

Military Reconstructive Simulation in the Cloud to Aid Battlefield Excavations

Peter Wittek

Swedish School of Library and Information Science
University of Borås
Borås, Sweden
Email: peterwittek@acm.org

Xavier Rubio-Campillo

Computer Applications in Science & Engineering
Barcelona Supercomputing Centre
Barcelona, Spain
Email: xavier.rubio@bsc.es

Abstract—Archaeological studies on battlefields may see great benefits from simulated military engagements: simulations help testing hypotheses based on historical data and may also help with validating methodologies used on the site. Such methods, however, require high-performance computing expertise and considerable computational power. With the emergence of on-demand computing instances in the cloud, distributed computations have become available to technically every organization or individual. This puts large-scale battlefield simulations within the reach of archaeologists, and the cloud paradigm also lowers the required technological expertise, potentially leading to a more widespread adoption of such simulation methods.

Index Terms—Cloud Computing, Computer Assisted Archaeology, Battlefield Archaeology, Military Simulation

I. INTRODUCTION

Constructive military simulation refers to a simulation that involves simulated people operating in a simulated environment [2]. Such simulations are designed to find out the possible outcomes of the courses of actions taken by the real people, akin to testing hypotheses based on a set of assumptions. The simulations are often constructed by stochastic processes that calculate the results of interactions between the entities or units in the simulated environment.

Reconstructive simulation is similar in its principles, but the objective is to gain a more accurate picture of what has actually happened in a military engagement in the past. Researchers will usually try to answer several questions, ranging from the development of the engagement itself to the individual experience of soldiers in the battlefield and the impact of terrain in the final outcome. There are different data sources that can be used to accomplish this task, having particular importance primary textual descriptions of the engagement and archaeological remains. Computer simulation will be useful in order to integrate this varied information inside a common framework, and explore research hypotheses through the variation of starting conditions [41].

Looking at the algorithmic part of simulations, they are often based on discrete event simulation (DES) or agent-based modelling (ABM). In the former approach, the operation of a system is described by a chronological sequence of events, where each event occurs at an instant in time and marks a change of state in the system [39]. In ABM, the agents follow a simple programmed behaviour, and their actions

and interactions affect the system as a whole [32]. Irrespective of the approach used, the number of steps to calculate in a large simulated system is considerable. This is where parallel and distributed computing help, ensuring that the workload is efficiently divided among the networked units in a high-performance computing (HPC) cluster. Such resources, however, are unlikely to be available to archaeologists. An attractive alternative to HPC units is the cloud infrastructure. By cloud, we simply mean a dynamically provisioned infrastructure or service [1]. Under this model, a user can dynamically provision any amount of computing resources from a (cloud) provider on demand, and only pay for what is consumed. Technically, this means that the user is paying for access to virtual machine instances that run a standard operating system. The virtualization technology enables the cloud provider to allocate available physical resources and enforce isolation between multiple users who may be sharing the same hardware. Once one or more virtual machine instances have been allocated, the user has full control over the resources and can use them for arbitrary computation. When the virtual instances are no longer needed, they are destroyed, thereby freeing up physical resources that can be redirected to other users.

Using this model, the upfront investment in HPC resources can be avoided, and maintenance of the sophisticated hardware is also unnecessary. Given the ad-hoc nature of the calculations, the costs can be kept within reasonable limits. Virtualization also allows scientific codes to be optimized and pre-installed on machine images, facilitating control over the computational environment, and ensuring that non-technical personnel can also benefit from the cloud paradigm [37].

Cloud computing in the humanities already has some applications, for instance, in language technology [46], digital libraries [47], and digital preservation [48]. Different types of simulations have also been executed in cloud systems. Climate models were one of the first to benefit [15], [21], and large-scale spatio-temporal data analysis and visualization of a multi-dimensional climate dataset resulted in reduced processing time in a cloud [31]. The U.S. Department of Energy investigated the cloud option for long-term scientific simulation, with particular attention to security and privacy [49]. Healthcare simulations have direct implications on patients, and such cal-

culations might be entirely infeasible without cloud resources [45]. More generic parallel and distributed DES has also been benchmarked in the cloud to identify potential bottlenecks [9].

