The Computer Journal Advance Access published December 5, 2012

© The Author 2012. Published by Oxford University Press on behalf of The British Computer Society. All rights reserved.
For Permissions, please email: journals.permissions@oup.com
doi:10.1093/comjnl/bxs160

Editorial

JACOB BEAL[1,*], STEFAN DULMAN[2], OLIVIER MICHEL[3] AND ANTOINE SPICHER[3]

[1] *Raytheon BBN Technologies, Cambridge, MA 02138, USA*
[2] *TU Delft, Delft, The Netherlands*
[3] *Universite Paris-Est Creteil, Créteil, France*
**Corresponding author: jakebeal@bbn.com*

Welcome to this special issue of *The Computer Journal*, focused on the emerging field of spatial computing. In this issue, we present six papers at the leading edge of research in this area, each grappling with problems of computation on large networks of spatially distributed devices. Although these papers have diverse domains and focus topics, they all evidence an important emerging research trend: the exploitation of manifold geometry to provide scalability, adaptability and ease of programming in the construction of distributed systems.

1. TOWARD A UNIFIED APPROACH TO COMPUTATION AND SPACE-TIME

Throughout the 20th century, the dominant view of computation was one that was essentially Platonic and dualist. In this view, the model of computation is grounded in descriptions of actions that manipulate pieces of information in a purely virtual realm comprising such abstract spaces as 'heap' and 'stack', or 'pointers' and 'values.' Outside of this perfect Platonic world of ideals is the real world of hardware and networks, which modifies and intrudes on this perfect realm, often splitting a computation into pieces and introducing problems such as concurrency, distributed execution and communication error.

The dualist view has been very powerful, and is quite understandable in origin, given the history of computation as a child of mathematics. Here in our new century, however, computation is becoming ever more entangled with the realities of space and time. The environment in which we live is becoming saturated with computers—in our pockets, our appliances, our vehicles, our buildings—all of which can communicate with others nearby. Decreasing cost and energy requirements are enabling ever larger deployments of sensor networks and autonomous vehicles, for purposes from scientific surveys to infrastructure maintenance, from environmental monitoring to public security. At the same time, as Moore's law is drawing to a close, the need for increased computing power requires that computations be spread over ever larger collections of hardware, from application-specific integrated circuits and reconfigurable computing devices, to massively multicore processors and giant data centers.

In all these ways and many more, we are moving to a world where the old dualist model of computation must be eclipsed by a monist, unified view of the world, in which computation and space-time are tightly bound, such that *what* is being computed is inseparable from *where* the computation is taking place.

Spatial computing is the study of computation on such systems: specifically, any collection comprising myriad computing devices distributed through space and where the difficulty of moving information between devices is strongly governed by their spatial separation and the functional goals of the system are generally defined in terms of the system's spatial structure.

2. EXPLOITING MANIFOLD GEOMETRY ACROSS DIVERSE SYSTEMS AND DOMAINS

This special issue presents six papers from the forefront of research in this area. Between them, these papers address many different domains, including parallel computer architecture, environmental monitoring, sensor networks, mobile and *ad hoc* networks, and smart buildings. Likewise, the topics of the papers in this collection span the range from algorithms aimed toward particular applications to more general architectural models, and even a general theory of function evaluation over space-time.

Three of the papers in this collection develop algorithms aimed at particular spatial computing applications. 'Spatial Computers for Emergency Support,' by Avgoustinos Filippoupolitis, Gokce Gorbil and Erol Gelenbe, presents a smart building evacuation system, in which events and measurements propagating through a network of sensors and indicators collectively compute the best evacuation routes given the location of hazards such as fire and smoke. Wider ranging and possibly mobile sensor networks are the focus of

Downloaded from http://comjnl.oxfordjournals.org/ by guest on December 18, 2012

both 'Decentralized Detection of Qualitative Spatial Events in Evolving Areal Objects', by Muhammad Jafar Sadeq, Matt Duckham and Michael Worboys and 'Decentralized monitoring of moving objects in a transportation network augmented with checkpoints', by Alan Both, Matt Duckham, Patrick Laube, Tim Wark and Jeremy Yeoman. The first paper focuses on the tracking of topological changes in a target whose extent spans many sensors, such as a chemical plume splitting into two disconnected components. The second focuses on distributed tracking of mobile objects through a meshlike network of 'checkpoints'.

The other three papers are more general, with two of these presenting composable mechanisms that can be used as a basis for building new algorithms and applications. A movable architecture for robust spatial computing', by David H. Ackley, Daniel C. Cannon and Lance R. Williams, presents a new architecture for tilable computing hardware based on asynchronous cellular automata, in which the structure of a computation is constructed and reconstructed dynamically in order to be able to tolerate frequent hardware faults. Meanwhile, the paper 'GDE: A Spatial Gradient-Based Algorithm for Distance Estimation in Large-Scale Networks', by Qingzhi Liu, Andrei Pruteanu and Stefan Dulman presents an improved mechanism by which mobile wireless nodes can estimate distance relations through their network without the need for an externally supplied coordinate system, and demonstrates several ways this capability can be used as a building block for more complex programs.

The final paper is a domain-independent work on the theory of spatial computing: 'On the Evaluation of Space-Time Functions', by Jacob Beal, Kyle Usbeck and Brett Benyo addresses the general problem of evaluating function calls across space-time manifolds, producing a well-defined model and proving its correctness.

Despite their diversity, these papers all share a common approach at a deep level, evidencing an emerging trend in spatial computing research. The models and algorithms presented in these papers are all ultimately based, not in graph theory or Euclidean geometry, but in a manifold geometry where the measurements and topology of the mathematical model directly reflect the physical constraints on mobility of information and objects. This alignment of the computational and physical worlds brings an inherent adaptability to the models presented, while the continuous nature of manifolds aids in the scalability and ease of description for the algorithms. Moreover, this shared exploitation of manifold geometry increases the portability of results from domain to domain, since the fundamental model of distortion-tolerant geometry remains the same.

This shared exploitation of manifold geometry is likely linked to the assumptions that are also shared by all six of these papers: that algorithms should be scalable to very large networks, that systems must be able to cope with widespread ongoing dynamism and failures, and that space-spanning abstractions can help make distributed systems both practical and easy to engineer effectively. Manifold geometry, it appears, is a useful tool for addressing these challenges.

We are proud to present this collection of papers, and hope that the reader will find them illuminating, useful and inspiring for new work in this increasingly important field of research.

Downloaded from http://comjnl.oxfordjournals.org/ by guest on December 18, 2012