

# Medieval Military Logistics: An Agent-based Simulation of a Byzantine Army on the March

Bart Craenen  
Georgios Theodoropoulos

School of Computer Science  
The University of Birmingham  
United Kingdom  
{ b.g.w.craenen, gkt } @cs.bham.ac.uk

Vincent Gaffney  
Philip Murgatroyd

Institute of Archaeology and  
Antiquity  
University of Birmingham  
United Kingdom  
{v.l.gaffney, psm703}@bham.ac.uk

John Haldon

History Department  
Princeton University  
U.S.A  
jhaldon@Princeton.edu

**Keywords:** Agent-based modelling, distributed simulation, historical studies, medieval, military logistics

## Abstract

Although historical studies are frequently perceived as clear narratives defined by a series of fixed events; in reality, even where critical historical events may be identified, historic documentation frequently lacks corroborative detail to support verifiable interpretation. Consequently, interpretation rarely rises above the level of unproven assertion and is rarely tested against a range of evidence. Agent-based simulation can provide an opportunity to break these cycles of academic claim and counter-claim.

This paper discusses the development of an agent-based simulation designed to investigate medieval military logistics so that new evidence may be generated to supplement existing historical analysis. It uses as a use-case the Byzantine army's march to the battle of Manzikert (AD 1071), a key event in medieval history. The paper focuses primarily on the design of the agents and the environment they interact with, as well as how the agent-based simulation as a whole can be used to generate new parameters with which historical evidence can be situated.

## 1. INTRODUCTION

The analysis of humanities data sets offers considerable challenges to computational science [1]. Historical studies are a good example of the difficulties associated with such research. For many, historical interpretation is associated with clear narratives defined by a series of fixed events or actions. In reality, even where critical historical events may be identified, contemporary documentation frequently lacks corroborative detail that supports verifiable interpretation. Consequently, for many periods and areas of research, interpretation rarely rises above the level of unproven assumptions, rarely or never tested against a range of evidence.

The application of agent-based modelling, as a means to explore the effect of individual action, has recently emerged

as an area of interest for the historical disciplines. Application has, however, been extremely limited both in the number of projects and their scale. Examples such as the analysis of resource exploitation by Mesolithic hunter gatherer groups [2] or Prehispanic settlements in North America [3] have generally involved the analysis of small-scale groups at individual or household level rather than larger societies [4].

This paper presents work as part of the "Medieval Warfare on the Grid" (MWGrid) project. MWGrid seeks to study behaviour dynamics at a larger scale, involving tens of thousands of agents, all within the context of modelling logistical arrangements relating to the march of the Byzantine army to the battle of Manzikert (AD 1071). The battle of Manzikert was a key event in Byzantine history, resulting in the collapse of Byzantine power in central Anatolia [5]. The MWGrid project brings together a series of parallel research developments relating to the analysis of medieval military logistics and the development of a novel computational infrastructure designed to investigate complex systems.

This paper provides an example of how an agent-based model for use in a historical context is created by describing in detail the design and use of the MWGrid agent-based model. The MWGrid model uses the MWGrid infrastructure and framework [6] [7] and can be seen as the single use-case of it. The MWGrid model itself consists of agents and the environment with which these agents interact. Design of both is based on large amounts of interpreted data, both historical and current, and this paper describes how the MWGrid model uses this as well as how it can be used to generate new parameters within which historical evidence can be situated. A short description of the MWGrid infrastructure and framework is provided to facilitate understanding the design of the model.

The paper is then structured as follows: Section 2 provides an overview of the problem domain and summarises the requirements on the MWGrid model. Section 3 presents a brief overview of the MWGrid framework. Section 4 describes the both the agent and the

environment aspects of the MWGrid model itself in detail. The paper concludes with a summary and an outline of future work.

## 2. MEDIEVAL MILITARY LOGISTICS: THE MANZIKERT CAMPAIGN

A key problem in historical studies is to analyse and understand the implications of the need for medieval states to collect and distribute resources to maintain armies. It is apparent from the historical record that these requirements affected all aspects of political organisation and when armies failed the results could prove disastrous to society as a whole. It is also clear that military studies seldom progress past the study of existing texts to bear out the pragmatic consequences of military behaviour, despite its decisive role in shaping pre-modern societies.

Study of the events associated with the Byzantine army's march across Anatolia to the Battle of Manzikert is particularly attractive in this context. This was a major logistical challenge that involved the largest Byzantine army for over 50 years travelling more than 700 miles across what is now part of the modern state of Turkey, from near Constantinople (modern Istanbul) to Manzikert (modern Malazgirt) just north of Lake Van (Figure 1).




