A Prototype Agent-based Model of Antimalarial Drug Compliance

Jez McKean, Sebastian Rahlf, Hayden Shorter, Guy Barnish, Ray Paton, Peter McBurney, Dyfrig Hughes, and Ian Hastings.

Jez McKean, Carbon360 Ltd, Liverpool, UK.
Sebastian Rahlf, formerly of Department of Computer Science, University of Liverpool, UK.
Hayden Shorter, AePona Ltd, Belfast, Northern Ireland, UK.
Guy Barnish, Liverpool School of Tropical Medicine, Liverpool, UK.
Ray Paton, formerly of Department of Computer Science, University of Liverpool, UK.
* Peter McBurney, Department of Computer Science, University of Liverpool, UK.
Dyfrig Hughes, Centre for the Economics of Health, IMSCaR, University of Wales – Bangor, Wales, UK.
Ian Hastings, Liverpool School of Tropical Medicine, Liverpool, UK.

** Corresponding author. Email address: p.j.mcburney@csc.liv.ac.uk*

1. Introduction

Malaria kills around 2 million people per year, predominantly attributable to *P. falciparum* infection of children in sub-Saharan African. Interventions such as the provision of bednets, or mass treatment with antimalarial drugs, which reduce the impact of malaria, generally reduce all-cause childhood mortality by 25-33%, suggesting that indirect malaria mortality may be huge. The main defence against this process is the provision of readily available, affordable and effective antimalarial drugs. Chloroquine was the mainstay for antimalarial drug treatment over the last 30 years but is now largely ineffective in much of Africa. Its replacement, sulphadoxine-pyrimethamine (SP) lasted for less than 5 years before widespread resistance emerged. As a consequence, mortality has risen, so the current priority is to deploy the replacement therapies in a manner that will delay the appearance of resistance as long as possible, and hence maximize their useful therapeutic lifespans. Field trials of new antimalarials can be used to assess their clinical effectiveness, but cannot provide information on how their widespread deployment will stimulate the evolution of resistance. This has been addressed through hypothetical mathematical models.

Malaria is a curable disease, so an obvious question to ask is why do people still die from it. They do so because they are not treated, or because they are treated too late, or because the treatment they receive or they take is not a complete course of medicine. Thus, to understand the diffusion of the disease in a community, we need to understand the human social and cultural context of treatment provision and compliance, along with individual (or family) decision-making behaviours. Such human behaviour, however, is not readily amenable to traditional mathematical representation and modeling, and so these investigations must rely on computer simulation.

In this note, we report the development of a prototype multi-agent system (MAS) simulation model of malaria treatment compliance decision-making in village Africa. Our model shares many features with recent MAS models in marketing, including a focus on individual-level decision-making and the inclusion of a network component to model word-of-mouth or other social influences (eg, see Urban and Hauser 1993). For health policy makers the strategic issues in this domain are very similar to those for commercial enterprises engaged in marketing: how to understand the dynamics of adoption (of a new malaria treatment or a new product) when such individual adoption decisions take place within a social context.

The main objectives for our work in developing a prototype MAS simulation model were (a) to better understand the domain-specific and computational issues involved, and (b) to gain experience of the software engineering challenges involved in a full-scale MAS model of malaria treatment compliance decisions. In this brief note, we describe the main components of our prototype model, and identify the key lessons which the prototype has provided us for software engineering and model use.

2. Model Description

Our prototype MAS model has two key components, a model of the decision-making process of a single household within a village, and a network diffusion model of the social influence of one household in a village on another. We describe each component here.

2.1. Individual household decision model

The first model aimed to identify and represent qualitatively all the factors likely to influence the decision-making process of a carer in deciding upon an antimalarial treatment for a sick child. The model adopted was hierarchical, and took the household as the basic decision-making unit. The household was the basic unit because it was assumed that all children in a household had the same carer; households vary in the number of adults and children, the level of traditional belief (which determines the likelihood of seeking a 'western' drug), and disposable income. Households occur within villages which differ in the distance to health care centre and thus the costs of consultation; moreover, communication with other households in the same village may affect a carer's behaviour. The final level is environmental; this is the epidemiological setting that determines factors such as incidence of fever (a key symptom of malaria), and 'true' malaria incidence rate (i.e. the proportion of fevers that are malarial in origin). The drug availability is a global entity and is assumed to be the same for all households/villages; its attributes are cost, and complexity of the regimen.

