FROM GPS AND VIRTUAL GLOBES TO SPATIAL COMPUTING – 2020

WHITE PAPERS

Table of Contents

SOCIETAL APPLICATIONS AND NATIONAL PRIORITIES

Spatial Computation and its Application to Disaster Management[1]

Dr. Nabil R. Adam
U.S. Department of Homeland Security, Science & Technology Directorate
Infrastructure Protection and Disaster Management Division

1. Introduction

The growing trend in the use of smart phones and other GPS-enabled devices has provided new opportunities for developing spatial computing applications and technologies in unanticipated and unprecedented ways. Spatial computing technologies which provide such capabilities as sensing, monitoring, and analysis, result in enhanced decision making.

For example, in a recent Intelligent Transportation project by IBM, researchers aim to help commuters avoid congestion and enable transportation agencies to better understand, predict and manage traffic flow [IBM11]. In this project traffic data is collected from various traffic flow sensors on roads, toll booths, intersections, and bridges. This information is combined with location based data from users' smart phones to learn their mobility pattern. Based on their preferred routes, the participating users would automatically receive traffic information and alerts on their phones; thus, resulting in reducing traffic congestion and accidents.

This project illustrates some of the capabilities of today's smart phones which highlight the potential of citizen sensors enabling the next generation of geo- informatics. An application area of such next generation of geo-informatics is *Social Media and its application to Disaster Management*.

2. Geoinformatics: application to disaster management

Social media, such as blogs, Twitter, and information portals, have emerged as the dominant communication mechanism of today's society. In the context of disaster management, exploiting such input to gain awareness of an incident is a critical direction for research in effective emergency management. Dynamic real-time incident information collected from on-site human responders about the extent of damage, the evolution of the event, the needs of the community and the present ability of the responders to deal with the situation combined with information from the larger emergency management community could lead to more accurate and real time situational awareness that allows informed decisions, better resource allocation and thus a better response and outcome to the total crisis.

DHS-S&T has just initiated the "Social Media Alert and Response to Threats to Citizens" (SMART-C) Program which fits within the bounds of the above DHS directive. This program aims at developing a citizen participatory sensing capabilities for decision support throughout the disaster life cycle via a multitude of devices (e.g., smart phones) and modalities (e.g., MMS messages, web portal, blogs, twitters, etc.) Specifically, the objective is to establish a bidirectional link between emergency response authorities and citizens that facilitates in: receiving early warning signals; detecting incidents and how they evolve; communicating alerts and advisories to citizens during and after the incident for response and recovery; and getting citizens' feedback for post-incident analysis and reconnaissance.

3. Challenges

Most currently available smart phones have been equipped with a variety of sensors, including GPS, accelerometer, gyroscope, microphone, camera and Bluetooth. This has been supplemented by new sensing applications across a wide variety of domains such as social networks, health, education, weather,

[1] Submitted to: the NSF/CCC (Computing Community Consortium) sponsored visioning workshop on Spatial Computing which outlines an effort to develop and promote a unified agenda for Spatial Computing research and development across US agencies, industries, and universities.

transportation, disaster management, gaming and entertainment. These applications and sensors built around smart phones and other devices (tablets, etc..) create huge volume of data with different modalities and types as listed in Table 1. Integration and analysis of such diverse and multi-modal data will help in observing and understanding the social media phenomena in our society and making further technological advances. However, there are several challenges that need to be address. Below we discuss three such challenges.

Table 1.

Data type	**Embedded Information**
Voice calls	Audio sample; caller/called number; date & time
SMS	Message transcript; caller/called number; date & time
MMS	Multimedia object (image, audio, video, etc); geo-tagged location; caller/called number; date & time
Social media feeds	Application type (e.g., twitter, facebook, etc.); type of event (e.g., posting or notification); media object (text message, video, audio, etc.); date & time
Geo-location data	GPS measurement of current location, accelerometer samples, gyroscope samples
Network connectivity data	Cell tower and WLAN access point observation and their location; Bluetooth observations

3.1 Event detection

Event extraction from unstructured data is an active area of research. Data from different sources when viewed in isolation may appear irrelevant, but when analyzed collectively may reveal interesting events [Ada07]. For the purpose of illustration, consider the following scenario in the context of disaster management.

Scenario : *Multiple residents post twitter messages about getting sick after eating at local restaurants in a given region (e.g., Southern New Jersey area) – the twitter feeds may reference different restaurants and may report different symptoms (e.g., fever, stomach ache, etc.). Based on these feeds, geo-spatial reasoning would be employed to automatically extract and characterize the event in both space and time. In this case the event is a health event and is progressing in the Southern New Jersey area. To assess the reliability of such event, the information from twitter feeds is corroborated with information from other sources such as hospitals, CDC alerts, and News feeds. This may also help in locating the source and likely cause of such event, e.g., outbreak of Salmonella. Based on assessed reliability of the event, local authorities would be alerted for further investigation. In addition, other restaurants in the region are also alerted as well as citizens (based on their location) who may have visited such restaurants or bought product from the local farm to seek medical help in case they develop related symptoms.*

As illustrated in the above scenario, some of the related challenges include:

- Integration and enrichment of multi-modal data (including unstructured data) from different sources. This becomes more complex when real time requirements are considered.
- Improved data quality is essential for robust event identification and characterization. In the spatial computing environment where data are often collected and assimilated automatically (e.g., from various type of sensors, social media) data quality (e.g., missing data, erroneous data, uncertainty, fidelity) issues are exacerbated.
- Validation and reliability of data are crucial to achieve higher accuracy for event identification and characterization.
- Semantic-based spatio-temporal reasoning for disambiguating events and tracking progression of events in space and time.

3.2 Data Privacy

The geo-spatial data retrieved from smart phones, sensors, and other smart devices often contains sensitive personal information. This data when combined with social media data significantly increases the risk to individual privacy breaches. The privacy concerns need to be addressed in all phases of spatial computing, including data collection, storage, analysis and dissemination. In the context of disaster management, social media (twitter, facebook, blogs, etc.) and mobile apps could be used for situational awareness and disseminating customized alerts and advisories based on users' location, language, and special needs. The challenge is how to achieve this targeted and customized alert and response while respecting individual privacy. For addressing this challenge, two inter-related issues need to be addressed: i) location privacy; and ii) protection of personal identifiable information (PII).

3.2.1 Location privacy

The current literature for location privacy can be categorized into following approaches: i) anonymization [Shin11, Liu09]; ii) mixing identifiers [Jad11]; iii) data perturbation [Hoh05]; iv) temporal obfuscation by adding random delays [Hoh07]; and v) and differential privacy [Che11]. However, such approaches have resulted in limited effectiveness with respect to data utility [CCC12]. The challenge here is how to achieve the right balance between location privacy and data utility? And how users can specify their privacy preference at the acceptable level while receiving the desired location based services.

3.2.2 Protection of personal identifiable information

There is a significant body of work addressing privacy of PII [Agg08, Zho08, Wan10, Ita09]. Most of this work, however, focuses on PII protection at the data storage and analysis phases. There is some work that addresses data privacy at the collection phase. This work is limited to specific application context, e.g., video surveillance [Wic04]; polling data [Gol06]. Given the large number of data sources and data modalities in the spatial computing environment, there is a need to develop new approaches for PII protection at the data collection phase. Moreover, such approaches need to take into account the real-time considerations for data collection.

3.3 Smart devices and the cloud

Today, smart devices, such as smart phones, tablets are connected to the cloud and use the cloud via RESTful web services for processing capabilities, storage, and security [Chr09, Art12]. This setting combined with the cloud constitutes a distributed global network. In this network, the cloud is aware of the state (e.g., idle/busy, battery, etc.) and resources (e.g., memory, computing power, etc.) of each device and the network topology in different geo-spatial regions. This environment present several research challenges, some being addressed in the context of traditional distributed computing and others are new that need attention, such as federated identity limitations on mobile platforms, discovering and composing services offered by smart devices (e.g., sensing services) [Chr09, Gar11].

Recently, a new generation smart devices is emerging with extensive computing power and memory. For example, the newly introduced inexpensive (within $200 range) 7-inch Goolge Nexus[2] tablet has Quad-core Tegra 3 processor, 1 GB RAM, 16 GB internal storage, and several sensors including, camera, microphone, accelerometer, GPS, magnetometer, and gyroscope. The powerful computing and memory of such devices extend their use beyond sensing to running computing tasks, especially when combined with the cloud. For example, can we use these mobile devices for Map Reduce jobs with the cloud provide the middleware for scheduling, coordination, and job migration (incase the device becomes unavailable due to user activity or network unavailability). In this environment the problem of discovery and composition of services offered by these smart devices and identity management is more challenging.

[2] http://www.google.com/nexus

References

[Ada07] N. Adam, V. Janeja, A. Paliwal, B. Shafiq, C. Ulmer, V. Gersabeck, A. Hardy, C. Bornhövd, J. Schaper: Approach for Discovering and Handling Crisis in a Service-Oriented Environment. ISI 2007: 16-24.

[Agg08] C. Aggarwal, P. S. Yu: A General Survey of Privacy-Preserving Data Mining Models and Algorithms. Privacy-Preserving Data Mining 2008: 11-52.

[Art 12] H. Artail, K. Fawaz, A.. Ghandour: A Proxy-Based Architecture for Dynamic Discovery and Invocation of Web Services from Mobile Devices. IEEE T. Services Computing 5(1): 99-115 (2012).

[CCC12] From GPS and Virtual Globes to Spatial Computing - 2020: The Next Transformative Technology. NSF/CCC Workshop Proposal (2012).

[Che11] R. Chen, B. Fung, B.. Desai: Differentially Private Trajectory Data Publication CoRR abs/1112.2020: (2011).

[Chr09] J. Christensen: Using RESTful web-services and cloud computing to create next generation mobile applications. OOPSLA Companion 2009: 627-634.

[Gar11] J. García-Macías, J. Alvarez-Lozano, P. Estrada, E. Avilés-López: Browsing the Internet of Things with Sentient Visors. IEEE Computer 44(5): 46-52 (2011).

[Gol06] P. Golle, F. McSherry, and I. Mironov: Data collection with self-enforcing privacy. In Proceedings of the 13th ACM conference on Computer and communications security (CCS '06): 69-78 (2006).

[Hoh05] B. Hoh and M. Gruteser: Protecting Location Privacy Through Path Confusion. In Proceedings of the First International Conference on Security and Privacy for Emerging Areas in Communications Networks (SECURECOMM '05). IEEE Computer Society, Washington, DC, USA, 194-205 (2005).

[Hoh05] B. Hoh, M. Gruteser, H. Xiong, A. Alrabady: Preserving privacy in gps traces via uncertainty-aware path cloaking. ACM Conference on Computer and Communications Security 2007: 161-171

[IBM11] IBM, Caltrans and UC Berkeley Aim to Help Commuters Avoid Congested Roadways Before their Trip Begins. Retrieved on July 31, 2012, from http://www-03.ibm.com/press/us/en/pressrelease/34261.wss

[Ita09] W. Itani, A. Kayssi, A. Chehab: Privacy as a Service: Privacy-Aware Data Storage and Processing in Cloud Computing Architectures. DASC 2009: 711-716.

[Jad11] M. Jadliwala, I. Bilogrevic, J.-P. Hubaux: Optimizing Mixing in Pervasive Networks: A Graph-Theoretic Perspective. ESORICS 2011: 548-567

[Liu09] L. Liu: Privacy and location anonymization in location-based services. SIGSPATIAL Special 1, 2: 15-22 (July 2009).

[Shin11] H. Shin, J. Vaidya, V. Atluri: A profile anonymization model for location-based services. Journal of Computer Security 19(5): 795-833 (2011).

[Wan10] Cong Wang, Qian Wang, Kui Ren, Wenjing Lou: Privacy-Preserving Public Auditing for Data Storage Security in Cloud Computing. INFOCOM 2010: 525-533.

[Wic04] J. Wickramasuriya, M. Datt, S. Mehrotra, and N. Venkatasubramanian. Privacy protecting data collection in media spaces. In Proceedings of the 12th annual ACM international conference on Multimedia (MULTIMEDIA '04): 48-55 (2004).

[Zho08] B. Zhou, J. Pei, and W. Luk: A brief survey on anonymization techniques for privacy preserving publishing of social network data. SIGKDD Explor. Newsl. 10(2) (December 2008).