Figure 1. Anatolia, extent of ABM and possible route

The army set out from near Constantinople in March of 1071 and arrived at the border fortress of Manzikert in August. The Byzantines' subsequent defeat by the Seljuk Turks and their Sultan Alp Arslan on August 26th was considered so catastrophic that the Byzantine chroniclers dubbed it "the dreadful day" [8], and it was the last time the Byzantine Empire exerted even minimal control over the whole of Anatolia.

Considering how important the battle was to the Byzantine Empire, the Seljuk Turks and the modern republic of Turkey, the historical sources leave several significant gaps. Historical sources are vague, contradictory or absent for such factors as how many people were in the Byzantine army or which route it took.

Understanding the context of how the Byzantine state supported the army and its progress underpins not only our understanding of the likely composition of the force and its capacity for movement, it also supports in some manner how we understand the momentous outcome of the

campaign. The less controversial aspects of the historical record are also useful in providing verifiable details against which to compare the model's outputs.

The MWGrid project aims to use an agent-based model of the Byzantine army's march to Manzikert to investigate the transport of tens of thousands of people, horses and mules along with tons of equipment over 700 miles through the Anatolian summer in order to provide an insight into this pivotal historic event. By modelling different scenarios based on historical records and modern interpretations of how the Byzantine infrastructure supported an army on campaign, we aim to draw valuable conclusions regarding how the transport, taxation, agricultural production and military organisation systems interconnect.

The project centres around agents representing all the members of the army. The commander through to the lowliest servant occupy part of an military structure with one clear goal; to arrive at a destination in a fit state to win a battle. It focuses on the movement and provisioning problems associated with moving tens of thousands of humans and animals across a pre-industrial landscape. This requires the army to be modelled at a 1:1 scale to provide a plausible movement model, requiring tens of thousands of agents. Multiple executions of the agent-based model are required with different numbers of people and animals, different levels of food availability and different types of organisation and route planning. The simulation will record both the state and progress of the army as well as the effects on the communities impacted by the progress of the army.

Clearly the processing power and the memory requirements needed for this simulation far exceed the capabilities of any sequential von Neumann machine. Distributed simulation and the harnessing of distributed computing resources emerge as the only viable approach to deal with a problem of such scale.

## 3. THE MWGRID FRAMEWORK

A high-level view of the MWGrid framework is provided in Figure 2. It can be seen as consisting of two major parts: the simulation system and the analysis environment.

The simulation system executes the model and produces detailed trace files that are fed into the analysis system for off-line post-processing. This is achieved via a range of packages depending on the output required. Statistics can be produced detailing movement rates, food consumption, agent health status, amount of time spent on the move and the state of the environment after the army has moved on. These can involve individual agents or the aggregation of statistics of the whole army or certain subgroups.

It is also possible to import a transformed trace-file into a 3D modelling package so that 3D visualisation can be created automatically. The ability to display the results of

the model in a 3D representation is an important tool in communicating the output of traditional agent-based modelling. 3D representations have their dangers though, as they can convey an artificial sense of authority due to their persuasive nature. The ability to produce different representations of the model's results therefore remains useful.




Figure 2. The MWGrid framework

The simulation system consists of three layers: the agent-based model (ABM); the distributed simulation kernel (PDES-MAS) [9]; and the Middleware as an interface between the two. The Middleware provides clearly described ways in which the ABM interacts with PDES-MAS using a number of specific interfaces. It also provides a support framework for running the ABM, by, for example, providing a Scheduler, and a Distributed Object class that can be extended into an agent [7].

#### 4. THE MODEL

The model consists of two main elements: the environment representing the terrain, infrastructure and resources of Anatolia; and the agents, representing the human and animal members of the Byzantine army on a one-to-one basis.

##### 4.1. The Environment

11th century Anatolia represents a large and rich environment for an agent-based model, the distance from Constantinople to Manzikert is over 700 miles as the crow

flies and contains a variety of different terrain types. Abstracting the right data at an appropriate scale is an important part of preparing the environment data. The environment is represented in a series of slices, each slice dealing with a different aspect of the Anatolian landscape. Each slice can either be an array of values covering the whole of the ABM area or a list of locations with associated values, used for sparse data sets.