A mother faced with a childhood fever has to weigh several considerations in deciding a course of action. In our model these were: Is the fever serious? Whose advice should she take? Is the proposed treatment likely to be effective? Will it have any side-effects? Is a source of treatment close enough to be accessible? Can she afford it? There are numerous factors that determine how she weighs these factors as are shown on Figure 1. Some of these factors may be contradictory, so decision rules needed to be specified to determine how the mother is likely to act. The mother is then assumed to act in five different ways:

1. She may buy an antipyretic drug (i.e. one that reduces fever) that has no antimalarial activity, paracetamol being the obvious example.
2. She may buy part of an antimalarial drug treatment, for example, a single tablet of chloroquine, rather than the full three-day regimen.
3. She may buy the full antimalarial treatment, leading to:
    - 3a. she gives part of the course, saving some tablets 'for the next time', or
    - 3b. she may give the full course.
4. She may take her child to a formal health care facility.

The qualitative decision rules under which a mother assesses the inputs (i.e., how she weighs considerations about the fever) and selects the appropriate action were quite complex, and will be reported in detail elsewhere.

We attempted to develop this model for a plausible sub-Saharan African setting, ostensibly for a hypothetical Kenyan village. In reality, the parameter values came from many different studies and, to the best of our knowledge, have not all been measured in the same place. Social structures, medical provision, etc, vary enormously over Africa and it is obviously unrealistic to describe a "typical'' village. However, the imaginary village appears at least plausible so is sufficient for the current purposes of prototype modelling. The chief assumptions were that transmission is sufficiently intense that only children suffered clinical malaria attacks (because adults are effectively 'immune'), and that we only considered a single first-line drug. Thus the model specifically investigates how a mother utilises the available first-line antimalarial drug in the treatment of childhood fevers.

Figure 1. The factors influencing a carer deciding how to deal with a childhood fever