In this paper, we argue that cloud-based computations are the most viable way for archaeologists to leverage on novel simulation techniques that would aid their field work. The rest of this paper is organized as follows. Section II provides a brief introduction to field of reconstructive military simulation and how they can aid excavations. Section III argues that HPC resources are welcome in such simulations. We point out the major obstacles to a more widespread adoption of such techniques in archaeological research in Section IV. We argue that cloud-based simulation can address all of the major issues and we propose an efficient architecture (Section V). Finally, Section VI concludes the paper.

II. SIMULATING BATTLEFIELDS TO AID EXCAVATIONS

Recent years have seen a growing number of projects in Humanities and Social Sciences that use computer simulation as their main research tool. Agent-based models (ABM) have been particularly successful within this regard due to its approach. An ABM defines the behaviour of any entity of a system that involves decision-making processes known as agents. These agents interact within a controlled environment that can be a real landscape with geographical features like vegetation, transport systems, etc., or an abstract domain. The researcher creates an initial scenario, giving values to certain parameters according to certain hypotheses that need to be validated. The simulation thus configured is executed in a series of regular time steps: an hour, a week, a year, or anything in between, depending on which time interval the experiment is appropriate. For each of these time steps the agents evaluate the situation, plan what actions they want to execute and interact with each other and also with the environment. Finally, the results of several of these simulations are analysed in order to answer the questions set at the beginning of the research.

A paradigmatic example of this methodology is SugarScape [27], specially designed for social scientists. The authors build a virtual society from ground up using simple, individual behavioural rules. The results portray the emergence of several universal rules that are not explicit in the initial model, but which are apparent in the agent-agent and agent-environment interactions (such as migration, conflicts, trade, hierarchies, cultural identities, and so on) [19], [18].

Archaeological science is one of these social disciplines where simulation has been more extensively used [14], [12], [30], [11], [29], [6]. The reason is that archaeology attempts to understand human behaviour based primarily on materials found in excavations. The archaeological methodology tries to detect spatio-temporal patterns related to the location and type of found structures and objects, coping with a high degree of uncertainty [8]. Within this context, computer simulation can be used as a virtual lab, suited to explore different hypotheses capable of explaining detected patterns, as well as to validate them, at least within the context of the simulation model.

Focusing our discussion on conflict or battlefield archaeology, it is important to note that a battlefield is different to other types of archaeological sites, such as settlements, since it results from the concentration of thousands of human beings in a small, delimited zone for a brief period of time, usually for one or two days. Consequently, no structures are usually available, and the type of observed patterns is extremely different to a settlement, where a smaller number of individuals interacted during a long period of time. Furthermore, the discipline has relatively recent foundations, as the first excavations that led to the development of this particular archaeology were made during the 1980s' in the battlefield of Little Big Horn [42].

ABM provides the battlefield archaeologist with a new technique that can replicate both the type of data generated by a military engagement, and the different ways the data can be collected. It would be complicated to plan such research using real fieldwork, as each battlefield is unique. Indeed, the usual reason for excavating a battlefield is to understand the events that took place there, thus making it difficult to compare different methods used in different battlefields. Moreover, we cannot work twice on the same area in order to test different strategies, as the results would be directly related to the order in which they were obtained.

Computer simulation can also be used to explore topics closely related to the spatial patterns that can be detected in battlefield archaeology, such as troops deployments, formation and battle tactics [40]. It is worth noting that while the use of ABMs is considered in current military training [25], it is almost unknown in military history research. Nevertheless, some interesting experiments have been published in the last decade, like the study of the WWII submarine campaign in the Bay of Biscay [35], [23], [22], [3], guerilla warfare [13] or the logistic system of a medieval army [7].