Each cell of the environment is approximately  $5m^2$ , giving a total size of  $280,700 \times 88,890$  cells. The area covered is so large that the easiest type of geographical data format to use (ESRI ASCII array) would require a file of around 140Gb per slice. For this reason, two other resolutions are supported within the model for environmental data;  $50m^2$  and  $500m^2$ , leading to much more manageable file sizes of 1.4Gb and 14Mb respectively. Although high resolution is required in the environment to accurately model the effects of crowd movement, such a high resolution is not necessarily required of all the data sets that describe the environment.

As the model is concerned with the practical aspects of feeding and moving the army, only certain characteristics of the historical environment are relevant to the model. These are: terrain, roads, land-use, water, and settlements.

##### 4.1.1. Terrain

The physical geography of Anatolia plays a large part in the process of route planning for the army. The nature of the terrain will affect movement rates as broken ground will force the army to move in narrower columns and tight mountain passes may even require the army to pass through in single file. Terrain also affects the energy expended during movement, thus requiring more food and water to be consumed.

The highest quality terrain data is the ASTER GDEM data (Figure 3), a joint satellite mapping project by NASA and Japan's Ministry of Economy, Trade and Industry [10]. It consists of height data in cells approximately  $30m^2$  which is converted into an ESRI ASCII array with cells  $50m^2$  in order to fit more conveniently into the MWGrid environment. As terrain is a continuous data set, it is stored in a zipped large raster file (1.4Gb unzipped, 275Mb zipped) of  $28,070 \times 8,889$  resolution. As this represents a large amount of data to keep in memory, especially when only a small part is required at any given time, the model splits this data into partitions and loads and discards partitions as required.

As the shape of the Anatolian terrain has not changed significantly in the last 900 years the use of modern data will not significantly affect the accuracy of the model.

##### 4.1.2. Roads

Roads facilitate both faster movement and more convenient route planning. The Byzantine road system was

a degraded version of the earlier Roman infrastructure [11] but it still offered a faster and more reliable option than trekking across country.

Roads are an example of a sparse data set, a more efficient way of storing the data is by recording only the location of cells where roads are present. This can be done at 50m<sup>2</sup> resolution in a file of around 3Mb. The *Tabula Imperii Byzantini* (TIB) maps [12] give locations for roads over the area that they cover but outside this area there are difficulties in extrapolating known data to fill in any gaps. Unlike settlements which are discrete data points, roads are linear features that actually lead from somewhere to somewhere, often by a meandering route. This makes them harder to generate in a plausible manner. This being the case, only major, well known routes are plotted outside the coverage of the TIB maps but, as the army will have tended to stick to major routes, this is anticipated not to greatly affect the model's accuracy.

#### 4.1.3. Settlements

The army was supplied with food and equipment by the communities through whose territories it passed. Any supplies provided were taken into account when the next tax bill was presented. This made settlements places to be visited along the march in order to pick up necessary supplies. They would also have tended to have been situated near water sources and would have served as suppliers of news, local guides, entertainment and souvenirs for the campaigning soldiers. This makes them important to both the route planning and provisioning elements of the model as the route will have tended to progress from settlement to settlement. Size is also an important factor as larger towns will have been the collecting point for larger areas, thus amassing more supplies.

Each settlement will have a surplus of food that, in addition to being affected by its surrounding landscape, is also affected by the time of year. The availability of grain is a commonly considered factor in the planning of military campaigns throughout antiquity and the amount of food surplus varies throughout the year by considerable amounts [13]. To facilitate this, settlements are not just static points with set amounts of resources but are their own special type of agent in order to take advantage of each tick of the simulation to update each type of resource as availability will change throughout the year.




Figure 3. Settlements and terrain data

As with roads, the TIB maps provide data for some areas. Based on the density of different types of settlement in the TIB-covered area, random points can be selected outside this area to create a reasonably plausible distribution of settlements. The TIB maps also separate settlements into various categories. These can be assigned population values that can be combined with research related to rural productivity to provide some indication of how much surplus food is available at each location.

#### 4.1.4. Land use

Transport of goods across Anatolia could be expensive due to the distances involved so the types of resources available to each settlement will have depended on local conditions. The produce available at each settlement depends on the agriculture of the area which in turn depends largely on the surrounding landscape. The main types of food available to the army are grain and meat. The availability of grain and meat have profoundly different implications for the movement and provisioning of the army. Grain was the staple food of antiquity, bread made from it would have formed a significant part of the diet at most levels of society. Bread however does not last forever and is bulkier to transport than the grain itself, making it standard for each unit to have handmills to grind their own grain and make their own bread.