EarthCube – Building Cyberinfrastructure in the Geosciences

A Whitepaper for the Spatial Data Computing Workshop
Washington, DC, September 10-11, 2012

M. Lee Allison, Chair, EarthCube Governance Steering Committee
Arizona Geological Survey, Tucson, Arizona

Overview

EarthCube is a process and an outcome, established to transform the conduct of research through the development of community-guided cyberinfrastructure for the Geosciences as the prototype for potential deployment across all domain sciences. EarthCube aims to create a knowledge management system and infrastructure that integrates all Earth system and human dimensions data in an open, transparent, and inclusive manner. EarthCube requires broad community participation in concept, framework, and implementation and must not be hindered by rigid preconceptions.

A fast-track process during spring, 2012 culminated in a Governance Roadmap delivered to the NSF-sponsored June charrette with an aggressive timetable to define and implement a governance structure to enable the elements of EarthCube to become operational expeditiously. The Governance Framework represents the implementation of initial recommendations laid out in the Governance Roadmap.
We discovered widely varying interpretations, expectations, and assumptions about governance among EarthCube participants. Our definition of governance refers to the processes, structure and organizational elements that determine, within an organization or system of organizations, how power is exercised, how stakeholders have their say, how decisions are made, and how decision makers are held accountable.

We have learned, from historical infrastructure case studies, background research on governance and from community feedback during this roadmap process, that other types of large-scale, complex infrastructures, including the Internet, have no central control, administration, or management. No national infrastructure that we examined is governed by a single entity, let alone a single governance archetype. Thus we feel the roadmap process must accommodate a governance system or system of systems that may have a single governing entity, particularly at the start, but can evolve into a collective of governing bodies as warranted, in order to be successful.

Our goal is to help ensure the realization of this infrastructure sooner, more efficiently, and more effectively, by providing a community endorsed Governance Framework. The Framework, and corresponding community outreach, will maximize engagement of the broader EarthCube community, which in turn will minimize the risks that the community will not adopt EarthCube in its development and final states. The target community includes academia, government, and the private-sector, both nationally and internationally.

Based on community feedback to-date, we compiled and synthesized system-wide governance requirements to draft an initial set of EarthCube governance functions and guiding principles. These functions will permit us to produce a Governance Framework based on an aggressive community outreach and engagement plan that we plan to finalize at the end of 2012.

Purpose

The overarching goals of EarthCube are to build a unified, adaptive, and scalable cyberinfrastructure framework for enabling transformative advances in geosciences research and education, thereby realizing the vision articulated in the NSF Geo Vision report.[1] In the process, EarthCube aims to create a knowledge management system and infrastructure that integrates all Earth system and human dimensions data in an open, transparent, and inclusive manner.

Developing a viable organizational and governance structure for any organization can be a challenge. Creating one for multi-disciplined, distributed, virtual collection of scientists, investigators, technologists, system operators, entrepreneurs, and administrators can be nearly impossible unless great care is taken to ensure that the proposed solution is flexible and responsive to meet participant's needs and institutional goals.

We believe that there is general agreement that "effective governance for EarthCube will:

- actively engage its diverse users
- provide leadership and oversight to forge close cooperation, coordination, and collaboration among distributed development activities and the principal EarthCube groups
- facilitate alignment of funding program plans and priorities with the needs of the community
- help the successful execution of the EarthCube mission, meeting stakeholder obligations"[2]

To be effective, the governance framework the community adopts is likely to be for a *system* of governance (a matrix of mechanisms for different elements and groups) that accommodates different practices and requirements among different elements of a large and diverse community. The governance roadmap also allows for a variety of mechanisms for how the governance mechanisms are chosen and implemented.

Challenges

The challenges we considered were not just to creating the governance roadmap per se but also to the role and impacts of a governance process and system on the overall viability and success of EarthCube as a community system. Challenges to the roadmapping process are inherent given the limited time frame. Among these challenges:

- Comprehensive background research review of governance topics from the domain sciences, IT, and social sciences is not yet complete.
- We identified many governance models, but have not been able to fully evaluate them.
- Further work is needed to evaluate the pros and cons of different models and determine which may be suitable for EarthCube.

---

[1] National Science Foundation, Advisory Committee for Geosciences, "GEO Vision Report." October 2009.

[2] Mohan Ramamurthy, "Unidata Governance: A Quarter Century of Experience," National Science Foundation EarthCube White Paper: Governance Category, 2011, 1.

- Our knowledge of the other EarthCube Working Group and Concept Team governance issues and needs is not yet complete.
- We have yet to fully engage the broader Earth, information, and IT science communities, thus our knowledge of their governance needs is limited.
- There is limited information available about problems and failures of past projects that we can incorporate as things to avoid.

Challenges to the viability of EarthCube were generated by community feedback and the governance research review. We divided them into:

- Conceptual and procedural challenges:[3] Time (short-term funding decisions versus the long-term time-scale needed for infrastructures to grow); Scale (choices between worldwide interoperability and local optimization); Agency (how to navigate planned versus emergent change), intellectual property rights, infrastructure winners and losers, agreement on data storage, preservation, curation policies and procedures, incentives to share data and data sharing policies, and trust between data generators and data users.
- Social and cultural challenges: Motivations and incentives, self-selected or closely-held leadership, levels of participation, types of organizations, and collaboration among domain and IT specialists)
- Technical challenges: From governance use cases.
- Trends and drivers: Federal government initiatives, cloud computing, internationalfforts such as the EU INSPIRE initiative, Australian National Data Service, etc, and commercial developments.

Requirements

To continue forward, we recommend building upon the process of community engagement and research review begun as a cornerstone of the Governance Roadmap process to identify and characterize the components of cyberinfrastructure. Community engagement is expected to occur in four steps (for a full description and graphic showing the progression of engagement seethe Governance Roadmap http://earthcube.ning.com/group/governance/forum/topics/earthcube-governance-roadmap-version-1-1):

- Identify cyberinfrastructure components of EarthCube
- Identify the cyberinfrastructure components' organizational paradigms and governance need.
- Identify the interaction among and between cyberinfrastructure components and systems within EarthCube.
- Identify the interactions between cyberinfrastructure components within and outside of EarthCube, and the needs of EarthCube consumers (including those comprising the "long tail" of science).

[3] Paul Edwards, Steven Jackson, Geoffrey Bowker, and Cory Knobel, "Understanding Infrastructure: Dynamics, Tensions, and Design - Report of a Workshop on "History & Theory of Infrastructure: Lessons for New Scientific Cyberinfrastructures," 2007, 24-33.

From GPS and Virtual Globes to Spatial Computing – 2020

This paper addresses aspects of computing related to analyzing and using spatial data. The variety and volume of data with spatial content afford us many opportunities to understand the world. They also challenge us to find effective means of "chaff removal" and of capturing and using relationships between data that is independently collected. Meeting this challenge will require progress on several fronts, starting with developing new ways to estimate location or to verify it. Such matters are often discussed with respect to social media; however, many modes of modern communication include latent, hidden, or indirect information that we could use—and we need to discover how to find them.

We need expanded approaches to discovering useful patterns in large spatial data sets, particularly data sets that reflect activity, behaviors, or movement, and to use them with complex data sets comprising billions of instances, such as large, dynamic graphs or collections of trajectories. One of our recent efforts reduced 400 quadrillion (10^{15}) available data relationships to 10 million relationships of potential interest; we currently have larger data collections of comparable complexity to analyze. We are especially interested in methods for parallel processing where the data contain many relationships and are not amenable to widely-used methods of partitioning.

We want to record and manipulate data about people, places, and activities not directly tied to the surface of the earth, e.g., sub-surface data (both objects and attributes), things that exist in buildings, tunnels, under water, in cyberspace and in hypothetical worlds. We need to capture and manipulate movement, change, and activities both by type and by instance (e.g., planned vs. actual routes or schedules), and to support generation and comparison of "geospatial narratives," such as boats leaving and entering ports, or staging and transport of supply chains, along with identification and monitoring of trends in quantifiable data. For many applications, all of the objects being represented (e.g., road network; transit schedule; vehicle trajectories) are subject to change, and we want the ability to invoke the state of affairs at any given point in time.

We seek new methods for using "related data" to validate or quantify reliability of data of unknown provenance or uncertain suitability for a task at hand. For example, there may be ways to compare volunteered place names with names used in social media and commercial or government publications to establish whether usage patterns are consistent with provider claims about connotations associated with the choice of one name over another. Do empirical data support reports that residents of the D.C. metro area generally refer to "the District" rather than "Washington." It will be useful to have an inventory of established, validated, methods that can be widely used, including methods to assign a "reliability score" to a source based on disposition of previous submissions.

We are looking for expanded capabilities to reason about data sets and for more effective ways to relate "discovered meaningful data" to other known or posited data. Such capabilities will support use of data reflecting different scales and accuracies, including changes in scale or accuracy across a single data set. They will be essential to reducing the search space for computational purposes and for human comprehensibility.

We continue to search for methods to integrate and conflate data, and some of the work described above might be applied to that end. Our goals include integrating about activities as well as data about places and things. We need to look at consistency across multiple sources at the object level, and we need to look for consistency between attributed features.

We want to record and examine variability in geometry, topology, and attribution, whether they reflect observed instances or rule-governed behavior. Examples include spatial footprints that vary with time of day (e.g., "high crime neighborhood"); attributes that vary (e.g., "dominant language" may vary temporally); even topology may vary (e.g., which side of the street the busses stop on may vary temporally). Capturing "general rules" that are spatially or temporally dependent (e.g., rules that apply to all instances within a municipality or during daylight savings time), and applying them efficiently will be important for data maintenance and verification and for establishing relationships that are valid. In addition to surmounting processing constraints, it will be important to overcome manageability issues that impede use of rule-based systems today.

In order to support responsible use of complex or highly-processed data, we need improved understanding of how to foster comprehension or to mislead with outputs of startling beauty or dizzying multi-modal effects. Progress on this front is essential not only for "end users" but also for researchers and analysts to view their own work critically. We need sophisticated-yet-simple methods to allow users to interact with data "in context," in large-scale settings and small ones, where "context" may refer to uncertainty, processing history, social/political/cultural/ or economic events, space, or other dimensions.

I have not addressed many important issues, such as sensors and direct exploitation of sensor data. Applying spatial computing to more heterogeneous collections of data will require algorithms to extract semantic content from "unstructured" data and translate it to a formal representation that supports rigorous manipulation, along with the use of relative as well as absolute positioning in time and space. I have not even mentioned privacy. I hope that other participants will elaborate on such topics.

Respectfully submitted --

Beth H. Driver

Spatial Surrogates to Forecast Social Mobilization and Civil Unrests[1]

Feng Chen[1,2], Jaime Arredondo[5], Rupinder Paul Khandpur[1,2],
Chang-Tien Lu[1,2], David Mares[4], Dipak Gupta[3], and Naren Ramakrishnan[1,2]

[1]Department of Computer Science, Virginia Tech
[2]Discovery Analytics Center, Virginia Tech
[3]Department of Political Science, San Diego State University
[4]Department of Political Science, University of California San Diego
[5]Latin American Studies Program, University of California San Diego

Introduction

Spatial computing is now widely pervasive in the engineering and science disciplines but we argue that there is an even bigger revolution happening in our ability to comprehend human behavior. Modern geo-tagged communication forms such as social media and microblogs are rapidly advancing the methods by which we can comprehend, and even influence, the progression of events as they unfold. The rise of "massive passive" data (e.g., tweets), in particular, has given significant impetus to being able to understand events across the globe.

Two key trends are manifest in the above developments. First, it has become possible to use public-domain, seemingly innocuous, aggregated data, to infer quantitative indicators of population level change. At the same time, as the scope of such inference enlarges, novel computational methods are becoming imperative for fusing data from such high-throughput sources. This position paper argues for a concerted effort to use "spatial surrogates" as an enabling mechanism to model and forecast social mobilization across the world.

Spatial surrogates are data reductions that we can exploit to aid in understanding population-level phenomena. As the name indicates, surrogates are cheap, easy-to-compute, statistics that are correlated with or that precede phenomena of interest. Surrogate modeling is an established practice in numerous domains such as multidisciplinary optimization and economic forecasting, and here we argue for its use in modeling key societal events. For instance, the idea of tracking flu activity geographically using search query data (in Google's FluTrends) is a modern example of knowledge discovery using surrogates. A second example is using spatial luminosity data to quantify economic output of countries [Chen and Nordhaus, 2010]. A final example is using Landscan data as a surrogate for population density.

Social Mobilization

Our domain of interest is social mobilization, i.e., how civilian populations mobilize to raise awareness of key issues or to demand changes in governing or other organizational structures. Protests, strikes and "occupy" events are part of such mobilizations. Such events occur in a variety of political systems, even authoritarian ones. Most of the time governing institutions can ignore, repress, or respond to social mobilization in ways that do not fundamentally challenge public policy and the political institutions that generated those policies. Nevertheless, at times

---

[1] Supported by the Intelligence Advanced Research Projects Activity (IARPA) via Department of Interior National Business Center (DoI/NBC) contract D12PC00337. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of IARPA, DoI/NBC, or the US Government.

governing institutions, even democratic ones, are overwhelmed by civil unrest generated by significant and repeated protests (e.g., Chile 1970s, Poland 1980s, Bolivia 2000-2005, Arab Spring 2010-present).

