The Knowledge Engineering Review, Vol. 31:4, 323–324. © Cambridge University Press, 2016
doi:10.1017/S0269888916000205

Spatial computing: introduction to the special issue

STEFAN DULMAN[1], JEAN-LOUIS GIAVITTO[2], ANTOINE SPICHER[3] and MIRKO VIROLI[4]

[1]*Hive Systems, Borneokade 117, 1019XC, Amsterdam, The Netherlands;*
e-mail: stefan.dulman@gmail.com;
[2]*IRCAM - CNRS UMR 9912 STMS, 1, place Igor-Stravinsky, 75004 Paris, France;*
e-mail: Jean-Louis.Giavitto@ircam.fr;
[3]*Université Paris-Est, 6-8 avenue Blaise-Pascal, Cité Descartes - Champs-sur-Marne, 77455 Marne-la-Vallée, France;*
e-mail: antoine.spicher@univ-paris12.fr;
[4]*Universita di Bologna, Via Zamboni, 33, 40126 Bologna, Italy;*
e-mail: mirko.viroli@unibo.it

Spatial Computing is an emerging field of research focusing on explicitly using the concept of space in computations. 'Spatial computers' are collections of local computational devices distributed through a physical space, in which direct interaction is bounded to some form of proximity, and the functional goals of the system are generally defined in terms of the system's spatial structure. Applications span many different domains, including parallel computer architecture, environmental monitoring, pervasive computing, sensor networks, mobile and *ad hoc* networks, music and smart buildings, to cite a few. A common theme among all these domains is the emphasis on the representation and the explicit handling of space. A key aspect of spatial computations concern how they manage knowledge, namely, how knowledge can be gathered from the spatial environment in which computation is embedded, how it can be conveniently stored in a distributed manner, how reasoning and situation recognition can be spatially carried on, and finally, how distributed agents can be fed with the knowledge produced by spatial computations.

Starting from the Spatial Computing Workshop, which reached its 7th edition in 2014 and was hosted through the years at the International Conference on Autonomous Agents and Multiagent Systems (AAMAS) and at the IEEE International Conference on Self-Adaptive and Self-Organizing Systems (SASO), we issued an open call soliciting submissions on topics related to spatial computing and knowledge engineering. After a thorough review process, among the received submissions we selected three high-quality papers: they altogether well review several facets of how spatial computing combines with knowledge engineering.

The first paper, by Jacob Beal, is entitled 'Trading accuracy for speed in approximate consensus', and addresses the problem of engineering the spatial computations needed to efficiently reconcile distributed knowledge into a single view, specifically, finding a consensus value. The main idea is to relax the requirements of accuracy in order to gain a significant increase in convergence time. Instead of waiting to converge to the mean value of the values in the network, the authors propose a routing overlay allowing some values to travel faster long-distances, so as to enhance the overall mixing process. Thus, starting from a well-known Laplacian-based consensus algorithm which converges in the order of the squared diameter of the network, the two new proposed algorithms succeed in achieving a convergence time in the order of the diameter of the network. The trade-off is the loss in precision—the convergence point will be around a value close to the average. As the author points out, this novel combination of traditional gossiping techniques with simple routing overlays opens an interesting new set of possibilities for a broader class of algorithms.

The second paper, by Simon Dobson and colleagues, is entitled 'Spatial Awareness in Pervasive Ecosystems', and addresses the issue of how knowledge can be managed in pervasive computing environments (which are intrinsically spatial) up to the point of achieving true self-organizing situation awareness. This is achieved by reviewing a number of recent results in the contexts of pervasive

https://doi.org/10.1017/S0269888916000205 Published online by Cambridge University Press

computing, service ecosystems, self-organization and knowledge representation, organized in a 'spectrum' of increasing awareness. This spectrum orderly features: (i) static and dynamic context representation (technologies used to model space, and data in space), (ii) local context perception (approaches for acquiring/organizing local context and for context-oriented mediation of interactions), (iii) context fusion (from local fusion of heterogenerous sensors to regional fusion of homogeneous sensors), (iv) distributed context perception (achieved by mechanisms of spreading, aggregation and evaporation, properly combined), and finally (v) self-organizing situation awareness (advanced self-organization techniques combining space and time to create processes supporting complex situation awareness).

The third paper, by Dominic Pacher, Robert Binna and Günther Specht, is entitled 'Optimizing Large Knowledge Networks in Spatial Computers', and deals with the problem of effectively and efficiently store large knowledge networks into spatial computers formed by regular one-dimensional, two-dimensional and three-dimensional spatial computers. Graphs are a primary structure for knowledge, and it is key to find ways of storing them in large spatial computers so that they properly distribute and adapt there, supporting efficient querying. The key factor that influences performance is the distance between linked nodes, and several techniques for minimizing it are presented, evaluated and compared. The key proposed technique is Mid Point Optimization, in which storing is evolved over time by moving nodes towards the mean position of its linked nodes.

We thank all the authors for their quality contributions, and the reviewers for their work in assessing and reviewing these papers. The increasing availability of computing devices spread in our living and working environments has the potential of dramatically impacting knowledge engineering, making it an even more spatial-aware process. In this regard, we hope this Special Issue can serve as an initial reference for many research studies evolving standard knowledge engineering techniques towards spatially distributed domains.

https://doi.org/10.1017/S0269888916000205 Published online by Cambridge University Press