selectfont

$$\texttt{attack} \stackrel{\text{def}}{=} \text{child}_a \parallel \text{locale}_m.(\text{lion}_a \mid \text{impala}_b \mid \text{id}_1)$$
$$\longrightarrow /\!l \, (\text{childattack}_a \parallel \text{localeattack}_m . (\text{lionattack}_{a,l} \mid \text{impalaseen}_{b,l} \mid \text{id}_1))$$

$$\texttt{timeout1} \stackrel{\text{def}}{=} /\!l \, (\text{childattack}_a \parallel \text{localeattack}_m.(\text{lionattack}_{a,l} \mid \text{impalaseen}_{b,l} \mid \text{id}_1))$$
$$\longrightarrow \text{child}_a \parallel \text{locale}_m.(\text{lion}_a \mid \text{impala}_b \mid \text{id}_1)$$

$$\texttt{join1} \stackrel{\text{def}}{=} /\!l \, (\text{child}_a \parallel \text{childattack}_{a'} \parallel (\text{lion}_a \mid \text{lionattack}_{a',l} \mid \text{impalaseen}_{b,l}))$$
$$\longrightarrow /\!l \, (\text{childgroup}_a \parallel \text{chilgroup}_{a'} \parallel (\text{liongroup}_{a,l} \mid \text{liongroup}_{a',l} \mid \text{impalaheld}_{b,l}))$$

$$\texttt{timeout2} \stackrel{\text{def}}{=} /\!l \, (\text{childgroup}_a \parallel \text{chilgroup}_{a'}$$
$$\parallel \text{localeattack}_m.(\text{liongroup}_{a,l} \mid \text{liongroup}_{a',l} \mid \text{impalaheld}_{b,l} \mid \text{id}_1))$$
$$\longrightarrow \text{child}_a \parallel \text{child}_{a'} \parallel \text{locale}_m.(\text{lion}_a \mid \text{lion}_{a'} \mid \text{impala}_b \mid \text{id}_1)$$

$$\texttt{join2} \stackrel{\text{def}}{=} /\!l \, (\text{child}_a \parallel (\text{lion}_a \mid \text{liongroup}_{a',l} \mid \text{liongroup}_{a'',l} \mid \text{impalaheld}_{b,l}))$$
$$\longrightarrow /\!l \, (\text{childgroup}_a \parallel (\text{liongroup}_{a,l} \mid \text{liongroup}_{a',l} \mid \text{liongroup}_{a'',l} \mid \text{impalaheld}_{b,l}))$$

$$\texttt{kill} \stackrel{\text{def}}{=} /b \, /\!l \, (\text{childgroup}_a \parallel \text{childgroup}_{a'} \parallel \text{childgroup}_{a''}$$
$$\parallel \text{localeattack}_m.(\text{liongroup}_{a,l} \mid \text{liongroup}_{a',l} \mid \text{liongroup}_{a'',l}$$
$$\mid \text{impalaheld}_{b,l} \mid \text{id}_1)$$
$$\parallel [xy]_b)$$
$$\longrightarrow \text{child}_a \parallel \text{child}_{a'} \parallel \text{child}_{a''} \parallel \text{locale}_m.(\text{lion}_a \mid \text{lion}_{a'} \mid \text{lion}_{a''} \mid \text{id}_1) \parallel 1.$$

The algebraic form of the reaction rules given in Figure 44 is

$$\texttt{enter-l} \stackrel{\text{def}}{=} /m \, (\text{area}_{A,m} \parallel \text{field} \parallel \text{locale}_m)$$
$$\longrightarrow /a \, /m \, (\text{area}_{A,m}.([xy]_a \mid \text{id}_1) \parallel (\text{field} \mid \text{aura}_a.\text{child}_a) \parallel \text{locale}_m.(\text{lion}_a \mid \text{id}_1))$$

$$\texttt{enter-i} \stackrel{\text{def}}{=} /m \, (\text{area}_{A,m} \parallel \text{locale}_m)$$
$$\longrightarrow /b \, /m \, (\text{area}_{A,m}.([xy]_b \mid \text{id}_1) \parallel \text{locale}_m.(\text{impala}_b \mid \text{id}_1))$$