Scholars of democracy, policymakers and most social activists are aware that significant levels of civil unrest make politics, economics and social relations difficult and unstable. Any efficient solution to civil unrest has to find ways to channel political power in the streets back into stable, institutionalized channels of interest representation where bargains can be negotiated. This requires understanding when significant social mobilization may occur and when that civil unrest will be of significant size and its likelihood of becoming violent.

Decades of intense social science research have shed important light into our understanding of the root causes of social mobilization and political violence. However, challenge for the scholars and policymakers alike, is the ability to forecast these events. Traditional empirical academic research develops forecasting models based on past data, collected and published by various governmental and non-governmental agencies. Yet, these efforts are time consuming and the resulting inquiries are akin to the astronomers looking at distant galaxies, which contain information about how they used to be rather than what they are now. The monitoring of the Internet and the social media has opened up a brand new area of social inquiry that is unique not only in it aspiration but also its inherent need to bring together scholars from all areas of academia like never before.

Studying and Forecasting Civil Unrests

In general, computational modeling of civil unrests is in its infancy. Recent research [Gonzalez-Bailon, 2011] has focused on how protest recruitment happens through an online network but comparatively little attention has been paid to forecasting civil unrests in society through information gleaned from online, geo-tagged, media.

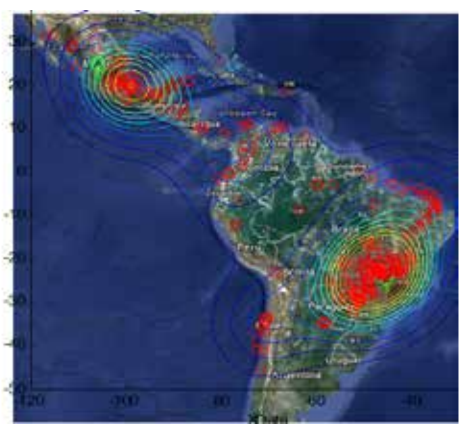

We have taken some initial steps along this direction by first organizing a dictionary of 726 terms related to protests, and redescribing a geo-tagged tweet stream in terms of frequencies of terms from this dictionary. The objective is to identify anomalous spatial regions based on Poisson mixtures. A linear time subset scan [Neill, 2012] is applied to identify anomalous spatial regions which are then scored using p-values computed by Monte Carlo simulation.

Our work focuses on countries in the Latin American region; the displayed map from June 30, 2012 provides an illustrative example where each (geo-tagged) tweet containing at least one term from our dictionary is plotted as a small red circle. Two anomalous spatial clusters are detected, as shown. One cluster is located in Mexico with 'país', 'trabajador', 'trabaj', 'president', and 'protest' as the top five frequent terms. These refer to the student-led protests that happened during the Mexican election held on July 1, 2012. The second cluster, located in Brazil, involves the high frequency terms: 'país', 'protest', 'empres', 'ciudad', and 'gobiern'. This cluster is related to the situation where approximately 2,500 people closed the Friendship Bridge linking Ciudad del Este (Paraguay) and Foz de Iguazu (Brazil), a demonstration held in support of Paraguay's president Fernando Lugo. Thus, initial results are encouraging.

Research Issues for Discussion

Spatial Surrogates: There are now a significant number of spatial data sources available, especially through the advent of location-based social networks such as Facebook, Twitter, and Foursquare. How can we leverage such a multiplicity of data sources to design accurate spatial surrogates? Although each data source by itself is unlikely to provide the desired specificity, it is possible that combinations of them will yield the desired quality of forecasting.

Machine Learning Models of Spatio-temporal Phenomena: Traditional models of spatial and spatio-temporal phenomena have been prohibitively expensive, e.g., involving the estimation of non-stationary covariance matrices. Modern methods such as the linear time spatial scan [Neill, 2012] promise to usher in significantly more efficient methods for detection. Can we establish an emphasis on both efficient and expressive algorithms for machine learning research?

Spatial Event Forecasting: Most current research focuses on event detection in the form of spatial or temporal bursts or clusters whereas the forecasting of events has not been well studied. However, significant domain knowledge can be harnessed in the form of how mobilization occurs on a spatial or temporal scale. How can machine learning algorithms exploit such prior knowledge effectively for spatial forecasting of civil unrests?

Geolocation: Only a small percentage of communications data harvested from social media are geo-tagged natively but it is possible to envision semi-supervised and transfer learning paradigms that enable a greater variety of data sources to be geo-tagged. What data sources provide corrobative and complementary evidence for geotagging purposes?

Integration of Crowdsourcing and Machine Learning: Concomitant with better geotagging capabilities, it is instructive to examine how a modicum of "active" crowdsourcing can augment "passive" data assimilation, and how such a data gathering loop can be integrated with a machine learning loop. This can yield systems that can systematically increase specificity of modeling by crowdsourcing data gathering in regions of most uncertainty.

Integrated Crisis Management: Finally, integrating the methodologies above for civil unrest modeling can lead to a powerful system for integrated crisis management, one that can quickly disseminate information spatially in the most efficient manner and reduce congestion and overload both in physical and in communication infrastructures.

References

X. Chen and W.D. Nordhaus, Using Luminosity Data as a Proxy for Economic Statistics, *PNAS,* May 16, 2011.

S. Gonzalez-Bailon, J. Borge-Holthoefer, A. Rivero, and Y. Moreno, The Dynamics of Protest Recruitment through an Online Network, *Nature Scientific Reports,* Vol. 1, Article 197, Dec 2011.

D.B. Neill. Fast Subset Scan for Spatial Pattern Detection. *Journal of the Royal Statistical Society (Series B: Statistical Methodology)* 74(2): 337-360, 2012.

Spatial Computing at the Topographic Engineering Center

Dr. James Shine

The US Army Engineer Research and Development Center's (ERDC) Topographic Engineering Center (TEC) is located in Alexandria, VA. Its mission is to provide the warfighter with superior knowledge of the battlefield, and to support the nation's civil and environmental initiatives through research, development, and the application of expertise in the topographic and related sciences. Since our work in the present day and age is largely fueled by computers, and since topographic data by definition has a spatial component, we perform a wide variety of spatial computing work in both the basic and applied research areas.

Our basic research portfolio consists of a limited number of small projects, 1 to 3 years in length, involving one or two government researchers and usually a partner in academia. We had two highly successful efforts involving spatial computing with the University of Minnesota which are now completed. The first one was entitled "Modeling Spatio-Temporal Co-occurrence patterns" and looked for object types that were within a specified neighborhood of each other. The project required the development of new spatial algorithms to perform the necessary computation and search. The second project was entitled "Spatio-Temporal Cascade Patterns" and again looked at events in certain locations and at certain times. However, this effort was looking for a starter event and then other related events at later times that "cascaded" from the original event. Again, spatial algorithms were developed, and the level of complexity required special attention to be paid to efficient search, as the search space was very wide.

Currently, we have several basic research projects at TEC. One involves looking at social media network data and trying to define a socially aware spatial spatial neighborhood. These neighborhoods allow for the exploration of the spatial and social distribution of communities. A combination of geographic information system software and other analytical software is being used to develop new models and metrics to map this previously unknown space.

Another effort seeks to store and organize large spatiotemporal data in a hierarchical fashion where each hierarchy corresponds to a different spatial resolution while minimizing the information loss when the data resolution is reduced. In order to minimize the information loss, this effort investigates new preservation techniques for various geographic features (e.g., topology) that are salient to tasks such as line of sight modeling. A related effort focuses on developing efficient techniques for mining massive spatiotemporal data. In this work unit, spatiotemporal data are transformed into symbolic sequences and techniques from statistical language processing are applied to efficiently mine patterns. Another effort uses spatial data from unmanned ground vehicles to compute properties of the surrounding terrain. There is also a work unit which analyzes spatial LIDAR data using an innovative computational approach.

In the applied research area, we have several projects which employ spatial computing methods. One project, Grapevine, is developing a geographic information retrieval and knowledge discovery system that focuses on creating new relevancy ranking algorithms and scalable spatiotemporal indexing methods for unstructured data. The search results will be analyzed ad-hoc by statistical learning techniques to discover relevant patterns. New statistical learning algorithms will be developed to rapidly process the (potentially large) search results in a real-time environment. Another project models socio-cultural data based on spatial parameters. As in the basic research social network project, there is lots of room to explore concepts such as water geography and cultural neighborhoods. This project analyzes field data for spatial patterns.

DISRUPTIVE TECHNOLOGIES IN SPATIAL COMPUTING

Position Paper Author: Nyeng Paul Gyang, Doctoral Student (Computer Science), Colorado Technical University (CTU), Colorado Springs, CO.

Date: Sunday, August 05, 2012.

Position Paper Title:

Novel Techniques for the Generation of DEM from LiDAR Point Cloud Data

LiDAR point cloud data sets are typically appreciably large such that processing them can use means of speeding up computer programs, including parallel processing using parallel computing architectures and systems like multi-core, many-core and GPU systems. The transition from point cloud data to a DEM comprises of a range of steps in a procedure involving a couple of major techniques, namely: Interpolation and surface reconstruction. There is no doubt that the processing of point cloud data in order to transition to a DEM can benefit from techniques for speeding up the algorithms for both interpolation and surface reconstruction – it is in light of this realization that it will be advisable to investigate the ways by which these algorithms may be parallelized and targeted for architectures and systems that are both requisite and apt for this purpose/application.

This researcher plans to work towards devising a parallel programming model that will enable and facilitate the development of parallel algorithms and writing of parallel code for DEM generation from point clouds using architectures and systems that are particularly suited for this application, especially shared memory architectures and systems such as multi-core, many-core and GPU systems. This research effort should be able to make significant contributions to the generation of digital models, which represent 3D surfaces in real life, ranging from bare terrain and covered terrain (including covering by vegetation, buildings/constructions, etc.) on planet earth to bare terrain on other worlds, such as planets, moons of planets, etc. Other applications that are not so far-fetched – i.e. closer to our everyday lives – include the generation of 3D Computer Aided Design (CAD) models for manufactured parts, metrology/quality inspection and a multitude of visualization, animation, rendering and mass customization applications (Wikipedia, 2012).

Evidence in the literature demonstrates that the trend of a shift towards the multi-core architecture is both real, existent and, in the foreseeable future at least, will be a key part of the dominant parallel processing architecture (Gepner and Kowalik, 2006), (Jin *et al.*, 2011), (Jost and Robins, 2010) and (Zhang *et al.*, 2007). Furthermore, it is envisaged and felt that the multi-core architecture will be combined in a hybrid architecture comprising of multiple nodes (with distributed memory across the nodes), each of which will consist of one or more multi-core processors (with shared memory within a single node). The envisaged future dominance of the multicore and many-core architectures will be grossly under-exploited, except the effort is exerted in order to fully exploit the power furnished by multi-core and many-core hardware in applications, including the generation of DEMs from point clouds; this presents a research opportunity as well as an open problem or challenge. For example, Guan and Wu (2010) note that their parallelization of the generation of DEMs with multi-core processors demonstrates the great potential of this parallel architecture and system for this Spatial Computing task.

This researcher imagines, envisages and hypothesizes that the contribution, to the proposed project on Spatial Computing, of the activity of the generation of DEMs from point clouds (using, particularly, multi-core and many-core systems), is not trivial – this proposed project, by the way, enshrines and is directed by the "unified agenda for Spatial Computing research and development across US agencies, industries and universities."

References

Gepner, P. and Kowalik, M.F. (2006). Multi-Core Processors: New Way to Achieve High System Performance. *Proceedings of the International Symposium on Parallel Computing in Electrical Engineering (PARELEC'06)*

Guan, X. and Wu, H. (2010). Leveraging the Power of Multi-core Platforms for Large-scale Geospatial Data Processing: Exemplified by Generating DEM from Massive LiDAR Point Clouds. *Computers & Geosciences* 36 (10), 1276-1282.