For these reasons we claim that computer simulation can be used to reconstruct and explore past military engagements using data gathered from textual sources and battlefield archaeology. However, a researcher who attempts to use such a system should be cautious, because certain pitfalls and risks must be addressed in order to guarantee valid results and scientific quality. The most important point is that we should not accept the results of a simulation as directly applicable to reality. This problem is the result of misunderstanding the concept of modelling itself [17]. An excessive emphasis on individualism should also be avoided, a problem seen in badly designed agent-based models [34].

In a more practical sense, it is important that the models constructed are understandable and useful for humanities researchers and social scientists. If this is not the case, these models will not be published in journals of the discipline and their impact on the community of historians will remain insignificant however brilliant the design.

III. THE BENEFITS OF HPC IN SOCIAL SIMULATION

As interesting as they are, ABMs have a number of issues that need to be addressed. Some of these issues are responsible for the low impact of these models in humanities and social

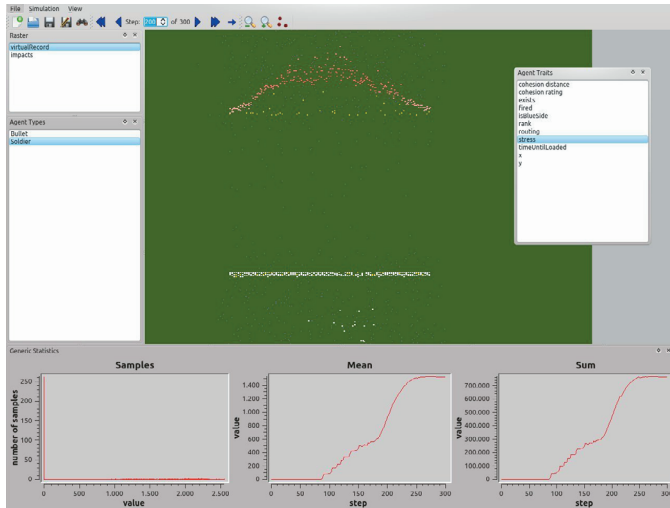


Fig. 1. Cassandra analysis tool

sciences, and if we manage to solve them, it could be an important step towards their more accepted use.

First of all, most of the existing ABM simulations tend to be simple and small. Despite the fact that lots of emergent behaviours can be detected in these environments, certain phenomena can only be observed in large-scale scenarios. For example, the tactics used in a given period are strongly related to the size of the armies that were deployed in a battlefield. Linear warfare of the 18th century cannot be correctly modelled while trying to simulate a few hundred individuals, as the different firing systems were thought to be used with thousands of soldiers, given the lack of accuracy of the firearms they used. Some of the hypotheses regarding command and control, and even individual behaviour follow the same reasoning, as do other types of traits. For this reason in some cases we will need to create larger simulations, more costly in terms of computer power and more difficult to analyse. The same problem is seen with the number of simulations. We have to work with a high degree of uncertainty, and for this reason, several simulations need to be executed in order to obtain valid results.

Both topics show that sometimes we will need to analyse larger datasets, and, as a consequence, the results can be more complex to understand and justify. Parallel to these issues there is another one, that, in our opinion, seems to be the most important one: the definition of behaviours. Classical ABMs often have a predefined behaviour, hard-coded inside the application. The major consequence is that critiques against ABM focus on the fact that they can be seen as self-fulfilling prophecies: emergent traits are not related to the problem, but to the way it was programmed, and if we can take a look at the code, we will learn the implicit or explicit assumptions the modeller had.

For simple models, it can be argued that the assumptions are less important, because such models can be replicated and understood without problems. Unfortunately the same

cannot be said about more complex models. If we want to analyse real-life data, or create agents with several traits, in the end it will become complicated to understand which system properties emerged from our model, and which ones from the way we programmed the agents. The solution to this issue is the development of agents with more elaborate behaviour, founded on artificial intelligence techniques instead of simple rules created by the researcher.