$$\texttt{mov1} \stackrel{\text{def}}{=} /m \, (\text{area}_{A,m}.([xy]_a \mid \text{id}_1) \parallel \text{locale}_m.(\text{lion}_a \mid \text{id}_1))$$
$$\longrightarrow /m \, (\text{area}_{A,m}.([x'y']_a \mid \text{id}_1) \parallel \text{locale}_m.(\text{lion}_a \mid \text{id}_1)) \quad \text{with } [xy], [x'y'] \in A$$

$$\texttt{mov2} \stackrel{\text{def}}{=} /m \, /m' \, ((\text{area}_{A,m}.([xy]_a \mid \text{id}_1) \mid \text{area}_{A',m'}) \parallel (\text{locale}_m.(\text{lion}_a \mid \text{id}_1) \mid \text{locale}_{m'}))$$
$$\longrightarrow /m \, /m' \, ((\text{area}_{A,m} \mid \text{area}_{A',m'}.([x'y']_a \mid \text{id}_1))$$
$$\parallel (\text{locale}_m \mid \text{locale}_{m'}.(\text{lion}_a \mid \text{id}_1))) \quad \text{with } [xy] \in A \text{ and } [x'y'] \in A'$$

---

[14]Entities of atomic controls may not contain other entities.

$$\texttt{aura-join} \stackrel{\text{def}}{=} \text{share} \,(\text{id}_1 \parallel \text{child}_a \parallel \text{child}_{a'} \parallel \text{id}_1) \,\text{by}\, \phi \,\text{in}\, ((\text{aura}_a \mid \text{aura}_{a'}) \parallel \text{id}_{2,\{a,a'\}})$$
$$\longrightarrow \text{share} \,(\text{id}_1 \parallel \text{child}_a \parallel \text{child}_{a'} \parallel \text{id}_1) \,\text{by}\, \psi \,\text{in}\, ((\text{aura}_a \mid \text{aura}_{a'}) \parallel \text{id}_{2,\{a,a'\}})$$

$$\texttt{aura-disjoin} \stackrel{\text{def}}{=} \text{share} \,(\text{id}_1 \parallel \text{child}_a \parallel \text{child}_{a'} \parallel \text{id}_1) \,\text{by}\, \psi \,\text{in}\, ((\text{aura}_a \mid \text{aura}_{a'}) \parallel \text{id}_{2,\{a,a'\}})$$
$$\longrightarrow \text{share} \,(\text{id}_1 \parallel \text{child}_a \parallel \text{child}_{a'} \parallel \text{id}_1) \,\text{by}\, \phi \,\text{in}\, ((\text{aura}_a \mid \text{aura}_{a'}) \parallel \text{id}_{2,\{a,a'\}})$$

$$\phi \stackrel{\text{def}}{=} [\{0\}, \{0,2\}, \{1,3\}, \{1\}]$$
$$\stackrel{\text{def}}{=} (join \parallel join \parallel \text{id}_2)(\text{id}_2 \parallel (\gamma_{1,2} \parallel \text{id}_1)(\text{id}_2 \parallel \gamma_{1,1}))(\text{id}_1 \parallel split \parallel split \parallel \text{id}_1)$$
$$\psi \stackrel{\text{def}}{=} [\{0\}, \{0,1,2\}, \{0,1,3\}, \{1\}]$$
$$\stackrel{\text{def}}{=} (merge_3 \parallel merge_3 \parallel \text{id}_2)(\text{id}_2 \parallel (\gamma_{1,1} \parallel \text{id}_1 \parallel \gamma_{1,1} \parallel \text{id}_1)(\text{id}_1 \parallel \gamma_{1,2} \parallel \gamma_{1,1}))$$
$$(\text{id}_1 \parallel split_3 \parallel split_3 \parallel \text{id}_1).$$

AUTHOR STATEMENT

This article builds on previously published user study of a ubiquitous computing game called Savannah that was published at the CHI 2005 conference [Benford et al. 2005]. That paper is referenced and also briefly summarised in this article, which then goes on to develop a formal mathematical model of the game that accounts for the problems that were observed. This model—which is the main part of this article—has not been previously published. The model builds on an approach called bigraphs. Previous papers have extended the formalism to allow shared locations [Sevegnani and Calder 2015] and applied it to the modelling and analysis of a wireless communication protocol [Calder and Sevegnani 2014] and policies for home networks [Calder et al. 2014]. This article develops those ideas further to define a general model for ubiquitous systems consisting of four perspectives and a novel invariant analysis technique based on bigraphical patterns.