Jin, H., Jespersen, D., Mehrotra, P., Biswas, R., Huang, L. and Chapman, B. (2011). High Performance Computing Using MPI and OpenMP on Multi-core Parallel Systems. *Parallel Computing* 37 (9), 562-575. doi:10.1016/j.parco.2011.02.002

Jost, G. and Robins, B. (2010). Experiences using hybrid MPI/OpenMP in the Real World: Parallelization of a 3D CFD Solver for Multi-core Node Clusters. *Scientific Programming* 18, 127–138. DOI 10.3233/SPR-2010-0308

Wikipedia, The Free Encyclopedia (2012). Point Cloud. *Retrieved from: http://en.wikipedia.org/wiki/Point_cloud*

Zhang, L., Y. Han, Huawei, L. and Xiaowi, L. (2007). Fault Tolerance Mechanism in Chip Many-Core Processors. *Tsinghua Science and Technology* 12 (Supplement 1[0]), 169-174.

Working in Virtual Spaces: Spatial Interfaces and Visualizations for Data Analysis and Creative Design

Daniel F. Keefe
Department of Computer Science and Engineering
University of Minnesota

The ability to *picture* and *interact with* concepts in new ways has always been intrinsic to the process of discovery. Muybridge's classic stroboscopic photographs of horses led to the discovery that all four of a horse's hooves leave the ground during a gallop; at the time this hypothesis was called, "unsupported transit". Da Vinci's hand-drawn studies of rushing water informed not only his art, but also the science of hydrodynamics. Today, engineers, scientists, and artists routinely rely upon physical models and 3D prototypes – often it is the physical act of touching, rotating, and annotating these models that brings forth new insights. Imagine if all these visual, physical, spatial human activities could take place in a virtual space, where powerful computational techniques could be combined with natural human interactions and visual communication. This could enable exciting new methods of computing that enhance creativity, enable discoveries, and open up countless new applications in the sciences, engineering, art, and design.

To realize the anticipated benefits of a powerful next generation of spatial interfaces and visualizations that can improve discovery across many fields, several major research challenges must be addressed. The following sections are organized around three of these challenges. Through describing the challenges along with recent examples of successful work in each challenge area, I hope to convey both the importance of continued research in each area and also the strong potential of new tools in this style.

Challenge 1: Create effective new methods for visualizing multidimensional time-varying spatial data.

The role of computing in society is increasingly defined by the way that computing enables access to and new ways of working with big data. Large spatial datasets (e.g., climate and population modeling data; medical imaging collections, simulations, and anatomical models) are some of the most challenging to analyze. Due to the complex spatial relationships in these data and the typical benefits of 3D visualization, it follows logically that visualization should be an effective tool for analyzing these data. However, we lack appropriate visualization tools for working with the complex spatial data that are now collected and generated. One reason is that the data that are most interesting for advancing science and engineering typically contain not only complex spatial information (e.g., geological features in a climate model) but also a wealth of other complementary data values (e.g., temperatures, wind speeds, cloud formations, and many other quantities); and, very importantly, all of these data may change over time. The resulting multidimensional time-varying spatial data visualization problem is extremely complex and is likely to be solved only through a holistic approach to visualization that draws upon knowledge from a variety of computer science sub-disciplines as well as related fields, such as psychology and cognitive science. Since the data are simply too large to be viewed directly, it is essential to develop new automated algorithms to mine these data that can work side-by-side with visualization to identify and then display the data to the user at varying levels, supporting overviews of large collections of data, clustered analysis, and detailed comparisons. To convey these data to users, tested, effective computer graphics algorithms need to be developed and optimized to best leverage human perception of 3D space and time. This is a major undertaking since creating a 3D visualization environment to convey just one or two data dimensions is itself a major challenge. Currently, it is not clear how best to design interactive visualization techniques for working with 10's to 100's of data dimensions visualized in changing spatial contexts; new research is needed to advance algorithms and techniques for both underlying data mining/management and perceptually accurate 3D data graphics.

An example of visualization research that begins to address these challenges comes from my group's recent efforts to visualize data that exist in complex 3D spaces and that also change importantly over

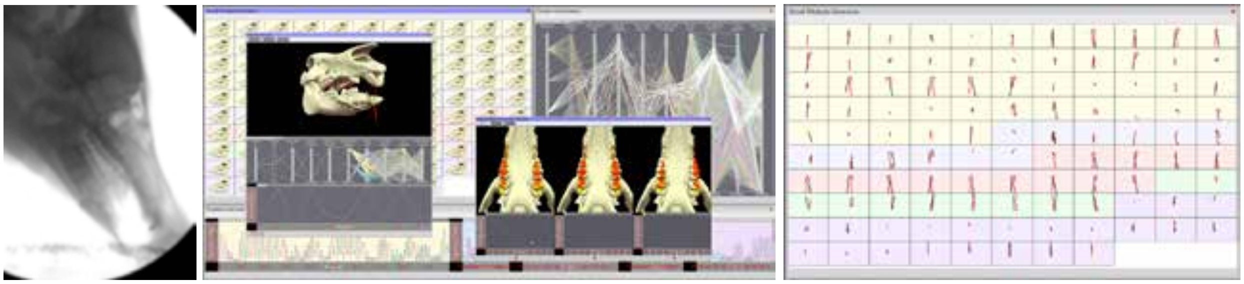

Figure 1: Interactive visualization system used to analyze > 100 high-res pig chewing motions collected by evolutionary biologists at Brown University studying historical diversification among animals. Video: http://ivlab.cs.umn.edu/generated/pub-Keefe-2009-MultiViewVis.php

time, such as the detailed moving relationships of bones and tissues studied in biomechanics. Here visualization is challenging because human perception of both 3D space and motion is complex. Figures 1 and 2 show two multidimensional spatial data visualization systems that have resulted from this work. The system in Figure 1 is the first interactive computer graphics tool to provide visual overviews of multidimensional data for an entire database of motions with scientific relevance. One of the primary goals of this system is to enable comparative analysis of large collections of motion data. Using this tool, scientists can filter and query multiple dimensions of the data to identify common patterns and anomalies, which are then examined in detail using interactive 3D visualization windows. This work is significant because visualizing just a single instance of these motions is challenging, but what scientists need today is a new way to visualize and compare motions across a large collection – this is much more difficult, and the tool shown here takes an important step in this direction. Figure 2 shows a next generation of these visualization strategies applied successfully to a completely different domain, motion data collected during minimally invasive surgery training exercises. Another current application targets visualizing in virtual reality tens of thousands of frames of spinal kinematics motion collected across patients with varying degrees of back pain. The motions of the vertebra are so spatially complex that some form of 3D visual analysis is necessary, but current clinical approaches (e.g., 2D statistics and videos) are not feasible for today's data-intensive studies.

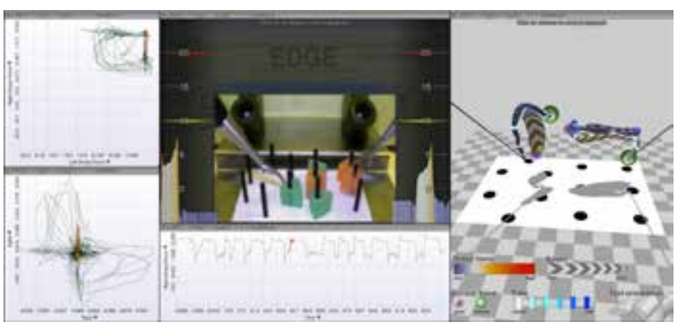

Figure 2: Visualizations of laparoscopic surgical training data. For the first time, surgeons can "see" their use of force, enabling new feedback and objective, data-driven evaluations of skill. Video: http://ivlab.cs.umn.edu/generated/pub-Schroeder-2012-Surgvis.php

Figure 3: Designing medical devices in a virtual heart. Touch gestures are used to investigate a 3D heart model created from imaging data. Video: http://ivlab.cs.umn.edu/generated/pub-Coffey-2011-InteractiveSliceWIM.php

Challenge 2: Enable new spatial computing workflows and boost the power of spatial visualizations through effective 3D user interfaces.

A major challenge in developing effective spatial data visualizations (e.g., within virtual reality environments) is supporting "real work" in these environments. This means enabling scientists to not only see their data in new ways but also to query, interrogate, and fluidly explore the data to both answer questions and generate new hypotheses. This is a challenge because the metaphors and techniques that we typically use to interface with computers (e.g., windows, icons, keyboard, mouse) do not translate effectively to spatial visualization environments. New spatial interfaces are needed. Intuitively, what we desire in spatial visualization environments are new ways of interacting with computers that are also spatial, matching the dimensionality of the visualization. In practice, new technologies (e.g., depth cameras, touch devices) show promise for this type of natural interface, but their application has been limited primarily to entertainment and home use. For science and engineering, precise inputs are needed for working with complex data. If we could create computer interfaces that are as fluid and natural as those we see today in games, phones, and other emerging devices but also support the precision and rich inputs needed to work with science and engineering applications, then our ability to *work* effectively with spatial data would be radically improved.

Some examples of research that aim to address this gap come from current work in my group. Figure 3 shows an interactive visualization of a virtual heart. Imagine an engineering team, which has designed a new medical device and wishes to analyze the change it makes in blood flow in the heart. The team asks, "Can we see a relationship between pressure and velocity in *this 3D region*". While defining "this 3D region" in a medical imaging dataset could take several hours using a keyboard and mouse, the novel 3D multi-touch visualization interface shown here makes it simple to rotate, zoom, and select anatomical structures or volumes of fluid flow; new 3D selections can be made in real time simply by moving ones fingers on top of the imaging data (see Figure 3 video). The approach uses custom-developed virtual reality hardware to make 3D renderings of data (e.g., a heart) appear to float in the air above the interactive table. By touching the shadow of (or a 2D slice through) the 3D data, which is projected onto the table, engineers can slice through the imaging data, plot 3D curves, measure volumes, and perform other intuitive physical operations in virtual space. This enables scientists to perform tasks, which used to only be possible in offline batch modes, now in real-time via visual workflows that integrate physical actions with virtual spaces, supporting creative spatial thinking and design processes.

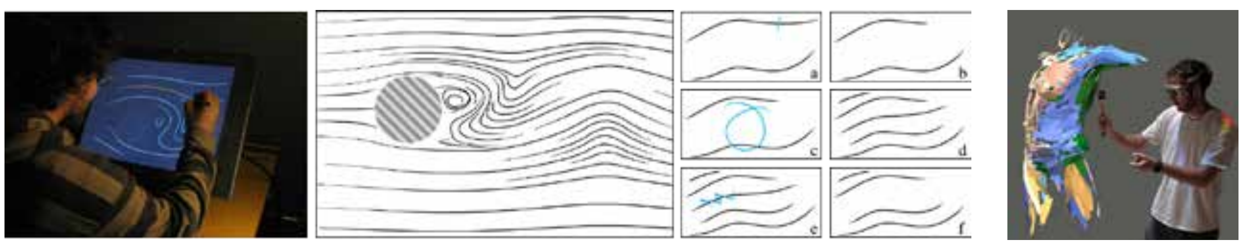

Figure 4: Interactive spatial visualization tools that enable collaboration between artists, designers, and scientists.
Videos: http://ivlab.cs.umn.edu/generated/pub-Schroeder-2010-DrawWithFlow.php
http://ivlab.cs.umn.edu/generated/pub-Keefe-2008-ScientificSketching.php

Challenge 3: Create new and expand existing exciting applications of interactive visual spatial computing across disciplines.

A final important challenge is to expand our thinking about spatial computing to embrace the broad wealth of potential applications of this area of computing. A focus on spatial visualizations and interactive techniques is especially important in this endeavor because these computational tools make spatial computing accessible to all segments of society, including K-12 educators and students, artists, and other creative minds.

Figure 4 shows two systems that combine spatial visualizations and interfaces to make computing accessible to artists, and have opened up new roles for artists in science. The system pictured in the left of Figure 4 places the intent of a traditionally-trained illustrator within the constraints implied by an underlying fluid flow dataset to produce accurate stylized hand-drawn renderings of flow patterns. The rightmost image portrays an innovative 3D tool for creating virtual sculpture by "painting in the air", used both for art practice and as a valuable sketchpad for prototyping 3D scientific visualizations. These new methods enable artists and designers trained in visual depiction to work creatively with the latest computer graphics technologies – without any knowledge of programming. Thus, interactive spatial visualizations can support scientific data analysis, creative engineering design processes, and new artistic explorations. There is great potential for advances in spatial computing to impact all segments of society.

Bridging a Spatial Data Gap: Incorporating Small-Scale Models into Large-Scale Systems
Position Paper for CCC Workshop: From GPS and Virtual Globes to Spatial Computing – 2020
John Keyser, Texas A&M University

As spatial data has become more ubiquitous, there has been a push to have additional detail available at smaller scales. Examples of this range from the desire for ever-higher resolution photos and GIS data to the use of "Street View" in Google Maps. At the same time as "large scale" mapping systems and applications built on them have improved, geometric modeling of smaller-scale data has become more common, with a variety of methods being used for capturing data. However, these two forms of data collection and storage have remained largely separate, with work on modeling small-scale geometric data staying largely separate from that used in large scale spatial data systems.