Summarizing these problems (size and number of simulations, and behaviour of agents) we can conclude that the future of ABM is strongly related to the use of high performance computational resources, capable of improving the scientific quality of the technique. Pandora¹ is a novel open-source framework designed to accomplish this task. It provides a C++ environment that automatically splits the execution of an ABM in different computer nodes. The package has support for distributed execution and serialization through HDF5 [16], several analysis techniques (spatial analysis, statistics and geostatistics, etc.), as well as advanced Artificial Intelligence agents. In addition, pyPandora is a Python interface to the framework, useful to design prototypes with minimal programming background. The ability to develop ABMs using Python is an important feature for social scientists, because this programming language is also used in other common modelling tools, like Geographical Information Systems (i.e. QGIS[36] and ESRI ArcGIS²). PyPandora allows researchers to create ABMs using a language that they already know, and enables the design of a common framework where one can combine simulation, spatial analysis and geostatistics.

Finally, the package is complemented by Cassandra, a visualization tool created to detect spatiotemporal patterns generated by the simulation (screenshot in Figure 1). This application allows any user to load the data generated by a Pandora parallel execution into a single computer and analyse it using a similar approach to sequential ABM platforms like NetLogo.

The capabilities provided by Pandora and other distributed ABMs packages like Repast-HPC [5] and GridABM [20] are solving the need for tools capable of exploiting HPC resources for Social Sciences and Humanities, but there still exists other technical issues that need to be considered, regarding the computational infrastructure needed to support these projects.

IV. BARRIERS TO ADOPTION

A widespread recognition of the importance of simulated battlefields in archaeological field research does not appear to be imminent. This has multiple reasons which all stem from the high barrier to entry considering the background and resources of a practising archaeologist.

First and foremost, large-scale agent-based simulations of battlefield engagements require considerable computational power. Parallel and distributed ABM asks for an HPC cluster to be able to efficiently test a number of hypotheses. Such

¹<https://github.com/miquelramirez/simulpast-cs1>

²<http://www.esri.com/software/arcgis/index.html>




Fig. 2. Overview of cloud simulation architecture

resources are unlikely to be available to an archaeological department.

It is also important to recognize the ad-hoc nature of these computations. Once the hypotheses are tested, the data are generated, visualization can be performed on an ordinary computer. Therefore it is hard to justify an investment in HPC resources, as they would idle most of the time. The associated maintenance cost is not negligible either.

Going beyond the hardware, the building of the software stack from scratch is a non-trivial task for a humanities researcher. The lack of technological expertise in managing the cluster, compiling code, and assembling the simulation workflow is probably the single biggest obstacle to adoption. Either a lengthy training or hiring a Computer Science professional would help.

Simulating a battlefield has an excess number of parameters. Misunderstanding the basics of the model might lead to incorrect parameter combinations, which may lead to misleading conclusions. An intuitive, non-technical interpretation of the parameters and results is essential.

V. ARCHAEOLOGICAL SIMULATIONS IN THE CLOUD

Cloud computing refers to a set of loosely related concepts that are typically manifested as dynamically provisioned infrastructure or service [1]. In what follows, we primarily view cloud computing as infrastructure-as-a-service, offering relatively bare bones systems on top of which a user or organization needs to deploy and manage their applications and data [38]. A public cloud makes a distributed infrastructure available in a pay-as-you-go manner to the public, the service being sold as utility computing. The idea behind utility computing is to treat computing resources as a metered service, like electricity or natural gas. Under this model, a user can dynamically provision any amount of computing resources from a cloud provider on demand and only pay for what is consumed. Resource consumption is measured

in machine-hours, breaking down to CPU hours, bandwidth usage, etc. The dynamically provisioned computer instance is also referred to as a virtual instance. The virtual instances behave like commodity-grade computers, but they may not directly correspond to one single computer. For instance, if the user requests two virtual instances with one CPU core in each, the virtual instances might be provisioned on a single computer with a quad-core CPU, where the remaining two cores are used for virtual instances requested by other users. It can also happen that the two virtual instances will run on two, physically distant computers.