Meat presents different challenges as it requires preserving in order to last any significant time but can be transported 'on the hoof' as live animals. This adds to the movement problems of the army by increasing its size but also provides it with food that effectively transports itself. The animals would be driven along with the army until needed and then would be butchered and prepared. Each type of food affects the movement of the army in different ways and the differential availability of each type of food across Anatolia will affect the army's progress in ways that cannot be accurately predicted in advance.

Any attempt to tailor the resources available at each settlement to the types of produce is bound to contain numerous assumptions. There is insufficient evidence about rural food production to enable us to put together a complete and convincing model of which types of agriculture were practiced in the many different environments of Anatolia. The best guide currently available is modern land use as this is partly based on climate and terrain, both of which are largely similar to 11<sup>th</sup> century conditions. Cultural and technical changes mean that differences exist but a comprehensive reconstruction of the flora, fauna and agricultural practices of Byzantine Anatolia is beyond the scope of this project.

The modern land-use data comes from the ESA's GLOBCOVER satellite programme [14]. This breaks land use down to a series of categories that can be converted into likely 11<sup>th</sup> century production. The amount of surplus

produce available to these communities is another topic where insufficient data exists to produce a comprehensive and plausible answer, but competing hypotheses exist that can be tested with the model [15].

As fine detail is unnecessary this data can be used at a resolution of approximately 500m<sup>2</sup> per cell.

#### 4.1.5. Water

Availability of water is the single most important issue for the army, without it the army will not survive for long in the hot Anatolian summer. Water sources are, like land use, an area where the difference between the modern and the ancient environment is largely unknown. Undoubtedly differences exist, especially where the modern Turkish state has turned to damming rivers to supply its irrigation and hydroelectric power needs. Despite this, with no way of determining how the situation has changed there is little alternative but to use modern data. This data has the benefit of containing flow rates, something unknowable for the 11<sup>th</sup> century even if a map of Byzantine Anatolian rivers could be found. None of the rivers of Byzantine Anatolia were navigable in the area covered by the ABM so the only effect of water on movement is to reduce speeds when crossing larger water courses.

#### 4.2. The Agents

The granularity and complexity of the agents is dictated by the project's aims. The practical problems involved with moving large numbers of people across broken terrain and through narrow ravines were known to medieval military leaders and terrain could greatly affect an army's progress [16]. Average movement rates have been calculated for armies based on historical itineraries but the relationship between army size and composition and the speed of march is one aspect that the model is ideally suited to address. With this in mind we have decided to model the people and pack animals of the army at a ratio of one agent per human or animal. This will result in a more convincing crowd movement model and simplify the modelling task as there will be no problems with deciding how to aggregate several individuals into one agent.

##### 4.2.1. Architecture

An agent consists of a plan queue, a messaging queue and a representation of a perception base along with a number of private and public variables modelling agent characteristics (Figure 4). The plan queue contains a list of the tasks the agent has to perform. The message queue contains a list of messages from other agents, including orders from superiors and messages from comrades. The perception base contains the information that the individual agent has regarding the world, including information gathered by the agent's own senses and information introduced by communication with other agents.

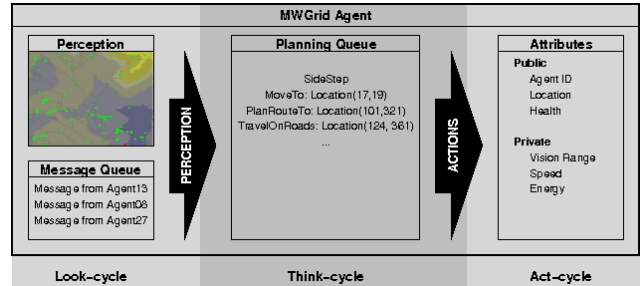


Figure 4. The MWGrid agent

Each agent has a series of variables, depending on its type. These are dictated by the need to model the organisation of the army and its movement and the effects of the march on each individual agent. An agent's variables can be public (those that are apparent to other agents such as health and location) and private (speed, vision range). It can perform range queries to access the public variables of every agent within a certain distance. Orders are given when another agent passes a message to an agent with an order as the content. Providing the message comes from a valid source (a superior or a trusted comrade) the content of the message is added to the agent's plan queue.