I believe there are a number of interesting research opportunities, with significant potential benefits, in efforts to effectively bridge the gap between spatial data at the larger scale with the geometric data that occurs at smaller scales. By "small-scale" geometry in this context, I am referring to information in scales measured from a few centimeters (e.g. hand-held objects) to a few meters (e.g. a vehicle or tree).

Significant Challenges with Smaller-Scale Geometry
There are several ways in which smaller-scale spatial and geometric data presents challenges that are less frequently seen as issues at larger scales. These include:

- *True 3D Nature of Geometry:* At larger scales, 2D or 2.5D (2D plus height) data is often sufficient for describing spatial aspects of geometry. However, at smaller scales the true 3D geometry of the spatial structure becomes more important. The topological connections between spatial locations are critical to understanding the spatial structure of an object or region, and much of the spatial computation with smaller-scale data will need to account for actual 3D geometry, rather than 2D footprints.
- *Occlusion and limited data:* The scale of the data, the 3D nature, and the typical positioning of small-scale objects and features in the world mean that spatial information will often be incomplete. Even with specific attempts to scan an object or a limited region, there will often be certain areas that are occluded or inaccessible. Thus, understanding the full spatial data will require some level of inference.
- *Changing Geometry:* Larger scale features tend to change infrequently or over longer periods of time. Small-scale data will often be mobile (e.g. vehicles), portable (e.g. furniture), or inherently changing (e.g. plants). This has implications for not only the representation of object geometry itself, but also for how to describe what might be very dynamic spatial information in a larger region.
- *Variety of Data Sources:* There are a wide variety of ways in which small-scale spatial data is being captured. This includes scanners designed to capture such data (e.g. laser scanners), image-based capture ranging from calibrated cameras to casual photographs, implied spatial information (e.g. from collections of GPS coordinates in a small area), and CAD models. Coming up with consistent and accurate models can be a challenge given this range of possible inputs, and the noise and error associated with each.
- *Size of Data Sets:* The raw amount of detail that can be provided when smaller-scale data is incorporated is potentially much higher than that available at larger scale. This means that simple data storage and access can become problematic if significant small-scale detail is maintained over a large region.

Research Opportunities

Combining small-scale spatial data with larger-scale systems can lead to richer, more accurate, and more useful spatial data systems. As one example of how this could be useful, consider an evacuation scenario from a structure or region. While some simulated behavior could be treated abstractly, having a better understanding of the small-scale details in the region would lead to more accurate models of evacuation behavior. The same would be true for many autonomous navigation or motion planning tasks, thus allowing the merging of the local and global phases of motion planning and navigation that are typical in current systems. As another example, having small-scale detail can help provide orientation and realism in VR/AR applications. To deal with the problems that this merger can create, research will need to be conducted from a variety of different fields.

There are a variety of specific research problems and major challenges that would need to be overcome to effectively integrate small-scale geometric data into large-scale applications and expand the range of possible applications for spatial data. A few of these are:

- *Capturing environments from massive data sets:* The expansion of sensor data (e.g. smartphones with both GPS and cameras) is producing very large amounts of data that could be mined for better understanding of spatial environments (as well as behavior within those environments). This presents a number of challenges, however, ranging from collecting the data (with the appropriate social and ethical issues that raises), through registering the data within a larger world, to geometric reconstruction from partial views and inference of structure from the captured data.
- *Data management:* The incorporation of small-scale data has the potential to massively grow the size of the spatial datasets that are being used. While level-of-detail techniques already exist for visualization, it is not clear that such methods will scale easily to datasets such as these. The size of the datasets probably makes local storage infeasible, meaning that distributed cloud-like storage and access paradigms will need to be used.
- *Simulation-ready models:* With large data sets that include significant detail over a large region, there are many possibilities for complex environmental simulations. This can include not only traditional tasks such as navigation and visualization, but also more novel techniques such as sound propagation, heat propagation, airflow, etc. Significant work will be needed to both make such spatial datasets usable for such applications, but on the application ends to realize the potential that such a dataset could provide.
- *Statistical modeling of environments:* Actually capturing full spatial data in various environments is not always reasonable, due to the way the environment can change, and the infeasibility of collecting a full data set. As a result, statistical descriptions of spatial data might need to replace precise descriptions of locations. Research will be needed into understanding the best ways to describe this data, to use the data in calculations, and to present to people the uncertainties and range of outcomes that such a model would create.
- *Predictive modeling of changing environments:* While much of the work in spatial data has assumed that locations remain basically static, this is not likely to be the case, especially at smaller scales. People, animals, nature, etc. can all have an effect on the environment, particularly at smaller scales. A person choosing to move one object in an environment could set off a chain reaction causing more significant changes elsewhere (as an extreme example, imagine moving a support for a building). Moving from a static description of a region to a dynamic representation that responds in physically-accurate ways to the interactions is a challenging problem, but one that needs to be solved for spatial data systems to become more dynamic.

Photogrammetry: A Foundational Technology for Geospatial Analysis

Edward M. Mikhail
Geomatics Engineering, Purdue University

Introduction

The root, “photo”, in photogrammetry may mislead people to restrict it to photography; it actually derives from a Greek word meaning “light”. Therefore modern photogrammetry deals with all manners of sensing within the electromagnetic spectrum. It complements Remote Sensing in that it concentrates on various aspects of precise positioning, mensuration, and extraction of metric information from sensed data for geospatial applications and analysis.

At the heart of its activity is the complete modeling of the functional and stochastic elements involved in relating the data resulting from sensing to the sensed objects. This complete model is then invoked during the myriad of tasks which yield reliable information and its equally reliable quality measures. In the world of geospatial analysis, photogrammetry has been recognized as emphasizing the rigorous propagation of quality (error covariances) from input data through the sensing models to the output information/products. Although photogrammetry has been extensively used during the past half-century with great success, there remain many challenges to be met as it continues to diversify its use as the backbone of geospatial analyses. Many of these will be enumerated in the following three sections.

Research Opportunities

Universal Modeling of Sensing:

Sensor models are fundamental to photogrammetric activities. Sensors fall into two broad categories, active sensors such as radar, sonar, laser scanners, etc., and passive sensors such as frame cameras (film and array based), push-broom linear array scanners, whiskbroom scanning systems, motion imagery, etc., within various ranges of the electromagnetic spectrum. In the past, each individual sensor was modeled uniquely, which resulted into inefficient stove-pipe approach to implementation. Recently, there has begun an effort toward seeking a single model for at least each single class of sensors. The most notable success has been with synthetic aperture radar, or SAR, which is currently being implemented and tested with several European systems (German and Italian). Following the SAR example, efforts are currently being expended on generalizing modeling of a significant subset of passive sensors. Once this is accomplished, the next opportunity is to develop a general sensor model for all passive sensors, which would be a challenging task. A more challenging task, perhaps, is a model that is general enough to accommodate data from both active, such as Lidar, and passive such as line scan imagery, or motion imagery. Significant savings would be accrued from such a generalization down stream in the data exploitation phase.

Fusion and Beyond

Fundamental to fusion is registration which in turn depends upon accurate correspondence. It is well known the more similar the entities, the more accurately determined the corresponding features. As the severity of the geometric conditions of sensing increases, such as highly convergent imagery; as the terrain character becomes more rugged, and in addition as the inputs become more diverse, such as of different modalities, correspondence becomes more difficult. Consequently registration, which is the basis for fusion becomes less accurate. The result that, if for example, change detection is one application of fusion, the apparent change may actually be due to the errors resulting from lack of proper registration. So, robust, reliable, and accurate means of determining correspondence under these stated conditions are critically needed. Related to this topic, the following are important research areas for investigation:

- Fast, accurate, and reliable registration of multiple sources, recognizing the highly non-linear nature of the underlying transformations involved.
- Absolute registration of a time series of video frames without the availability of a priori sensor attitude information or ground control information.
- Incorporation of image semantics into image registration and modeling
- More meaningful generalized fusion, not just pan-sharpening and band layering.

Open problems

Automated Feature Extraction

This has probably been one of the most difficult and persistent problems. In the words of one colleague, "we have always been promised that it will be completely accomplished 'within the next 10 years'.. and we are still waiting". Many automated techniques can achieve ~ 80% performance autonomously or with minimal supervision. However, reliably getting that final 20% has proven to be elusive and incredibly difficult. To solve this, approaches must be developed that, by necessity, deal with different types of scenes and content, sensing conditions, etc.

Quality Assessment Measures, and Propagation

When geospatial information is placed in the hands of the user, its power is greatly enhanced if reliable and standardized measures of its quality are provided. Several different terms are frequently used to define the quality of spatial data and information such as: accuracy, precision, error estimates, confidence, uncertainty, and provenance. Confusion can often result from such variety of terminology, particularly when one tries to implement quality standards that are intended for the promotion of interoperability.

A 2010 National Research Council (NRC) report written for the NGA (National Geospatial-Intelligence Agency) on "New Research Directions for NGA" recognized data uncertainly as a long-term issue that cuts across all NGA core technology areas. It must be clear that the proper course of action taken on the basis of geospatial information can be materially affected/altered because of improper handling of error propagation, and conveyance and visualization of the information quality. Whereas "uncertainty" is well understood and properly treated in the basic photogrammetric operations, there remains significant research work to deal with the variety of tasks, such as registration, fusion, and beyond, including aggregation, integration, and conflation of geospatial data across time and space to facilitate spatial analysis. Topics for investigation include:

- Proper and practical statistical modeling of correlated errors. Correlation is typically temporal in the sensor support data for "same pass" images or a video sequence of frames, or it can be spatial correlation between 3D ground point coordinates.
- Replace the current very limited quality reporting by circular and linear errors, or CE and LE. Devise effective means of conveying the original 3X3 covariance matrices which result in error ellipsoids.
- Efficient characterization of very large/dense lidar point data for use in "down-stream" processing, and similarly for 3D voxel representation.

Grand Challenges

Geospatial Database

We presently have continuous image streams, (satellites, aerial, terrestrial) and related information from many sources. What is needed now is an efficient automatic means to use these data to maintain a geospatial database. The availability of large control database with full covariance between all 3D point pairs provide information for an unlimited number of photogrammetric applications. In addition to the database, the following is also required: (1) A powerful method for its optimal/rigorous sequential generation and update; (2) An efficient means for dissemination to users; and (3) Techniques to represent effectively the various cross-covariance information for ground point pairs in a very large control base covariance matrix.

What if No-GPS?

GPS is currently indispensable to all photogrammetric/geospatial activities. As a contingency, should it become unavailable, there is a need to consider alternatives. The power of photogrammetry and high quality extensive, ground control database can be exploited for the purpose.

Global Orientation System (GOS)

An outstanding development would be to establish a GOS with one arc second real time accuracy about each of the three axes. Given that, combined with GPS, real time accurate geospatial applications would become a reality.

Acknowledgements

Many of these ideas resulted from discussions and exchanges with the following colleagues and friends: Todd Johanesen, and Tom Ager, NGA; Chris McGlone, SAIC; Gregg Kunkel, BAESystems; John Dolloff, Pete Doucette, Tim Nagy, and Hank Theiss, Integity Apps, Inc.

Augmented Reality Everywhere: the Last ~~Kilometer~~ ~~Centimeter~~ Pixel

Greg Welch*

The University of Central Florida
Institute for Simulation & Training
Computer Science Division, EECS
Orlando, FL 32826-3281

The University of North Carolina at Chapel Hill
Department of Computer Science
Chapel Hill, NC 27599-3175

Abstract

Thanks to today's relatively high-speed wireless data networks, and location/orientation-aware mobile phones, we are tantalizingly close to realizing a common dream of Augmented Reality (AR) everywhere. But even with appropriate content and concepts of operation, things are not quite right: icons and annotations are mis-registered and/or jiggle around. We can get the data to the phone (the last kilometer) and perhaps localize the phone relatively accurately (the last centimeter) but we can't seem to realize stable and accurate registered imagery (the last pixel). We need to marshall our significant networking, modeling, computing, and other resources to "close the loop" on AR registration.

1 The Reality of Augmented Reality

Augmented Reality (AR), once offering the promise of Super Man's X-Ray vision, wants to be all the rage with today's camera-equipped smart phones and tablet computers. Even assuming ubiquitous wireless data, and reasonably accurate position and orientation estimation, we are still not seeing widespread adoption of AR on camera-equipped mobile phones. Why is this?