A Beowulf cluster is a computer cluster of what are normally identical, commodity-grade computers networked into a small local area network with libraries and programs installed which allow processing to be shared among them. A number of identical virtual instances can make up a Beowulf cluster³⁴. Several cloud providers offer specific HPC cluster instances. It is ensured that such cluster instances that are launched simultaneously are physically close and they are connected with a high-speed network, and, as a result, they are more expensive than comparable standard instances.

Amazon Web Services (AWS) is one of the major cloud providers, and several research teams have successfully deployed scientific workflows on their infrastructure. AWS Elastic Compute Cloud (EC2) instances can be a possible solution to on-demand, small sized HPC applications [15], [24], [28], [26], [44].

A cloud-based Beowulf cluster clearly solves the computational problem for archaeological simulations. Since it is provisioned dynamically and paid for by the hour, there is no need for upfront investment and there is no maintenance cost either. Extensive simulations can be performed, testing a

³<http://www.datawrangling.com/on-demand-mpi-cluster-with-python-and-ec2-part-1-of-3>

⁴<http://aws.amazon.com/hpc-applications/>

large number of hypotheses, and, having obtained the results, local desktops can store the data for visualization. Calculations of such scale would simply not be possible for archaeological departments.

Figure 2 illustrates the general layout of our proposed architecture⁵. The end-user, the archaeologists is provided a virtual image of a Linux distribution, containing all supporting libraries and an initial configuration. A simulation scenario runs on the top of this software stack. The virtual image is launched on a number of cloud computing instances, forming an HPC cluster. Since storage in the computing instances is volatile, the simulation results are streamed to a persistent cloud storage. From here, the results can be shared with fellow archaeologists, and they can also be downloaded for local processing, visualization, and interpretation of results.

The virtual image can be downloaded and manipulated locally. This is important when developing new simulations, as a locally run virtualised computer instance would not occur additional costs. Once the development is completed, the image can be re-imported to the cloud. Exploration of the parameter space and large-scale experiment can be subsequently executed on cloud clusters.

Note, however, that simulations may not be the most ideally suited application for cloud computing. Communication requirements are significant in DES [44] and also in agent-based modelling [4]. The interconnect on the EC2 platform severely limits performance and causes significant variability when communication dominates the running time of an algorithm [26]. As pointed out above, specific HPC instances can be reserved with guaranteed high-speed network connection, renting these instances significantly increases the costs. The cost-efficient solution to the problem is to launch a lower number of instances and run the computations longer. If the right amount of computational power is provisioned, the cost can be significantly reduced without impact on application performance [10].

Beyond computational aspects, the cloud paradigm also helps with the lack of technological expertise. The virtual image in the cloud, configured to run Pandora, saves the trouble of building the software stack in a software environment that is likely to be unfamiliar to an archaeologist. Identical launches of this virtual image can constitute a cluster. Sharing this image with the broader community shifts the focus from command line expertise to the menu-driven interface of a cloud provider. While this could still be improved to entirely mask the technical details from the user, it is an important step forward.

VI. CONCLUSION

The key benefits of cloud-based archaeological simulations can be summarized as follows:

- Lowering the technological barrier: sharing the virtual image of a configured system enables non-technical people to utilize ABM simulation methodology in their work.

⁵Figure based on Amazon Web Services reference architectures: <http://aws.amazon.com/architecture/>

- Feasibility: Given the computational requirements, traditional or cloud-based HPC resources are a must. Since traditional HPC clusters are unlikely to be available to humanities departments, cloud computing is the only option.

Future work includes further lowering the need for technical expertise. While the benefits of an infrastructure-as-a-service model are clear, archaeologists would probably still need training to use such systems. We envision delivering battlefield simulation in a software-as-a-service model [43] that would completely obliterate the need for technical training and would also help the archaeologists in understanding the intricate network of parameters that are inherent in simulation. Ultimately, we believe that the cloud paradigm could enable simulation systems to aid archaeologists on the field using handheld devices [33].