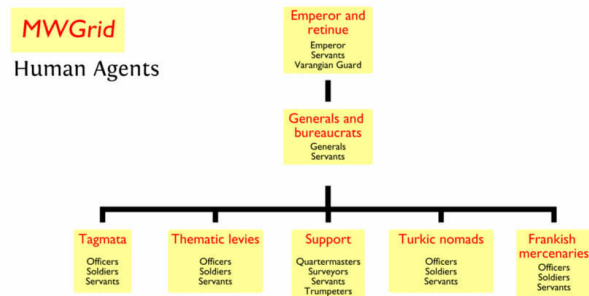
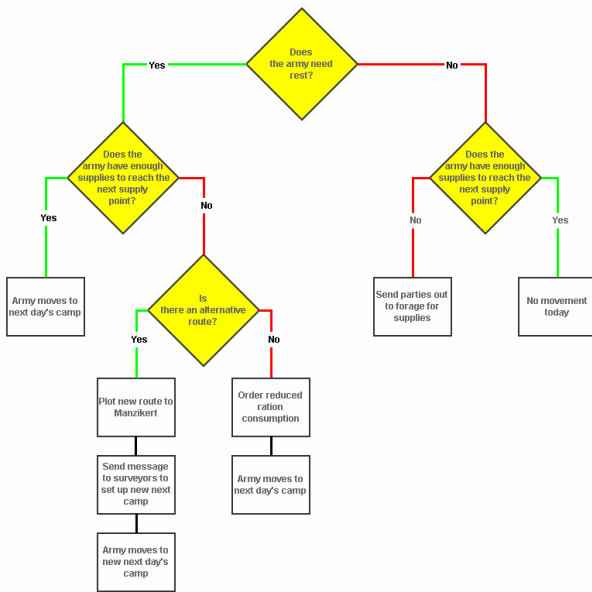


Figure 5. Agent organisation

The organisation of the Byzantine army provides a rigid hierarchy into which the agents can fit (Figure 5). The Emperor occupies the apex of the pyramid and makes the decision of which route to take (Figure 6). His order to move, along with the location of the next day's camp, will be propagated down the command structure from officer to officer until it reaches the officers of each individual squad.



**Figure 6. The Emperor's decision making process**

Agents lower in the hierarchy contain less decision making processes. The main task of each agent is to transport itself and any equipment it is responsible for from the previous camp to the next one. Once at the camp some soldiers, primarily from the lower status Thematic troops will be tasked with digging a bank and ditch around the camp while others ensure their unit has sufficient supplies. These supplies will be collected from their unit's baggage train. The camp will have been deliberately sited near a water source where possible to ensure a decent supply each day and to avoid the need to transport too much water between camps. The command structure of the army again provides a useful tool to regulate behaviour with resources flowing up and down between baggage handlers at both unit and squad level and information on what is running low flowing up to those with responsibility to organise resupply.


#### 4.2.2. Movement

The movement of the army and its individual agents is a critically important part of the model. The model specifically aims to help assess how the multitude of micro-level movements of individuals affects the macro-level movement of the army. How these interact and what repercussions they have for the environment the army passes through is the central thrust of the project. This distinction between the macro-level movement of the army and the micro-level movement of the individuals within it is reflected in the systems designed to facilitate this movement.

The emperor will have decided on a route in advance of the campaign. This was normal procedure and allowed the areas being passed through enough warning to stockpile the

requested supplies to feed and equip the army. The planning of this overall route will have depended on ease of movement and the availability of supplies and settlements in which to store them. This aspect of planning is dealt with by a modification of standard Probabilistic Roadmap (PRM) movement. In standard PRM movement a random pattern of nodes is created over the environment and edges are created between them [17]. Agents can plan routes and move between nodes along edges. This makes the process of navigating across large numbers of individual cells considerably easier as route planning decisions are taken from node to node rather than from cell to cell. In the MWGrid ABM, the random pattern of nodes is replaced by the pattern of settlements. This is more appropriate to the model than a random series of nodes as the army needs to move from settlement to settlement to pick up resources and take advantage of water sources.

Once an overall route is planned by the Emperor, the settlements along the route are treated as waypoints between which an A\* route planning algorithm can be used [18]. Due to the fine resolution of the environment, the distance between settlements often results in an excessive tree depth which increases the processing time of the A\* planner. An intermediary stage is therefore introduced where, instead of planning the route between settlements from environment cell to environment cell, the route is planned from terrain height data point to terrain height data point (Figure 7). The terrain data is at roughly 50m<sup>2</sup> resolution, 10 times coarser than the environment. This reduces the tree depth by a factor of 10 as a result, ensuring the route plan completes in an adequate timescale. Each agent can then use a further instance of the A\* route planner to move from terrain point to terrain point, a distance of 10-15 environment cells. At this stage the route planner can take into account cells that are already full of agents, avoiding them if possible.