One problem is that of content and use concepts. We (society) still haven't quite figured out what we should see/show with our AR-equipped phones. Do we need to see graphical navigation indications and annotations superimposed on live imagery of the real world around us? Perhaps. Do we want to see restaurant reviews visually registered on the wall of the restaurant across the street? Perhaps. Do we want to have to hold our phone up to an historic landmark to read about it? Perhaps for information that has spatial relevance (e.g., annotations associated with parts of the landmark), but otherwise perhaps not. It would seem that AR is a wanna-be commodity paradigm that we still have not quite figured out. Apps on today's mobile phones are more a novelty than useful tools.

A second problem is that there remain non-trivial technical challenges in getting it right. We are almost there but not quite. Despite remarkable progress in our ability to get data to mobile phones (the last kilometer) and to estimate the position and orientation (pose) of a phone relatively accurately (the last centimeter), we can't seem to realize stable and accurate registered imagery (the last pixel). Augmented imagery jiggles around as my hand (inevitably) shakes, sloshing back and forth as I pan around, and/or is completely mis-registered. (My wife, as pretty as she may be, is not a Vanda orchid.) Why is it that with ever-increasing accuracy and resolution, and more sophisticated localization and orientation algorithms, we can't make imagery "lock on" to the appropriate objects in the real world around us?

*e-mail: welch@ucf.edu, welch@unc.edu

2 Open Loop Augmented Reality

Part of the problem is that most AR systems today have no understanding of the "real world around us." They don't recognize the restaurant, historic landmark, nor orchid. Why should they? We all have a tendency to assume that new and better hardware will solve the problem. We see things like the Kinect changing the world of natural user interfaces (and anything else researchers can think of) and assume that the next generation of GNSS and other technologies will provide the necessary accuracy (popularly confused with *resolution*) to fix our mobile phone AR registration problems. But we need to wake up. If we try and wait for accuracy and stability in mobile phone pose estimates to become sufficient to make augmented imagery appear "locked on" to real world objects around us, we will be waiting a long, long time—perhaps forever.

The problem is that today's "open loop" AR paradigm is almost certainly doomed to never succeed like we dream it will. Component accuracies and resolutions, delays, and dynamic variations in the various components and parameters conspire against us. This is compounded by errors in the models for the objects/scenes we are trying to augment. As impressive as such models are, they are never going to be perfect in all respects. These inevitable errors and perturbations are magnified by distance and other factors, and manifest themselves as mis-registered and/or unstable imagery. Even in a controlled environment such as a laboratory it is exceedingly difficult to realize acceptably accurate and stable augmented reality imagery. The typical approach is to place specially designed visual markers *near* the object of interest (e.g., next to the little AR man we want to render on the table), or to make the markers the object of interest (render the AR man *on* the marker). But we can't put markers all around us as we move about during our daily lives. Nor can we expect that our municipalities will adorn our cities with AR markers. (Would we want a plethora of markers around our historic landmarks anyway?)

3 Closing the Loop

Aside from content and use concepts, we need to give up on the mistaken notion that the "natural" evolution of localization and orientation technology will somehow result in acceptable registration of augmented imagery with objects/scenes. Instead we need to recognize that for AR "everywhere," e.g., visual annotations outdoors with our mobile phones or magic eyeglasses, the primary goal should be **visual registration**, as opposed to accurate pose estimation. We know our models of the world will be imperfect—do we really care if the estimated pose is imperfect, *if* in exchange the AR annotations are locked onto the object/scene of interest? To be clear, I'm not arguing for giving up on *seeking* accurate pose estimation—of course we want that. I'm arguing for giving up on *relying* on it as the silver bullet that will solve our AR registration problems.

So how do we do this? Rather than putting visual tracking markers everywhere *near* the objects/places of interest, we need to *make* the places of interest our markers. That is, rather than *assuming* our open-loop rendering will be registered, we need to *ensure* it is registered via final-stage image (or audio or ...) processing that seeks to minimize error between (a) a *simulated* image of the AR annotations overlaid on the object/scene models, and (b) a *real* image of the AR annotations as they appear on the actual imagery of the object/scene. For example, a desired (and simulated) augmented image might include an overlaid graphical line that correctly appears collinear with the edge of a building in the modeled scene; while in the *real* final augmented image, pose estimation errors can cause the same overlaid graphical line to *intersect* the edge of the building in the real image. We humans would know that the non-collinear lines and corresponding intersection point indicates an error, but today's typical AR system would not. It should. It should attempt adjust the pose (or other uncertain parameters) to minimize the differences between the simulated final augmented image and the real augmented image. In other words, it is the combination (e.g., optical masking or superposition) of real and virtual imagery that should be the "signal" that is optimized, not the real imagery (or audio, ...) alone.

This idea is not unlike that of *map matching* in navigation—if my vehicle continues to go straight in spite of the navigation system's belief that the road I am on is curving, most navigation systems will recognize that a nearby straight road (that matches my straight trajectory) is probably where I am, and will change its estimate of my location and heading to match. (If I was going straight on a curved road, an erroneous "correction" to my navigation path is probably the least of my worries!)

Most navigation systems "carry" their maps with them in entirety. For humans moving about arbitrarily in the world it is unreasonable to expect that we will be able to carry models of all of the objects/scenes around us at all times. Instead we need to leverage the ever-expanding reach of high-speed wireless data and cloud computation to deliver appropriate object/scene model data "just in time" as we're moving about.

This idea depends on extensive modeling of our world, dynamic objects, etc. But that is already happening—the models are continually evolving—improving and/or expanding via controlled data collections (e.g., large mapping companies) and crowd sourcing. We need to foster these efforts, and make sure the data is available everywhere as we move around. In fact we (as AR users/consumers) can help with the modeling as we use our AR tools. We are, after all, capturing images and comparing them to expectations based on a combination of AR elements rendered according to them (the models) and the effect as seen in the real augmented imagery. As misregistration is minimized, the residuals can be attributed to both pose and model errors in a form of crowd-sourced continuous automatic calibration.

4 Dynamic Cooperative Feedback

Certainly geometric and photometric models of physical objects and scenes (buildings, roads, etc.) should provide one source of data for closed-loop registration. But why stop there? There is other "ambient" information to be had all around us, and in particular the "us" can provide information for others.

For example, moving targets (people, vehicles, robots...) that are simultaneously attempting to estimate their own pose and enforce their own registration, can provide information to other moving targets nearby. That information could include for example a visual (or aural or...) model for the moving target, its estimated pose, and associated uncertainties. That would allow one target to use another target as a reference for pose estimation and closed-loop registration. The shared information could also include the raw measurements from the sensors of another nearby moving target—the latter target providing remote sensing capability for the former. Information from moving targets cooperating in this way, combined with information about static objects/scenes, can offer increased leverage in each system's attempt to estimate pose and drive the registration error to zero. During the ongoing process of pose estimation and registration minimization, targets would of course feed back the parameters they had to adjust (to minimize the registration error) back into the cloud, as part of the continuous cloud-sourced automatic calibration of the world models and state.

5 Conclusions

For cases where our goal is to "attach" overlaid information to objects/scenes, we cannot rely on magical improvements in pose estimation technology in mobile phones to be the silver bullet that will solve our registration woes. We need to re-define the problem as one of minimizing the registration errors. This does not mean giving up on pose estimation (it is necessary, of course)—in fact the optimization that seeks to minimize registration errors can itself act as a source of feedback to adjust parameters/models in a continuous automatic way, as we use our devices for AR. Now if only we could figure out *what* we want to register and *why*....

SPATIAL COMPUTING SCIENCES

Human Interaction in Space: Proximal, Virtual, Distributed

A Position Paper for the NSF/CCC Visioning Workshop on Spatial Computing

Thomas Erickson
IBM T. J. Watson Research Center
snowfall@acm.org

August 2012

INTRODUCTION

I am interested in how humans communicate and collaborate. For the most part, as I will discuss shortly, this is deeply entwined with environment and space.

My principal interest *vis a vis* spatial computing is how the human communicative and collaborative practices that have developed in the context of proximal (i.e., face to face) interaction are extended to supporting interactions between spatially distributed (i.e., non proximal) groups. This involves both studying the new forms of interaction that arise in spatially distributed groups, as well as understanding how to better design systems in light of these new forms of spatially distributed interaction.

In this position paper, I will suggest three opportunities for research, and call out one additional area for discussion. First, I will provide some background on human interaction.

BACKGROUND: HUMAN SPATIAL INTERACTION

Human interaction takes place in space. People face one another when they speak, and communicate not simply with words but with expressions, gestures and bodily postures and positions. Human interaction as carried out by dyads and among small groups has long been a topic of research in social science; many researchers have devoted considerable attention to the ways in which people use space (e.g., [6, 7, 9]) to structure their communication.

It is important to recognize that such interaction embraces far more than conversation with known others. People navigate crowded sidewalks, cross streets, form queues, and exhibit other forms of remarkably orderly interaction with strangers as they inhabit urban spaces (e.g., [8, 14]). Architects and urban designers (e.g., [1, 5, 8, 14]) draw on such observations to offer design guidelines for buildings, public spaces, and cities.

What ties all of this interaction together is that it takes place in close proximity—what is generally referred to as face to face interaction. In such situations, human interaction is governed by a combination of *cues* embedded in the environment, *social norms* particular to the locality, and interactants' *mutual observations* of one another. An example would be a crowd at a street corner waiting to cross the street: a crossing signal and cross walk (environmental cues) designate the timing and spatial locus of crossing; social norms having to do with obeying lights govern the crowd's behavior (such norms being stronger in some places than others); and mutual observation of others' behavior reinforces or undermines the norms (as when one person decides to cross against the light and the rest of the crowd follows).

Having laid out this perspective on the spatial dimensions of human interaction, I will discuss three areas that seem to offer prospects for important research.

NEW TOOLS TO STUDY PROXIMAL INTERACTION

First, and most briefly, the emergence and growing ubiquity of technologies for tracking the spatial location of people and gathering other sorts of biometrics offers the prospect of new methods and deeper understandings of how traditional proximal interactions play out. I will say no more about this, as Pentland and his colleagues (e.g., [11]) have been exploring this direction, and it is not a principal focus of my research.

LARGE-SCALE INTERACTION IN VIRTUAL SPACES

Second, I am interested in how proximal interactions are carried out in virtual environments. While a recent wave of interest in such environments appears to have subsided, I am convinced that we have not seen the last of them. I believe that the combination of the pull of increasing bandwidth and processing power, with the push of increased restrictions on our ability to expend time and resources in travel, will result in spatialized virtual environments becoming a common venue for interaction.

While I was initially skeptical about the prospects of virtual environments for supporting graceful human interaction, a few years ago I had an experience that changed my mind. IBM held a 500 person conference in a customized version of Second Life, and I took the opportunity to do a field study of the event [4]. While there was no shortage of problems—both straightforward and subtle—some aspects of the conference worked remarkably well. In particular, I was struck by the success of the poster sessions and the ways in which they supported graceful interactions among the participants. In some cases the interactions mirrored those seen in face to face poster sessions (e.g., self-organization of small groups; social navigation; opportunistic engagements), and in other cases they surpassed what was possible in their face to face analogs (e.g., the ability to search and navigate the space).

I came away from the study with several of conjectures.

- 3D avatar-based virtual environments have real prospects of becoming a useful mode of interaction among large numbers of people.
- Such environments are not—and are not likely to develop into—faithful analogs of physical space. The arbitrary nature of their 'physics' destabilizes the interactions that occur in them (e.g., how far voices carry; teleportation), and we need to understand how people—particularly those who have come of age playing massively multi-player online games—adjust their interactions accordingly. (There is, of course, already a rich vein of research in MMPOGs to draw upon.)
- We do not, as a field, understand how to design such environments to enable 'mundane' interactions. The needs of work-a-day communication do not always fit well with the ludic roots of these environments. A combination of usage studies and design research is needed to remedy this.

LARGE-SCALE INTERACTION IN REAL SPACE

For most of their history humans have interacted in face to face situations. It is only with the emergence of modern technologies like the printing press, telephony and now the internet that people have been able to interact across distance. Most recently, the growing ubiquity of mobile devices equipped with locative sensors is enabling distributed groups of people to communicate and coordinate in real time from wherever they happen to be. From flash mobs to crowdsourcing, we need to understand how humans are successfully interacting in distributed spatial contexts.

One domain that illustrates these issues is what I call geocentric crowdsourcing. Crowdsourcing is the use of the perceptual, cognitive and enactive abilities of a distributed group to achieve some purposeful end (e.g., Wikipedia; Mechanical Turk). Geocentric crowdsourcing is when the locations of the members of the crowd matter, as they do when crowdsourcing is applied in domains like smart cities or citizen science.