ACKNOWLEDGMENTS

REFERENCES

Gregory D. Abowd and Alan J. Dix. 1992. Giving undo attention. *Interact. Comput.* 4, 3, 317–342.

H. Alexander. 1987. *Formally-Based Tools and Techniques for Human-Computer Dialogues*. Ellis Horwood Limited.

Oana Andrei, Muffy Calder, Matthew Higgs, and Mark Girolami. 2014. Probabilistic model checking of DTMC models of user activity patterns. In *Quantitative Evaluation of Systems*, Gethin Norman and William Sanders (Eds.). Lecture Notes in Computer Science, Vol. 8657. Springer International Publishing, 138–153. DOI:http://dx.doi.org/10.1007/978-3-319-10696-0_11

Thomas Ball, Byron Cook, Vladimir Levin, and Sriram K. Rajamani. 2004. SLAM and static driver verifier: Technology transfer of formal methods inside Microsoft. In *Integrated Formal Methods*, Eerke A. Boiten, John Derrick, and Graeme Smith (Eds.). Lecture Notes in Computer Science, Vol. 2999. Springer, Berlin, Germany, 1–20. DOI:http://dx.doi.org/10.1007/978-3-540-24756-2_1

Benjamin B. Bederson and James D. Hollan. 1994. Pad++: a zooming graphical interface for exploring alternate interface physics. In *Proceedings of the 7th Annual ACM Symposium on User Interface Software and Technology (UIST'94)*. ACM, New York, NY, 17–26. DOI:http://dx.doi.org/10.1145/192426.192435

Victoria Bellotti, Simon Buckingham Shum, Allan MacLean, and Nick Hammond. 1995. Multidisciplinary modelling in HCI design . . .in theory and in practice. In *Proceedings of the SIGCHI Conference on*

Human Factors in Computing Systems (CHI'95). ACM Press/Addison-Wesley Publishing Co., New York, NY, 146–153. DOI:http://dx.doi.org/10.1145/223904.223923

Steve Benford, Duncan Rowland, Martin Flintham, Adam Drozd, Richard Hull, Josephine Reid, Jo Morrison, and Keri Facer. 2005. Life on the edge: supporting collaboration in location-based experiences. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI'05)*. ACM, New York, NY, 721–730.

M. L. Bolton, E. J. Bass, and R. I. Siminiceanu. 2013. Using formal verification to evaluate human-automation interaction, a review. *IEEE Trans. Syst. Man Cybern. A, Syst. Humans* 99, 1–16.

Pascal Bruegger, Denis Lalanne, Agnes Lisowska, and Béat Hirsbrunner. 2009. Tools for designing and prototyping activity-based pervasive applications. In *Proceedings of the 7th International Conference on Advances in Mobile Computing and Multimedia (MoMM'09)*. ACM, New York, NY, 129–136. DOI:http://dx.doi.org/10.1145/1821748.1821778

M. Calder. 1998. What use are formal design and analysis methods to telecommunications services? In *Feature Interactions in Telecommunications and Software Systems*, K. Kimbler and L. G. Bouma (Eds.), Vol. V. IOS Press, 23–31.

Muffy Calder, Alexandros Koliousis, Michele Sevegnani, and Joseph Sventek. 2014. Real-time verification of wireless home networks using bigraphs with sharing. *Sci. Comput. Program.* 80, 288–310.

Muffy Calder and Michele Sevegnani. 2014. Modelling IEEE 802.11 CSMA/CARTS/CTS with stochastic bigraphs with sharing. *Formal Aspects Comput. Sci. (FACS)* 26, 3, 537–561.

Stuart K. Card, Thomas P. Moran, and Allen Newell. 1980. The keystroke-level model for user performance time with interactive systems. *Commun. ACM* 23, 7 (Jul. 1980), 396–410. DOI:http://dx.doi.org/10.1145/358886.358895

Giovanni Conforti, Damiano Macedonio, and Vladimiro Sassone. 2005. Spatial logics for bigraphs. In *Automata, Languages and Programming*. Lecture Notes in Computer Science, Vol. 3580. Springer, Berlin, Germany, 766–778.

Alan J. Dix and Colin Runciman. 1985. Abstract models of interactive systems. *People and Computers: Designing the Interface*, P. J. and S. Cook (Eds.). Cambridge University Press, 13–22.