**Figure 7. The three levels of route planning**

The Emperor plans a route between settlements and breaks the route down into marches of a day's length. The location of the next day's camp is communicated to the Emperor's immediate subordinates who pass this message down until it reaches individual squad leaders. This location is used as the destination of each individual squad leader's intermediate route plan. The soldiers follow their squad



leader using normal flocking behaviours unless told to perform some other task [19].

### 4.2.3. A\* route planning

Even chopping the route into discrete sections, an unmodified A\* route planner will not do everything needed in a reasonable timescale. One way in which performance can be improved is to have a dynamic method of calculating the A\* route planner's heuristic. The heuristic calculates how many steps it takes to reach our destination and assigns a cost of 1 per cell. If actual movement costs are 1 or higher per cell this means the result will be the lowest cost route however as the difference between the heuristic and the actual movement cost increases the processing time also lengthens as the planner has to examine more of the possible routes. It can be assumed that the average cost of movement will be greater than 1 per cell and increase the heuristic cost accordingly. This will speed up the route planning but if the heuristic cost exceeds the actual cost then the route may not be the one with least cost.

There is clearly a trade off between performance and quality. Due to the variety of terrain types involved in the journey across Anatolia, ranging from large flat plains to steep mountain passes, static costs for both movement and the heuristic would result in cases of either suboptimal routes or long processing times.

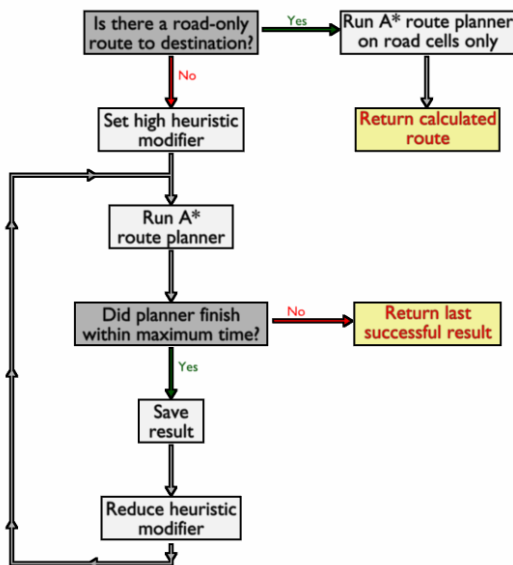


Figure 8. The dynamic A\* route planning process.

One solution is to have a dynamic method of calculating the heuristic cost. This process starts the route planning with a high heuristic cost per cell, likely to create a fast but suboptimal route. The route is then recalculated with decreasing values for the heuristic modifier until it exceeds a set time limit for the route planning process. It

then takes the route calculated by the last successful run and uses that as the route. This method ensures all routes fall within an acceptable area of accuracy and performance.

A separate step can be added which attempts to find a route along a road from the start location to the destination. If such a route exists then this route will be taken, working on the assumption that a road route is preferable to an off road route even if the road route is longer (Figure 8).

### 4.2.4. Planning

Each agent has a plan queue where an agent's designated tasks are stored in the order in which they need to be performed. The current plan to be executed consists of a series of actions. Plans further down in the queue consist of one symbolic action that is expanded into a series of appropriate actions when it is executed (Figure 9).

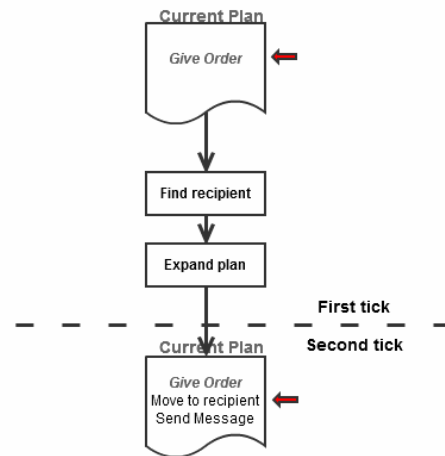


Figure 9. The expansion of a plan into a series of actions

If these subsequent actions need revising then the action queue can be cleared back to the original first action and then expanded again. This is useful if an agent has a plan to pick up a resource, creates a plan to move to the nearest example of the resource then finds that someone else has taken the resource in the meantime. In this case all subsequent actions can be cleared and the original action can be reprocessed, leading to the next nearest resource being found.

This planning process works well with the limited number of actions required. Due to its highly logistical nature, the majority of the tasks involve moving somewhere and interacting with some environment object. With the addition of a 'priority' attribute and a 'plan to execute in event of failure' in each plan, each agent should have enough information to be able to prioritise its own tasks and fail in an intelligent way.