Two examples illustrate the concept. FixMyStreet [3] is an application that allows urban inhabitants to report potholes and other street-related problems on a publicly visible map, which are then brought to the attention of the appropriate governing body. As individuals' reports appear on the shared map, it creates a powerful aggregate representation of the state of the streets—areas with lots of problems become quite apparent. Another example is Cyclopath, a user-editable street map intended to help bicyclists find bicycle-friendly routes around the city (e.g., [2, 10, 13]. Cyclopath relies on the cycling community to add data – road surface conditions, off-road paths, location of coffee shops – that is useful in determining a good bike route, but not found on conventional maps. Both of these systems have been quite successful, and rely on both the local knowledge and local motivation of their users.

However, these sorts of systems are in their infancy, and they would be more useful if the behavior of their crowds could be orchestrated. For instance, it would be useful to be able to assess how well-covered a particular region was, either by repeated sampling by users, or via integration with other digitally-sourced data. Similarly, if a particular area needs work, it would be useful to understand how to focus the work of the crowd on a particular area (e.g., [12]). These, and similar ends, could be achieved by a variety of means ranging from enabling sub-groups to purposely collaborate to providing global incentive mechanisms to shape crowd behavior to particular ends.

Crowd orchestration is just one example of the issues that large scale interaction raises. As cities get smarter and inhabitants are increasingly connected, new possibilities arise for mining and shaping mass behavior. Might we be able to detect concerns about the level of public safety of a neighborhood by tracking changes in pedestrian behavior over time? Might the distribution and dynamics of a taxi fleet serve as a distributed sensor of economic growth or contraction in particular areas of a city? In the event of traffic congestion, how might a flow of commuters be re-routed over multiple routes to ameliorate the problem (even while assuming that not all will simply follow instructions issued via their smart devices)?

A CLOSING QUESTION: WHAT ABOUT ROBOTICS?

In wrapping this up, I'll raise an issue that I hope the workshop will address. While I have no particular expertise in the area, I was surprised that there was no mention of robotics. Surely autonomous devices that can move and act in spatial environments are relevant to spatial computing. And with the growing popularization of robotic toys, the consumer acceptance of robotic utility devices, and the advent of drones for domestic and foreign

REFERENCES

1. Alexander, C., Ishikawa, S., Silverstein, M., Jacobson, M., Fiksdahl-King and Angel, S. *A Pattern Language*. Oxford University Press, 1977.
2. Cyclopath. http://cyclopath.org
3. FixMyStreet. http://Fixmystreet.org.uk

4. Erickson, T., Shami, N. S., Kellogg, W.A. and Levine, D.W. Synchronous interaction among hundreds: An evaluation of a conference in an avatar-based virtual environment. *Proc. CHI 2011*. ACM Press, 2011.

5. Gehl, J. *Life between buildings: Using public space*. Van Nostrand Reinhold, 1980.

6. Goffman, E. *Behavior in public places: Notes on the social organization of gatherings*. Doubleday Macmillan, 1963.

7. Goodwin, C. *Conversational organization: Interaction between speakers and hearers.* Academic Press, 1981.

8. Jacobs, J. *The death and life of great American cities.* Random House, New York. 1961.

9. Kendon, A. *Conducting interaction: Patterns of behavior in focused encounters*. Cambridge University Press, 1990.

10. Panciera, K., Preidhorsky, R., Erickson, T., and Terveen, L. Lurking? Cyclopaths? A Quantitative Lifecycle Analysis of User Behavior in a Geowiki. *Proc. CHI 2010*. ACM Press, 2010, 1917-1926.

11. Pentland, A. *Honest signals: How they shape our world.* MIT Press, 2008.

12. Priedhorsky, R., Masli, M., and Terveen, L. Eliciting and focusing geographic volunteer work. *Proc. CSCW, 2010.* ACM Press, 2010.

13. Priedhorsky, R. and Terveen, L. The computational geowiki: What, why, and how. *Proc. CSCW 2008*. ACM Press (2008), 267–276.

14. Whyte, W.H. *City: Return to the center*. Anchor Books, New York.1988.

SPATIAL SIMILARITY

Mike Goodchild, UC Santa Barbara

Measurement of similarity is a time-honored practice in science, forming the basis of all taxonomies. Placing things into categories is one of the most basic and elementary scientific functions, based on minimizing differences between things in the same category and maximizing the differences between things in different categories. Nevertheless the measurement of similarity between locations is remarkably under-exploited, in part because we normally expect locations to be grouped contiguously into regions based on a combination of proximity and similarity, rather than similarity alone. Groupings of locations based only on similarity are likely to be fragmented.

It seems appropriate at this point in the development of GIScience to consider the question of similarity anew. It is easy to identify use cases for geographic *analogs*, locations which are similar on one or more dimensions to a location of interest; proposed here is a research program to support search for analogous locations based on similarity, with the user exercising control over how similarity is measured. Analogs are commonly used in marketing, where an assessment of a potential site for a retail outlet is made by searching for existing outlets with similar marketing-oriented attributes, on the argument that performance of the potential site will be similar. Researchers might use such a service of analog search in the early stages of a project, when it might be important to identify similar locations. Analogs might be useful in marketing to tourists ("This location is like…") or in location scouting for movies.

Research on spatial analogs would be timely, given the very rapid growth in the volume of available geospatial data that could be used for defining similarity. The fact that so much of these data are online is another advantage, since it would allow search to be conducted rapidly.

Geographers have long distinguished between two somewhat orthogonal perspectives, those of space and place. A spatial perspective is characterized by coordinates, distances, geometry, and the functions commonly found in GIS. A platial perspective on the other hand is dominated by placenames that may or may not be well defined, static, and spatially bounded. The planimetrically controlled maps of a spatial perspective become the schematic maps of a platial perspective. Interest in place and in the space/place duality has grown rapidly in recent years because of the rise of consumer-oriented geospatial services and practices, including volunteered geographic information (VGI), point-of-interest databases, and online wayfinding aids. A research program on spatial similarity and analogs should include both spatial and platial perspectives: search for similar locations, and search for similar named places.

The research will need to address several key issues:

- How to identify the set of dimensions on which similarity will be measured, how to present the set to the user, and how to allow the user to build similarity metrics.
- How best to organize the vast amount of relevant information to support similarity measurement and search for analogs. How should scale be addressed, how should

data be organized (by layer, by feature type, etc.), and how should it be indexed for rapid search?

- How is the problem of search for analogs distinct from the problem of similarity measurement, and what aspects of the latter literature are relevant?
- What is the full set of use cases? How are general trends in spatial computing likely to affect the problem in the next few years?
- What previous work has been conducted on the problem? One PhD dissertation at Penn State addressed some aspects of it in 2008 (Banchuen, 2008).

References

Banchuen, T., 2008. *The Geographical Analog Engine: Hybrid Numeric and Semantic Similarity Measures for US Cities*. PhD Dissertation, Pennsylvania State University.

Spatial Cognition

Stephen C. Hirtle
School of Information Sciences
University of Pittsburgh
Pittsburgh, PA 15260

Background

Geographic Information Science (GIScience) is a field that emerged roughly 20 years ago (Goodchild, 1990) with the inclusion a cognitive principles as a fundamental component of the science (Mark, et al.1999). This philosophy came about from earlier precursors, including the specialist meetings of the National Center for Geographic Information and Analysis (NCGIA) and the Conference on Spatial Information Theory (COSIT), both of which explicitly brought together geographers, computer scientists, cognitive psychologists, among others, to examine the nature of spatial information.

The inclusion of spatial cognition as part of the larger GIScience movement has resulted in two distinct benefits. First, there has been a large scale effort to identify and to quantify the cognitive elements of spatial knowledge and spatial communication. That is, the notions of regions, landmarks, hierarchical reasoning, geographical scale, the interplay of visual, verbal and spatial knowledge into a cognitive collage, and other cognitive elements, have been identified as crucial to human understanding of space (Hirtle, 2011; Kuipers, 2000; Montello, 2009; Tversky, 1993). Second, modern spatial tools have been built upon the knowledge that we have gained in studying the cognition of spatial concepts. This include the use of multiple modalities and selective information presentation that are found in at the heart of current navigation systems and route-finding algorithms, as well as crowd-sourcing solutions to spatial problems and the spatialization of large datasets (Agrawala, Li, & Berthouzoz, 2011; Goodchild, 2007; Jones, 2007; Skupin & Fabrikant, 2008)

Open Problems and Research Opportunities

Given the advances of the past two decades of research, in addition to the current state of technology, one can identify several open problems and research opportunities that have emerged. These are listed below in no particular order and together make up a grand challenge of providing real-time, spatially relevant information on integrated platforms for spatial planning and decision making.

Acquisition of geographical knowledge. Recent studies have shown that people not only use GPS navigation systems, but also become dependent on them for navigating repeated times to the same location (Parush, Ahuvia, & Erev, 2007). This raises an interesting question about how technology can present accurate geographic information to a user, but at the same time support the acquisition of geographic knowledge. Ideally, the repeated use of the technology would increase geographic awareness, rather than leading to impoverish knowledge of the surrounding environment.

Communication of geographical information. Spatial communication involves the matching the description of the environment with the physical environment (Hirtle, Richter, Srinivas, & Firth, 2010). In human-to-human communication this might involve landmarks, road objects and topography, such "Turn left at the stop sign, just past the McDonalds at the top of the hill." The ability to automatically extract landmarks, visible objects, and difficult navigational maneuvers is an open problem for the development of user-friendly navigation systems.

Attributes of space. Related to the previous challenge is the identification of salient objects in the environment. Unique or useful objects for identifying spatial locations vary from region to region and depend on both cultural norms and the variation within the environment (Klippel, Hansen, Richter, & Winter, 2008). Ongoing work on geographic ontologies could provide a theoretical framework, but additional research is needed on how to automatically extract salient objects and how to best use those objects in the development of navigation systems.

Cognition of dynamic phenomena. Environments are not static and particularly in mission-critical applications, such firefighting and public safety, there is a need to comprehend, represent, and model dynamic phenomena in real-time (Hornsby & Yuan, 2008). Not only is the modeling of dynamic geographic phenomena critical, but the communication of the parameters is also of great importance.

Crowd-sourcing and VGI. The ability to use crowd-sourcing and other forms of volunteered geographic information (VGI) will lead to a new generation of spatial tools (Goodchild, 2007) For example, while it is theoretically possible to identify potentially safe and efficient bicycle routes in the United States from road network data, a more profitable approach might be to automatically track routes taken by bicycle riders over a period of time. This approach would generate the preferred paths, regardless of the underlying database constraints on the space (Panciera, Priedhorsky, Erickson, & Terveen, 2010). The explosion of crowd-sourced data along with the growing number of explicit VGI projects will lead to vast new data sources, many with an implicit cognitive bias, which can be data-mined for new and useful information.

Together, these five specific topics are examples of the kind projects that would link the spatial cognition of the user with technological tools. In each case, additional work would be needed to understand (1) the explicit cognitive constraints on the user, (2) the required computational models to incorporate these constraints, and (3) the user interface issues to implement a cognitively-aware geographical information tool.

References

Agrawala, M., Li, W., & Berthouzoz, F. (2011). Design principles for visual communication. *Communications of the ACM, 54*(4), 60-69.

Goodchild, M. F. (1990). Spatial information science. In *Proceedings of the Fourth International Symposium on Spatial Data Handling* (pp. 3-12). Zurich, Switzerland.

Goodchild, M. F. (2007). Citizens as sensors: the world of volunteered geography. *GeoJournal, 69*(4), 211-221.

Hirtle, S. C. (2011). *Geographical design: Spatial cognition and geographical information science*. San Rafael, CA: Morgan & Claypool.

Hirtle, S. C., Richter, K.-F., Srinivas, S., & Firth, R. (2010). This is the tricky part: When directions become difficult. *Journal of Spatial Information Science, under review.*

Hornsby, K. S., & Yuan, M. (Eds.). (2008). *Understanding Dynamics of Geographic Domains*. Boca Raton, FL: CRC Press.

Jones, M. T. (2007). Google's geospatial organizing principle. *IEEE Xplore, 27*(4), 8-13.

Klippel, A., Hansen, S., Richter, K.-F., & Winter, S. (2008). Urban Granularities - A Data Structure for Cognitively Ergonomic Route Directions. *GeoInformatica, 13*, 223-247.

Kuipers, B. (2000). The Spatial Semantic Hierarchy. *Artificial Intelligence, 119*, 191-233.

Mark, D. M., Freksa, C., Hirtle, S. C., Lloyd, R., & Tversky, B. (1999). Cognitive models of geographical space. *International Journal Of Geographical Information Science, 13*(8), 747.

Montello, D. R. (2009). Cognitive Research in GIScience: Recent Achievements and Future Prospects. *Geography Compass, 3*(5), 1824-1840.