VII. ACKNOWLEDGEMENT

This work was carried out while the first author was visiting the Department of Computer Applications in Science & Engineering at the Barcelona Supercomputing Center, funded by the "Access to BSC Facilities" project of the HPC-Europe2 programme (contract no. 228398). The second author is supported by the SimulPast Project (CSD2010-00034), funded by the CONSOLIDER-INGENIO2010 program of the Ministry of Science and Innovation – Spain.

REFERENCES

- [1] M. Armbrust, A. Fox, R. Griffith, A. Joseph, R. Katz, A. Konwinski, G. Lee, D. Patterson, A. Rabkin, and I. Stoica, "Above the clouds: A Berkeley view of cloud computing," EECS Department, University of California, Berkeley, Tech. Rep. UCB/EECS-2009-28, Tech. Rep., 2009.
- [2] E. Cayirci, "NATO's joint warfare centre perspective on CAX support tools and requirements," in *Proceedings of ITEC-06, the International Training and Education Conference*, 2006, May 2006.
- [3] L. Champagne and R. R. Hill, "A simulation validation method based on bootstrapping applied to an agent-based simulation of the Bay of Biscay historical scenario," *Journal of Defense Modeling and Simulation: Applications, Methodology, Technology*, vol. 6, no. 4, pp. 201–212, 2009.
- [4] L. S. Chin, C. Greenough, and D. J. Worth, "Optimising communication routines in parallel X-agents," Software Engineering Group, Rutherford Appleton Laboratory, Tech. Rep., January 2007.
- [5] N. Collier and M. North, "Repast HPC: A platform for large-scale agent-based modeling," in *Large-Scale Computing Techniques for Complex System Simulations*, W. Dubitzky, K. Kurowski, and B. Schott, Eds. Wiley, 2011, pp. 81–110.
- [6] A. Costopoulos and M. W. Lake, *Simulating Change: Archaeology Into the Twenty-First Century*. University of Utah Press, 2010.
- [7] B. Craenen, G. Theodoropoulos, V. Suryanarayanan, V. Gaffney, P. Murgatroyd, and J. Haldon, "Medieval military logistics: A case for distributed agent-based simulation," in *Proceedings of SIMUTools-10, 3rd International ICST Conference on Simulation Tools and Techniques*, Malaga, Spain, March 2010.
- [8] E. Crema, A. Beva, and M. Lake, "A probabilistic framework for assessing spatio-temporal point patterns in the archaeological record," *Journal of Archaeological Science*, vol. 37, no. 5, pp. 1118–1130, 2010.
- [9] G. D'Angelo, "Parallel and distributed simulation from many cores to the public cloud," in *Proceedings of HPCS-11, International Conference on High Performance Computing and Simulation*, Istanbul, Turkey, July 2011, pp. 14–23.
- [10] E. Deelman, G. Singh, M. Livny, B. Berriman, and J. Good, "The cost of doing science on the cloud: the Montage example," in *Proceedings of SC-08, 21st Conference on Supercomputing*, Austin, TX, USA, November 2008, pp. 1–9.