The scenarios we wish to model as part of the Byzantine army's march across Anatolia can be modelled with relatively basic object handling plans (Collect

Resource, Pickup Object, Drop Object), basic message passing instructions (Give Order, Request Information, Pass Message) and some plans referring to specific situations needed during the army's march (Set Fire, Setup Tent, Dig Ditch, Patrol). The limited number of plans required and the restricted set of circumstances in which they will be used means this approach in which the process of performing tasks is largely hard-coded doesn't increase the time involved with programming the model unreasonably.

#### 4.2.5. Messaging

Every agent in the army has an appreciation of where they are in the army's organisational structure. They know who their superiors are and who is under their command. They also know the rank of all other agents. In this way an agent can prioritise messages received from other agents. Each message consists of the Agent ID of the sender, the Agent ID of the intended recipient and the content itself. Content can be a piece of information about the environment, some information about another agent or an order. It is placed in the recipient's message queue to await processing. On being processed, either the information is added to the agent's perception base or the order (if from a valid source) is added to the agent's plan queue.

Orders are propagated down the chain of command until they reach the intended recipient. Information is also contained in messages and can be filtered up the chain of command to the emperor. If the emperor is to know whether the army needs to rest or forage for supplies he must gather information from the units and the quartermasters about the status of the provisions. The information exchange within an army is relatively tightly constrained, a soldier who discovers his squad is without food reserves will not take this information straight to the emperor, he will inform his superior, who will in turn do the same and the information will reach the decision makers in that way.

#### 4.3. Middleware

Interaction between the ABM and PDES-MAS is handled by the Middleware layer of the framework. Time in a MWGrid simulation is partitioned into discrete 'time-steps', these are maintained by the Scheduler. During each time-step the Scheduler calls the step-method of each agent in the simulation. This step-method then goes through a look-think-act cycle of the agent (Figure 4). During the look-cycle the agent senses its environment and reads its messages from the message queue. Sensing takes the form of reading public variables for all agents through the use of range-queries. In the think-cycle the messages read from the message queue will be interpreted and the plan-queue is then updated. This then results in an action in the act-cycle which is then expressed as writing public and private variables of the agent itself, or sending messages to other agents. The simulation ends when the allotted time-steps

have run out and the step-method of all agents have been called for each of these time-steps.

The simulation itself implements an update-method. The functionality of the update-method is to turn internal data in the simulation into external data in the form of a trace-file. The trace-file is then used for post-processing to produce statistics, and 2D and 3D representations or animations of the simulation (see also Figure 2). To this end, the Scheduler calls the update-method each time-step thus providing read access to all agents in the simulation.

## 5. SUMMARY AND FUTURE WORK

This paper has described a large scale agent-based model, designed to study medieval military logistics. An important historical event, the Manzikert campaign, is used as a case study. Details in the historical record provide important, well established details that can be used to validate the model's outputs. The model's outputs in turn can be used to define parameters within which the historical evidence can be framed.

Historical sources for the Manzikert campaign give inflated numbers for the size of the Byzantine army yet, despite an agreement on the unreliable nature of the sources, contemporary historians are unable to provide evidence for even the range of possibilities. It is the fundamental uncertainty over how many people could be transported across Anatolia and under what circumstances that the model will address. The model will examine the relationship between army size and speed of movement where previous work has relied on average movement rates [13]. It can help us determine the amounts of agricultural surplus required for armies of differing sizes. Different models of agricultural production can be tested against different models of army size and movement to help eliminate implausible scenarios. Some models of agricultural production may result in insufficient surplus to support even a modest army. Some army sizes may be unsupported with even the most optimistic models. In modelling these scenarios with different sizes of army and levels of food availability some models derived from historical sources may be determined practically impossible.

By visualising the movement of the army and the effect it has on the stores of the settlements visited it can provide easy to understand examples of a process previously only modelled with top down, systemic approaches. This will enable communication of the model's results to an audience unused to interpreting agent-based models.

The model itself can be expanded beyond the current scope to include the spread of disease using epidemiological models. Disease, more than famine, was an ever present threat on campaign and medieval armies could be seriously weakened by it. A agent-based approach to modelling the introduction and spread of disease through an army would



help provide parameters within which historians could evaluate existing theories.

## 6. ACKNOWLEDGEMENTS

This project is funded by a joint AHRC/EPSRC/JISC e-Science research grant.