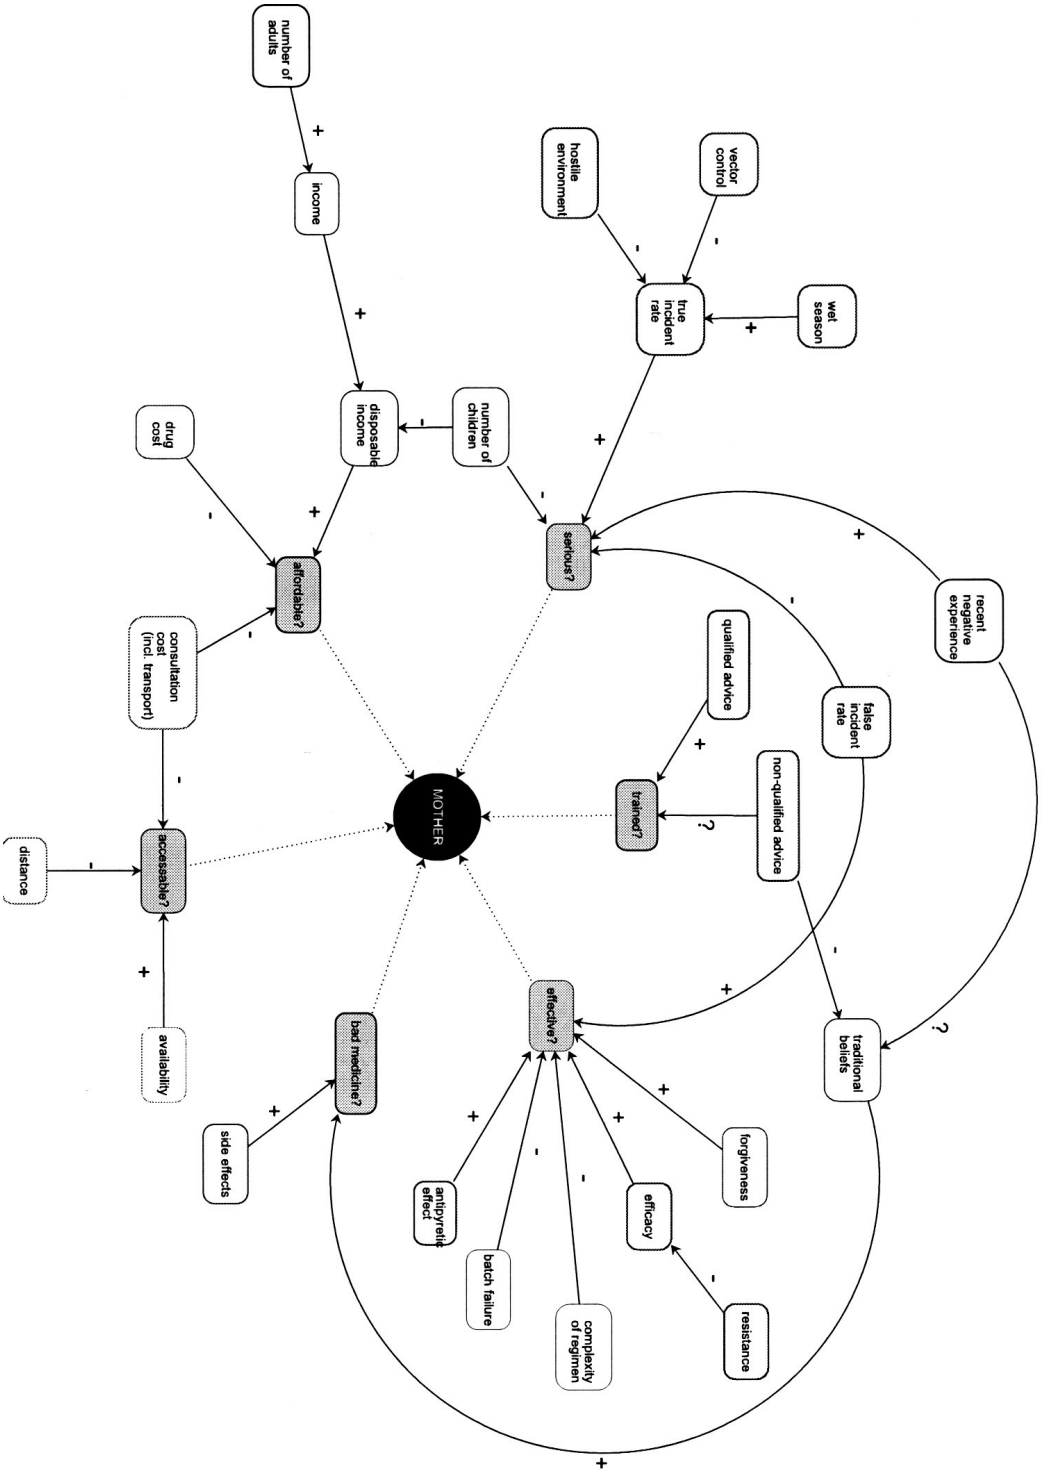

2.2. Network diffusion model

The second component of the prototype MAS model of malaria treatment decision-making was a network model of the diffusion of treatments through a village. Individual households were represented as nodes in the network, shown as circles in the graph of Figure 2. The influence of one household over another household's decision-making is indicated by an arrow from the node representing the first household to the node representing the second. Some households, such as those of traditional community leaders or health-workers may influence a great number of other households, and not be influenced in turn.

Each household was assumed to undertake its malaria treatment decisions according to the Individual Household Decision model described above. For any new treatment, households could experience one of five different levels of trust, from No Trust through to Complete Trust. Different levels of trust are indicated by different nodal colours, and the percentage of households at each level of trust at any time is shown by the table of percentages on the left-hand side of Figure 2.

The model simulates the diffusion of belief in a new treatment through a village network of households. The nature and speed of this diffusion will depend both upon the network topology of influence-links between households (ie, the shape of the graph assumed) as well as community-wide and individual-household-level factors which influence decision making about malaria treatment. An example of such a community-wide factor is the current rate of fevers (including malarial fevers) in the community. An example of a household-level variable is the degree of authority of the household in the village hierarchy. The prototype model allows the user of the model to alter these factors using the sliders shown on the right-hand side of Figure 2. The model was constructed to allow additional factors to be added readily. The prototype model may be run one step at a time, or as quickly as computer processors permit, and use of the model will be demonstrated in the presentation.