- [11] J. Diamond, "Life with the artificial Anasazi," *Nature*, vol. 419, no. 6907, pp. 567–569, 2002.
- [12] J. Doran, "Prospects for agent-based modelling in archaeology," *Archeologia e Calcolatori*, vol. 10, pp. 33–44, 1999.
- [13] —, "Modelling a typical guerrilla war," in *Proceedings of DIS-06, IEEE Workshop on Distributed Intelligent Systems: Collective Intelligence and Its Applications*, Prague, Czech Republic, June 2006, pp. 285–290.
- [14] J. Doran, M. Palmer, N. Gilbert, and P. Mellars, "The EOS project: Modelling Upper Palaeolithic social change," in *Simulating Societies: Computer Simulation of Social Phenomena*, N. Gilbert and J. Doran, Eds. UCL Press, 1994, pp. 195–221.
- [15] C. Evangelinos and C. Hill, "Cloud computing for parallel scientific HPC applications: Feasibility of running coupled atmosphere-ocean climate models on Amazon's EC2," in *Proceedings of CCA-08, 1st Workshop on Cloud Computing and Its Applications*, Chicago, IL, USA, October 2008.
- [16] M. Folk, A. Cheng, and K. Yates, "HDF5: A file format and I/O library for high performance computing applications," in *Proceedings of SC-1999, 12th Conference on Supercomputing*, Portland, OR, USA, November 1999.
- [17] V. Gaffney and G. Bradley, "Who's in command here? the digital basis of historical, military logistics," in *General Issues in the Study of Medieval Logistics: Sources, Problems and Methodologies*, J. Haldon, Ed. Brill Academic Publisher, 2006, pp. 37–69.
- [18] N. Gilbert, *Agent-Based Models*. SAGE Publications, 2008.
- [19] N. Gilbert and K. Troitzsch, *Simulation for the Social Scientist*. Open University Press, 2008.
- [20] L. Gulyás, A. Szabó, R. Legéndi, T. Máhr, R. Bócsi, and G. Kampis, "Tools for large scale (distributed) agent-based computational experiments," in *Proceedings of CSSSA-11, Computational Social Science Society of America Annual Conference*, Santa Fe, NM, USA, October 2011.
- [21] Q. He, S. Zhou, B. Kobler, D. Duffy, and T. McGlynn, "Case study for running HPC applications in public clouds," in *Proceedings of HPDC-10, 19th ACM International Symposium on High Performance Distributed Computing*, Chicago, IL, USA, June 2010, pp. 395–401.
- [22] R. R. Hill, R. Carl, and L. E. Champagne, "Using agent-based simulation to empirically examine search theory using a historical case study," *Journal of Simulation*, vol. 1, no. 1, pp. 29–38, 2006.
- [23] R. R. Hill, L. E. Champagne, and J. C. Price, "Using agent-based simulation and game theory to examine the WWII Bay of Biscay U-boat campaign," *The Journal of Defense Modeling and Simulation: Applications, Methodology, Technology*, vol. 1, no. 2, pp. 99–109, 2004.
- [24] Z. Hill and M. Humphrey, "A quantitative analysis of high performance computing with Amazon's EC2 infrastructure: The death of the local cluster?" in *Proceedings of GRID-09, 10th IEEE/ACM International Conference on Grid Computing*, Banff, Canada, October 2009, pp. 26–33.
- [25] A. Ilachinsky, *Artificial War: Multiagent-Based Simulation of Combat*. Center for Naval Analyses - World Scientific, 2004.
- [26] K. Jackson, L. Ramakrishnan, K. Muriki, S. Canon, S. Cholia, J. Shalf, H. Wasserman, and N. Wright, "Performance analysis of high performance computing applications on the Amazon Web Services cloud," in *Proceedings of CloudCom-10, 2nd IEEE International Conference on Cloud Computing Technology and Science*, Indianapolis, IN, USA, December 2010, pp. 159–168.
- [27] E. J.M. and A. R.L., *Growing Artificial Societies: Social Sciences from the Bottom Up*. MIT Press, 1996.
- [28] G. Juve, E. Deelman, K. Vahi, G. Mehta, B. Berriman, B. Berman, and P. Maechling, "Scientific workflow applications on Amazon EC2," in *Proceedings of eScience-09, 5th IEEE International Conference on E-Science Workshops*, Oxford, United Kingdom, December 2009, pp. 59–66.
- [29] T. Kohler and S. A. van der Leeuw, *The Model-based Archaeology of Socionatural Systems*. School for Advanced Research Press, 2007.