Paul Dourish. 2004. What we talk about when we talk about context. *Pers. Ubiquitous Comput.* 8, 1, 19–30. DOI:http://dx.doi.org/10.1007/s00779-003-0253-8

Paul Dourish. 2006. Re-space-ing place: "Place" and "space" ten years on. In *Proceedings of the 2006 20th Anniversary Conference on Computer Supported Cooperative Work (CSCW'06)*. ACM, New York, NY, 299–308.

Emmanuel Dubois, Christophe Bortolaso, Damien Appert, and Guillaume Gauffre. 2014. An MDE-based framework to support the development of mixed interactive systems. *Sci. Comput. Program.* 89, 199–221. DOI:http://dx.doi.org/10.1016/j.scico.2013.03.007

Emmanuel Dubois and Philip D. Gray. 2007. A design-oriented information-flow refinement of the ASUR interaction model. In *Engineering Interactive Systems - EIS 2007 Joint Working Conferences, EHCI 2007, DSV-IS 2007, HCSE 2007, Salamanca, Spain, March 22-24, 2007. Selected Papers*. 465–482. DOI:http://dx.doi.org/10.1007/978-3-540-92698-6_28

S. Dupuy-Chessa, G. Godet-Bar, J.-L. Prez-Medina, D. Rieu, and D. Juras. 2010. A software engineering method for the design of mixed reality systems. In *The Engineering of Mixed Reality Systems*, Emmanuel Dubois, Philip Gray, and Laurence Nigay (Eds.). Springer, London, 313–334. DOI:http://dx.doi.org/10.1007/978-1-84882-733-2_16

D. Gentner and A. L. Stevens. 1983. *Mental Models*. Taylor & Francis.

Jeremy Gow and Harold Thimbleby. 2005. MAUI: an interface design tool based on matrix algebra. In *Computer-Aided Design of User Interfaces IV*. Springer, 81–94.

Thomas R. G. Green and Marian Petre. 1996. Usability analysis of visual programming environments: A 'cognitive dimensions' framework. *J. Vis. Lang. Comput.* 7, 2, 131–174. DOI:http://dx.doi.org/10.1006/jvlc.1996.0009

Saul Greenberg. 2011. Opportunities for proxemic interactions in ubicomp. In *Proceedings of the 13th IFIP TC 13 International Conference on Human-Computer Interaction - Volume Part I (INTERACT'11)*. Springer-Verlag, Berlin, Germany, 3–10.

Chris Greenhalgh and Steven Benford. 1995. MASSIVE: A collaborative virtual environment for teleconferencing. *ACM Trans. Comput.-Hum. Interact.* 2, 3 (Sep. 1995), 239–261. DOI:http://dx.doi.org/10.1145/210079.210088

Jonathan Grudin. 1990. The computer reaches out: The historical continuity of interface design. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 261–268.

Edward T. Hall. 1966. *The Hidden Dimension*. Anchor Books, New York.

Bill Hillier and Julienne Hanson. 1984. *The Social Logic of Space*. Cambridge University Press. http://dx.doi.org/10.1017/CBO9780511597237 Cambridge Books Online.

Charles Antony Richard Hoare. 1978. Communicating sequential processes. *Commun. ACM* 21, 8, 666–677.

Kristina Höök and Jonas Löwgren. 2012. Strong concepts: Intermediate-level knowledge in interaction design research. *ACM Trans. Comput.-Hum. Interact.* 19, 3, Article 23 (Oct. 2012), 18 pages. DOI:http://dx.doi.org/10.1145/2362364.2362371

Edwin Hutchins. 1995. *Cognition in the Wild*. MIT press.

Bonnie E. John and David E. Kieras. 1996. The GOMS family of user interface analysis techniques: Comparison and contrast. *ACM Trans. Comput.-Hum. Interact.* 3, 4 (Dec. 1996), 320–351. DOI:http://dx.doi.org/10.1145/235833.236054

Jean Krivine, Robin Milner, and Angelo Troina. 2008. Stochastic bigraphs. *Electron. Notes Theoretical Comput. Sci.* 218, 73–96.

Paul Marshall, Yvonne Rogers, and Nadia Pantidi. 2011. Using F-formations to analyse spatial patterns of interaction in physical environments. In *Proceedings of the ACM 2011 Conference on Computer Supported Cooperative Work (CSCW'11)*. ACM, New York, NY, 445–454.

