

Decentralized Spatial Computing

Foundations of Geosensor Networks

Bearbeitet von
Matt Duckham

1. Auflage 2012. Buch. xxii, 322 S. Hardcover

ISBN 978 3 642 30852 9

Format (B x L): 15,5 x 23,5 cm

Gewicht: 678 g

Weitere Fachgebiete > Geologie, Geographie, Klima, Umwelt > Geologie und
Nachbarwissenschaften > Geoinformatik

schnell und portofrei erhältlich bei


DIE FACHBUCHHANDLUNG

Die Online-Fachbuchhandlung beack-shop.de ist spezialisiert auf Fachbücher, insbesondere Recht, Steuern und Wirtschaft. Im Sortiment finden Sie alle Medien (Bücher, Zeitschriften, CDs, eBooks, etc.) aller Verlage. Ergänzt wird das Programm durch Services wie Neuerscheinungsdienst oder Zusammenstellungen von Büchern zu Sonderpreisen. Der Shop führt mehr als 8 Millionen Produkte.

Preface

This book was born out of the simple conviction: that there is a right way and a wrong way to design decentralized spatial algorithms. This conviction grew from many frustrating months designing decentralized spatial algorithms in, what I now believe to be, the *wrong* way. Conventional centralized approaches to algorithm design cannot be reliably applied to decentralized systems because they do not account adequately for the critical features of these systems: the interaction between nodes, and the limited, local knowledge of individual system components. The issues of decentralization have long been a topic of study in the domain of distributed computing. However, advances in technology, such as geosensor networks and smart phones, are bringing these issues to the fore of spatial computing too. Like so many other books in the field of geographic information science, the key question is once again: “What’s special about spatial?” More specifically, why should we study the problems of designing decentralized *spatial* algorithms separately from the more general problems of decentralized algorithms? The answers to this question begin in Chapter 1 (on the second page of Chapter 1), but continue throughout the book, to the final chapter.

How to read this book

This book has a simple story to tell, and so is best read in chapter order. The chapters form a logical progression, each chapter building on the previous: Chapter 1 setting the scene; Chapter 2 specifying an abstract formal model of decentralized spatial information systems; Chapter 3 introducing our decentralized spatial algorithm design technique; Chapter 4 exploring the design of increasingly sophisticated decentralized algorithms, with minimal spatial information; Chapter 5 introducing more sophisticated spatial capabilities; Chapter 6 adding time into the mix; Chapter 7 verifying empirically the efficiency; Chapter 8 verifying the veracity and robustness of those algorithms; and Chapter 9 finally outlining some of the important topics for further work in this area.

The book is written for two (intersecting) groups of students and researchers: computer scientists and geographic information scientists. GIScience is a big tent that delights in welcoming others from a plethora of disciplines, from geography to geostatistics, economics to ecology, philosophy to physics, and cognitive science to cartography as well as from traditional strongholds of computing—computer science, computational geometry, and mathematics. (These examples are not chosen at random: I have enjoyed close research collaborations with individuals from each and every one of these disciplines.) Some familiarity with basic discrete mathematics (sets, relations, functions, graphs) is certainly an advantage in reading this book. To help those who have forgotten, or even never encountered discrete mathematics before, the book contains a concise appendix to the key discrete mathematics topics required for this book, as well as recommended references to further readings in this area (Appendix A). There is also a brief introduction to SQL, the standard language for manipulating and querying databases, in Appendix B. For those new to GIScience, the core expertise in geographic information that unites GIScientists—what makes spatial special—can be found in a number of books on the topic, including [76] and a book I coauthored with Mike Worboys[119]. Those who wish to delve deeper into the more theoretical aspects of distributed systems and decentralized computing, recommended texts include [77, 82] and [99], the latter forming part of the foundations of this current book.

Throughout this book I have tried to write to my “student”—someone who wants to learn about the *process* of decentralized spatial algorithm design, what works and what doesn’t, what the basic principles are, and what the “take-home messages” are. This “big picture” can easily be obscured by the technical details. Where these two—technical rigor and procedural understanding—are in conflict, I have always sacrificed the focus, depth, and detail found in an original research article in favor of the broader message. A more comprehensive coverage of these technical details can be found in the many scholarly articles this book draws upon.

Understanding distributed and decentralized spatial algorithms requires one to be exposed not only to the theory, but also to the practice. To assist in the transition from theory to practice, from reading to internalizing, simulation models and associated source code for every algorithm presented in this book can be found (freely available) on the Web site that complements the book (<http://book.ambientspatial.net>).

The simulation models are written in NetLogo, a popular and easy-to-use agent-based simulation system that rewards handsomely the small investment in time required to learn how to use it. Despite the simplicity of NetLogo, developing a complete code corpus was a major effort. We have attempted to verify carefully all the code developed. However, given the scale and complexity of the task of developing the protocols and simulation models, it is entirely possible that some mistakes or problems may come to light. My research group and I will continue improving and extending the code base

at least until 2014, and hopefully even further into the future. Consequently, if you think you have found an error in the code, we would be keen to hear from you (contact details available through the Web site) and we will do our best to fix any errors.

At the end of each chapter you will also find a few review questions. These questions are designed to help you test your own understanding of the material in the chapter, as well as to provide some direction for further work or research. At the time of writing, I have not compiled a list of worked answers to these questions, but as and when these become available, they will also be placed on the book's Web site.

Finally, I welcome any comments or feedback in connection with this book and the Web site (for example, alternative or improved protocols or NetLogo code; examples of worked answers to review questions; and other comments on the book). I will endeavor to respond as quickly as my other commitments allow, and will post any errata on the book Web site.

Matt Duckham

Department of Infrastructure Engineering
University of Melbourne, Parkville, Australia

25 April 2012