- [30] M. Lake, "MAGICAL computer simulation of Mesolithic foraging," *Dynamics in Human and Primate Societies: Agent-based Modelling of Social and Spatial Processes*, pp. 107–143, 2000.
- [31] S. Lu, R. Li, W. Tjhi, K. Lee, L. Wang, X. Li, and D. Ma, "A framework for cloud-based large-scale data analytics and visualization: Case study on multiscale climate data," in *Proceedings of HPCCloud-11, Workshop on Integration and Application of Cloud Computing to High Performance Computing*. Athens, Greece: IEEE, December 2011, pp. 612–617.
- [32] C. Macal and M. North, "Tutorial on agent-based modelling and simulation," *Journal of Simulation*, vol. 4, no. 3, pp. 151–162, 2010.
- [33] E. Mancini, G. Wainer, K. Al-Zoubi, O. Dalle *et al.*, "Simulation in the cloud using handheld devices," in *Proceedings of MSGC-12, 1st IEEE Workshop on Modeling and Simulation on Grid and Cloud Computing*, Ottawa, Canada, May 2012.
- [34] D. O'Sullivan and M. Haklay, "Agent-based models and individualism: is the world agent-based?" *Environment and Planning*, vol. 32, no. 8, pp. 1409–1425, 2000.
- [35] J. Price, "Game theory and U-boats in the Bay of Biscay," Master's thesis, Air Force Institute of Technology, 2003.
- [36] Quantum GIS Development Team, *Quantum GIS Geographic Information System*, Open Source Geospatial Foundation, 2009. [Online]. Available: <http://qgis.osgeo.org>
- [37] J. Rehr, F. Vila, J. Gardner, L. Svec, and M. Prange, "Scientific computing in the cloud," *Computing in Science & Engineering*, vol. 12, no. 3, pp. 34–43, 2010.
- [38] T. Rings, G. Caryer, J. Gallop, J. Grabowski, T. Kovacicova, S. Schulz, and I. Stokes-Rees, "Grid and cloud computing: opportunities for integration with the next generation network," *Journal of Grid Computing*, vol. 7, no. 3, pp. 375–393, 2009.
- [39] S. Robinson, *Simulation: the practice of model development and use*. Wiley, 2004.
- [40] X. Rubio-Campillo, J. M. Cela, and F. X. Hernández, "The development of new infantry tactics during early eighteenth century: a computer simulation approach to modern military history," *Submitted*.
- [41] X. Rubio-Campillo, J. M. Cela, and F. X. Hernández Cardona, "Simulating archaeologists? Using agent-based modelling to improve battlefield excavations," *Journal of Archaeological Science*, vol. 39, no. 2, pp. 347–356, 2012.
- [42] D. Scott, R. Fox, M. Connor, and D. Harmon, *Archaeological Perspectives on the Battle of Little Bighorn*. University of Oklahoma Press, 1989.
- [43] M. Turner, D. Budgen, and P. Brereton, "Turning software into a service," *Computer*, vol. 36, no. 10, pp. 38–44, 2003.
- [44] K. Vanmechelen, S. De Munck, and J. Broeckhove, "Conservative distributed discrete event simulation on Amazon EC2," in *Proceedings of MSGC-12, 1st IEEE Workshop on Modeling and Simulation on Grid and Cloud Computing*, Ottawa, Canada, May 2012.
- [45] H. Wang, Y. Ma, G. Pratz, and L. Xing, "Toward real-time Monte Carlo simulation using a commercial cloud computing infrastructure," *Physics in Medicine and Biology*, vol. 56, pp. N175–N181, 2011.
- [46] G. Wilcock, "Cloud computing for the humanities: Two approaches for language technology," in *Proceedings of HLT-10, 4th International Conference on Human Language Technologies – The Baltic Perspective*, Riga, Latvia, October 2010, pp. 241–248.
- [47] P. Wittek and S. Darányi, "Leveraging on high-performance computing and cloud technologies in digital libraries: A case study," in *Proceedings of HPCCloud-11, Workshop on Integration and Application of Cloud Computing to High Performance Computing*, Athens, Greece, December 2011, pp. 606–611.
- [48] —, "Digital preservation in grids and clouds: A middleware approach," *Journal of Grid Computing*, vol. 10, no. 1, pp. 133–149, June 2012.
- [49] K. Yelick, S. Coghlan, B. Draney, and R. Canon, "The Magellan report on cloud computing for science," U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Tech. Rep., December 2011.