Paolo Masci, Paul Curzon, Dominic Furniss, and Ann Blandford. 2015. Using PVS to support the analysis of distributed cognition systems. *Innovations Syst. Softw. Eng.* 11, 2, 113–130. DOI:http://dx.doi.org/10.1007/s11334-013-0202-2

Robin Milner. 2009. *The Space and Motion of Communicating Agents*. Cambridge University Press.

Thomas P. Moran. 1981. The command language grammar: A representation for the user interface of interactive computer systems. *Int. J. Man-Mach. Stud.* 15, 1, 3–50.

Chris Newcombe, Tim Rath, Fan Zhang, Bogdan Munteanu, Marc Brooker, and Michael Deardeuff. 2015. How Amazon web services uses formal methods. *Commun. ACM* 58, 4 (Mar. 2015), 66–73. DOI:http://dx.doi.org/10.1145/2699417

Jakob Nielsen. 1986. A virtual protocol model for computer-human interaction. *Int. J. Man-Mach. Stud.* 24, 3, 301–312.

Donald A. Norman. 1983. Some observations on mental models. *Mental Models* 7, 112, 7–14.

Donald A. Norman. 1988. *The Psychology of Everyday Things*. Basic books.

Sam Owre, Sreeranga Rajan, John M. Rushby, Natarajan Shankar, and Mandayam Srivas. 1996. PVS: combining specification, proof checking, and model checking. In *Computer Aided Verification*. Springer, 411–414.

David Lorge Parnas. 2008. Connecting good theory to good practice: Software documentation: A case study. In *Software Engineering 2008. Fachtagung des GI-Fachbereichs Softwaretechnik, 18.-22.2.2008 in München*. 17–20. http://subs.emis.de/LNI/Proceedings/Proceedings121/article1978.html.

Yvonne Rogers. 2012. HCI Theory: classical, modern, and contemporary. *Synth. Lect. Hum.-Centered Inform.* 5, 2, 1–129.

J. Rushby. 2007. What use is verified software? In *12th IEEE International Conference on Engineering Complex Computer Systems*. IEEE, 270–276. DOI:http://dx.doi.org/10.1109/ICECCS.2007.47

Albrecht Schmidt, Michael Beigl, and Hans W. Gellersen. 1998. There is more to context than location. *Comput. Graph.* 23, 893–901.

Michele Sevegnani and Muffy Calder. 2015. Bigraphs with sharing. *Theoretical Comput. Sci.* 577, 43–73. DOI:http://dx.doi.org/10.1016/j.tcs.2015.02.011

Simon Buckingham Shum and Nick Hammond. 1994. Transferring HCI modelling and design techniques to practitioners: a framework and empirical work. In *Proceedings of HCI '94, People and Computers IX, Glasgow*. (Aug. 1994), Cambridge University Press, 21–36.

Jeffrey Star and John Estes. 1990. *Geographic Information Systems: An Introduction*. Prentice Hall, Englewood Cliffs, New Jersey.

Susan Leigh Star and James R. Griesemer. 1989. Institutional ecology, 'translations' and boundary objects: Amateurs and professionals in Berkeley's museum of vertebrate zoology, 1907-39. *Soc. Stud. Sci.* 19, 3, 387–420.

Lucy A. Suchman. 1987. *Plans and Situated Actions: The Problem of Human-Machine Communication*. Cambridge University Press.

Bernard Sufrin. 1982. Formal specification of a display-oriented text editor. *Sci. Comput. Program.* 1, 3, 157–202.

H. Thimbleby. 2015. Safer user interfaces: A case study in improving number entry. *IEEE Trans. Softw. Eng.* 41, 7 (Jul. 2015), 711–729. DOI:http://dx.doi.org/10.1109/TSE.2014.2383396

Peter C. Wright, Bob Fields, and Michael D. Harrison. 1996. Distributed information resources: a new approach to interaction modelling. In *Proceedings of ECCE8: 8th European Conference on Cognitive Ergonomics*. European Association of Cognitive Ergonomics, 10–13.

Peter C. Wright, Robert E. Fields, and Michael D. Harrison. 2000. Analyzing human-computer interaction as distributed cognition: the resources model. *Hum.-Comput. Interact.* 15, 1, 1–41.