## 7. REFERENCES

- [1] Gaffney, V. 2008. "In the Kingdom of the Blind: Visualisation and E-Science in Archaeology, the Arts and Humanities". In M. Greeenrass and L. Hughes (Eds) *The Virtual Representation of the Past*. 125-135 Ashgate, London
- [2] Lake, M. W. 2000. "MAGICAL computer simulation of Mesolithic foraging". In Kohler, T. A. and Gumerman, G. J., (Eds), *Dynamics in Human and Primate Societies: Agent-Based Modelling of Social and Spatial Processes*, 107-143. Oxford University Press, New York.
- [3] Kohler, T.A., C. D. Johnson, M. Varien, S. Ortman, R. Reynolds, Z. Kobti, J. Cowan, K. Kolm, S. Smith, and L. Yap. 2007. "Settlement Ecodynamics in the Prehispanic Central Mesa Verde Region". In Kohler, T. & van der Leeuw, S. E. (Eds) *The Model-Based Archaeology of Socionatural Systems*. School for Advanced Research Press, Santa Fe.
- [4] Lake, M. 2007 "Agent-based modeling: a question of scale. AHRC ICT Methods Network Workshop on Space and Time: Methods of Geospatial Computing for Mapping the Past" [http://www.nesc.ac.uk/talks/772/Lake\\_e\\_sci\\_07\\_v1.pdf](http://www.nesc.ac.uk/talks/772/Lake_e_sci_07_v1.pdf)
- [5] Haldon, J.F. (Ed) 2005 *General issues in the study of medieval logistics: sources, problems, methodologies*. Brill, Leiden.
- [6] Craenen, B., G. Theodoropoulos, V. Suryanarayanan, V. Gaffney, P. Murgatroyd and J. Haldon. "Medieval Military Logistics: A Case for Distributed Agent-based Simulation". Distributed Simulation & Online gaming (DISIO 2010); ICST/CREATE-NET; March 2010.
- [7] Craenen, B. and G. Theodoropoulos. "Interfacing Multi-Agent Models to Distributed Simulation Platforms: The Case of PDES-MAS". Winter Simulation Conference (WSC 2010); December 2010.
- [8] Friendly, A. 1981 *The Dreadful Day*. Hutchinson & Co. Ltd. London.
- [9] Logan, B. and G. Theodoropoulos. 2001. "The distributed simulation of multi-agent systems". *Proceedings of the IEEE* 89(2), 174–186.
- [10] NASA Jet Propulsion Laboratory, California Institute of Technology. <http://www2.jpl.nasa.gov/srtm/>
- [11] French, D. *Roman roads and milestones of Asia Minor*. BAR international series 105. Oxford
- [12] Hellenkemper, H. & Hild, F. 2004. *Lykien und Pamphylien*. Verlag der Österreichischen Akademie der Wissenschaften, Wien
- [13] Engels, D.W. 1978. *Alexander the Great and the Logistics of the Macedonian Army*. University of California Press.
- [14] Bicheron P., P. Defourny, C. Brockmann, L. Schouten, C. Vancutsem, M. Huc, S. Bontemps, M. Leroy, F. Achard, M. Herold, F. Ranera and O. Arino. 2008. GLOBCOVER Products Description and Validation Report. [http://ionia1.esrin.esa.int/docs/GLOBCOVER\\_Products\\_Description\\_Validation\\_Report\\_I2.1.pdf](http://ionia1.esrin.esa.int/docs/GLOBCOVER_Products_Description_Validation_Report_I2.1.pdf)
- [15] Goodchild, H. "Modelling Agricultural Production. A Methodology for Predicting Land Use and Populations". In J.F. Haldon, (Ed) *General issues in the study of medieval logistics: sources, problems, methodologies*. Brill, Leiden
- [16] Dennis, G.T. 1985. *Three Byzantine military treatises*. Dumbarton Oaks texts ; 9. Washington.
- [17] Kavradi, L. E., P. Svestka, J.C. Latombe and M.H. Overmars 1996. "Probabilistic roadmaps for path planning in high-dimensional configuration spaces", *IEEE Transactions on Robotics and Automation* 12 (4): 566–580
- [18] Hart, P. E., N.J. Nilsson and B. Raphael. 1968. "A Formal Basis for the Heuristic Determination of Minimum Cost Paths". *IEEE Transactions on Systems Science and Cybernetics SSC4* 4 (2): 100–107.
- [19] Reynolds, C. 1987. "Flocks, herds and schools: A distributed behavioral model.", *SIGGRAPH '87: Proceedings of the 14th annual conference on Computer graphics and interactive techniques* : 25–34