Figure 2: Screenshot of Malaria Treatment Diffusion Model

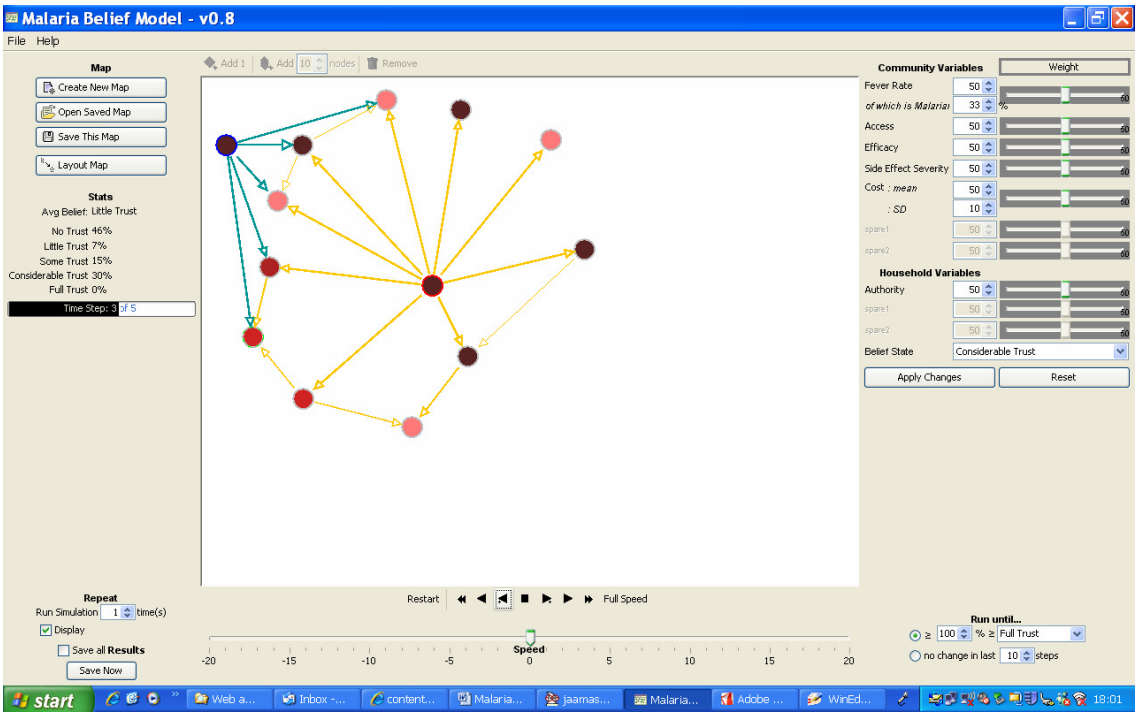

3. Implications for MAS Modellers

The model developed was a prototype, created with the express intention of understanding the likely challenges in engineering a full-scale model. As such, the exercise provided valuable lessons for future development, some of which will be well-known to agent-simulation modellers and some of which may be new. Briefly, these issues were:

1. *Degree of granularity:* In any modelling exercise, design choices are made concerning the level of detail of modelling. For this prototype, we have assumed a single decision-maker within each household, and we have ignored intra-household decision processes. Such processes may well be important, particularly in non-western cultures, e.g. see McBurney 1988.
2. *Treatment of time:* Modelling time is always problematic in multi-agent simulation models. For simplicity, we assumed for the prototype that each household reconsiders its decision at each clock-tick, and immediately informs its downstream neighbours. In addition, these instantaneous decisions and communications take no account of the timescales of malaria infection or development in a patient, or in a community.
3. *Modelling of inter-node communications and influence*: For simplicity, we assumed that influence between nodes was uni-directional, not bi- or multi-directional. We also assumed that inter-node communication was on one level only, ie, that there was no reflective discussion by households about the communications or influence they undertook. In real social networks, such reflective behaviours may greatly influence the beliefs and behaviours of the participants (e.g., see Hardie and MacKenzie 2006).
4. *Model reality versus usability:* The more realistic a model is, typically the more complex it is, and hence the less understandable by those who were not involved in its development. Any modeler has to strike a balance between these two dimensions, and an appropriate trade-off reached. The result may be domain- or use-context-specific.
5. *Model calibration:* As with any agent-based model, model calibration against real-world data is a key challenge. Because published field data on household decision-making regarding malaria treatments is scattered (pertaining to different countries, communities and time points), approximate and incomplete, we have relied on expert judgment, informed by what published literature exists, to develop the Household Decision model.
6. *Model assessment:* As with any simulation model, verification of model fit with the real phenomena it seeks to represent, and assessment of model outputs under what may be hypothetical assumptions, are problematic (see Marks 2007, Midgley *et al.* 2007). We have not addressed these issues with this prototype.

Many of our assumptions here would need to be revisited in full-scale MAS model of malaria treatment decision-making. We hope to do so in future work.

Acknowledgments:

The authors gratefully acknowledge financial support received from the Wellcome Trust under a VIP Award to the University of Liverpool and from the DFID-funded *Malaria Knowledge Programme* of the Liverpool School of Tropical Medicine.

References:

I. Hardie and D. MacKenzie [2007]: Assembling an economic actor: the *agencement* of a Hedge Fund. *The Sociological Review*, **55 (1)**: 57—80.

R. E. Marks [2007]: Validating simulation models: a general framework and four applied examples. *Journal of Computational Economics*, In press.

P. J. McBurney [1988]: On transferring statistical techniques across cultures: the Kish Grid. *Current Anthropology,* **29 (2)**: 323-5.

D. F. Midgley, R. E. Marks and D. Kunchamwar [2007]: Building and assurance of agent-based models: an example and challenge to the field. *Journal of Business Research*, **60 (8)**: 884—893.

G. L. Urban and J. R. Hauser [1993]: *Design and Marketing of New Products*. Englewood Cliffs, NJ, USA: Prentice-Hall.