Panciera, K., Priedhorsky, R., Erickson, T., & Terveen, L. (2010). Lurking? Cyclopaths? A Quantitative Lifecycle Analysis of User Behavior in a Geowiki. *Proceedings of CHI 2010*, 1917-1926.

Parush, A., Ahuvia, S., & Erev, I. (2007). Degradation in Spatial Knowledge Acquisition When Using Automatic Navigation Systems. In S. Winter, M. Duckham, L. Kulik & B. Kuipers (Eds.), *COSIT 2007, Melbourne, Australia* (Vol. LNCS, vol. 4736, pp. 238-254). Heidelberg: Springer

Skupin, A., & Fabrikant, S. (2008). Spatialization. In J. Wilson & S. Fotheringham (Eds.), *The Handbook of Geographical Information Science* (pp. 61-79.): Blackwell Publishing.

Tversky, B. (1993). Cognitive maps, cognitive collages, and spatial mental models. In A. U. Frank & I. Campari (Eds.), *Spatial information theory: Theoretical basis for GIS*. Berlin: Springer-Verlag.

Spatial Cognition for Robots

Benjamin Kuipers*

3 August 2012

Spatial Knowledge is Everywhere

Spatial knowledge is fundamental to almost all of human knowledge. This is easily seen by considering how ubiquitous spatial metaphors are in human thought and communication [8]. (Note that the meanings of the words "fundamental" and "ubiquitous" in the previous sentences both involve spatial metaphors!)

Since spatial knowledge is so central to human knowledge representation and problem-solving, it will be important to provide similar capabilities for spatial knowledge representation and inference for intelligent computational systems. A robot is a special case of an intelligent computational system, in that it is situated in, and interacts with, the physical world, so spatial knowledge is particularly important for robots.

Representations for spatial knowledge that are quite different from human representations, such as GPS coordinates, can certainly be useful computational tools. However, to solve problems as formulated by humans, and especially to communicate effectively with humans, human-like knowledge representations are essential.

Scales of Spatial Knowledge

The knowledge an agent (human or robot) may have about the spatial structure of a situation depends on the relation between the agent's sensory capabilities and the scale of that spatial structure. This leads to a variety of quite different representations, with quite different learning and problem-solving capabilities.

At the bottom, spatial knowledge starts with **2D sensor structure**. The receptors in the human retina, or the pixels in a camera, have a physical 2D structure. Likewise, touch sensors on the skin have a locally 2D structure. The camera's pixel structure has a simple specification and is engineered in. The biological structure of receptive fields or touch sensors is partly innate and partly learned or calibrated from experience. Connections among sensory modalities such as sight, sound, and touch must also be learned, and can demonstrably be re-learned as circumstances change. Useful models of these adaptive capabilities will be important for long-lived robots and other computational systems [11].

3D object shape and appearance models help factor the highly variable sensory images of an object perceived over time into a relatively constant shape model and relatively simple time-varying pose and configuration variables. (In the early months of life, babies spend a lot of time observing their own hands as they move them and change their shapes. Surely they are learning something very important about space, as well as about hands.) Computer vision researchers have developed multiple approaches to representing 3D object shape, ranging from compositions of 3D primitives such as generalized cylinders or *geons* [2], to the *view-sphere*, a 2D manifold of 2D images as the object is perceived from any surrounding pose [3].

Small-scale space is the nearby space within the sensory horizon of the agent, described in terms of relations among objects, and between objects and the observer. Relations among objects can include near/far,

*Computer Science & Engineering, University of Michigan, Ann Arbor, Michigan, USA. kuipers@umich.edu.

right/left, forward/back, above/below, and so on. Solid objects participate in simple relations such as support and touching, and make up more complex assemblies such as a clock [4]. Methods such as the Region Connection Calculus [12] categorize the ways that 2D regions can relate, including overlapping and containment. Configuration space methods such as probabilistic road-maps [5] make it possible to compute trajectories involving closely fitting parts, but it is not clear what relation these methods have to human spatial knowledge.

Large-scale space is knowledge of spatial structure beyond the sensory horizon of the agent, such as the *cognitive map* of a building or a city [9, 7]. Knowledge of large-scale space must be acquired and integrated over time and travel. The Spatial Semantic Hierarchy [6, 1] is an integrated hierarchy of different representations for large-scale space, supporting a flexible set of states of incomplete knowledge.

Abstract spaces are spaces whose structure is not observed directly by the agent, but is inferred from other sources of evidence. These include both very large spaces such as the structure of the solar system, and very small spaces such as the shape of the DNA molecule.

These different spatial scales do not have sharp boundaries. However, knowledge in each scale has enough common spatial structure to be worth considering as a separate type.

Communicating Spatial Knowledge

In addition to having spatial knowledge of its own, an agent must also be able to communicate its spatial knowledge to other agents, and to understand spatial knowledge that is communicated to it. This applies to all three scales of space discussed above. The problem of communicating spatial knowledge can be decomposed according to the medium of communication.

- linguistic description, including route directions and assembly instructions;
- graphical communication, including displays and drawn and printed maps;
- gesture, sound, haptic, and other communication modalities.

Technology certainly exists (e.g., Mapquest and Google Maps) for generating graphical and linguistic route directions, though significant improvements remain possible. The more challenging problem is *understanding* natural human-generated route directions and other spatial instructions [10]. This is partly due to the fact that people shift seamlessly between different representations for spatial relations when describing a situation [13].

Grand Challenges

These challenges require learning, problem-solving, and communication in both directions with various media.

Learning and using a cognitive map of large-scale space. Certain mobile robots (e.g., an intelligent wheelchair or autonomous car) have a job to do, taking a human from one place to another. Can it learn a useful cognitive map from observations obtained while it is doing its job, without the opportunity to explore the environment autonomously? Can it represent the states of incomplete knowledge it will necessarily have? Can it learn parts of the cognitive map from natural human instructions given to it, as verbal directions, as sketch maps, or as joystick guidance? Can it provide useful feedback to a person about inadequate, incorrect, or incomplete instructions?

Building complex structures from jumbled parts in small-scale space. Consider a jumbled collection of parts from which a complex assembly can be created, for example a water pump for a car or a LEGO toy structure. The robot's task is to sort through the collection, learning the appearances, shapes, structures, and properties of the individual parts, and experimenting with how they might fit together. Instructions may be provided as verbal language, as written text, as diagrams, or as a multi-modal combination. Can the robot learn to understand references in the instructions? Can it learn the skills required to create the specified assembly? Can it learn fundamental skills that make it easier to create other assemblies in the future? What does it learn about the properties of objects and their relation to the object's shape?

References

[1] P. Beeson, J. Modayil, and B. Kuipers. Factoring the mapping problem: Mobile robot map-building in the Hybrid Spatial Semantic Hierarchy. *International Journal of Robotics Research*, 29(4):428–459, 2010. doi:10.1177/0278364909100586.

[2] I. Biederman. Recognition-by-components: a theory of human image understanding. *Psychological Review*, 94:115–147, 1987.

[3] H. H. Bülthoff, S. Y. Edelman, and M. J. Tarr. How are three-dimensional objects represented in the brain? *Cerebral Cortex*, 5:247–260, 1995.

[4] K. Forbus, P. Nielsen, and B. Faltings. Qualitative spatial reasoning: the CLOCK project. *Artificial Intelligence*, 51(1–3):417–471, 1991.

[5] L. E. Kavraki, P. Svestka, J.-C. Latombe, and M. H. Overmars. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Trans. Robotics and Automation*, 12(4):566–580, 1996.

[6] B. Kuipers. The Spatial Semantic Hierarchy. *Artificial Intelligence*, 119:191–233, 2000.

[7] B. J. Kuipers. Modeling spatial knowledge. *Cognitive Science*, 2:129–153, 1978.

[8] G. Lakoff and M. Johnson. *Metaphors We Live By*. The University of Chicago Press, Chicago, 1980.

[9] Kevin Lynch. *The Image of the City*. MIT Press, Cambridge, MA, 1960.

[10] M. MacMahon, B. Stankiewicz, and B. Kuipers. Walk the talk: connecting language, knowledge, and action in route instructions. In *National Conference on Artificial Intelligence (AAAI-06)*, Boston, July 2006.

[11] D. M. Pierce and B. J. Kuipers. Map learning with uninterpreted sensors and effectors. *Artificial Intelligence*, 92:169–227, 1997.

[12] D. A. Randell, Z. Cui, and A. G. Cohn. A spatial logic based on regions and connection. In *Proc. 3rd Int. Conf. on Knowledge Representation and Reasoning (KRR)*, pages 165–176. Morgan Kaufmann, 1992.

[13] J. G. Trafton, N. L. Cassimatis, M. D. Bugajska, D. P. Brock, F. E. Mintz, and A. C. Schultz. Enabling effective human-robot interaction using perspective-taking in robots. *IEEE Trans. on Systems, Man and Cybernetics – Part A: Systems and Humans*, 35(4):460–470, 2005.

Representation and Analysis of Spatial Dynamics in the Era of Ubiquitous and Abundant Spatial Information

Focus: Spatial Databases

Position paper by May Yuan, University of Oklahoma

Spatial dynamics has long been a topic of interest to geographers, geoscientists, ecologists, and many other disciplines because a system understanding of the complex interactions of processes in space and time is central to scientific inquiries. There is a wealth of literature in spatial or space-time representation, which is most common categorized as field-based or object-based. However, spatial dynamics cannot be simply modeled in schemes of fields or objects alone due to its needs to address multi-level structures, multi-scalar processes, and the interactions across scales.

Two open questions are fundamental to representing and analyzing spatial dynamics with ubiquitous and abundant geospatial information: (1) analytical methods for data collected without any statistical sampling schemes; (2) representation and analysis for individuals, aggregates, agglomerates, and spatial narratives (spatial stories).

1. How to develop analytical methods for data without statistical sampling schemes? How to remedy potential data biases, under- or over-sampling issues, and other data errors?

Many analytical and computational methods have been developed for spatial dynamics, such as system analysis, Monte Carlo simulation, spatial diffusion models, spatial interaction models, cellular automaton, and agent-based modeling, just to name a few. Generally speaking, these existing analytical and computational methods are subject to specific assumptions of data representation, sampling schemes, or expected relationships among variables. In the era of ubiquitous and abundant geospatial information, ever-growing spatially aware devices and sensors provide ambient spatial data free of any systematic sampling schemes. For example, some spatial data are geocoded from social media, such as site characterization at flickr or web blogs, travel of dollar bills at Where is George site, or personal whereabouts at Four Square, Facebook, or Waze.

Since these data are created socially, there is no assurance for unbiased and representative samples to satisfy data requirements in established spatial statistics. Moreover, the proliferation of spatially aware technologies also provide wealth of individual data with great potential for real-time modeling of population dynamics (the number of occupants in a building at any given time of a day), analysis of human activities at a micro scale (such as entering and existing patterns at a store entrance), or human-oriented environmental monitoring (such as temperature sensors on GPS vehicles for monitoring heat island effects). Most current applications of such data appear common in identifications of clusters, outliners, and network connections in space and time. While exploration of space-time patterns is important, analysis of spatial dynamics is critical to a deeper understanding of what is going on. How do surges of shoppers relate to different types of store sales events? How do different groups of people utilize a building throughout a day or a week? Existing methods can identify space-time patterns and clusters. A logical next step is to develop methods that can reveal how patterns or clusters relate and interact.

2. How to represent, query, and analyze spatial dynamics with individuals, aggregates, agglomerates, and spatial narratives in space and time seamlessly in a database?

Volumes of data are from environmental sensor networks, such as in-situ observation networks, satellite systems, or mobile sensor networks. These sensor networks are designed to monitor environmental events and processes. While massive data have been collected, the events and processes of interest are rarely organized into a database for query and analysis at a level beyond observations. The lack of database representation schemes is one key barrier to develop a database for events and processes. Also lacking are mechanisms to assemble observation data and structure the data in ways to reflect: (1) when and where all necessary components, as observed in the data, are in place and the respective process emerges; and (2) when and where the process develops into different phases. With databases capable of representing events and processes, analytical methods can then be developed to explore when and where processes interact with other features and processes, what environmental settings may be facilitating or constraining the process development, and how human behaviors and activities may influence or be influenced by environmental processes. We need databases that can represent high level space-time abstracts by forming aggregates, agglomerates, and narratives from space-time data.

Aggregates are space-time clusters based on the objects of the same type. For example, traffic jams are clusters of vehicles stopping or moving at a snail pace on roads. Data models should have the ability to represent these identifiable clusters and to reference individuals, when appropriate, to their respective clusters. Another form of aggregates is based on activities or routines. GPS tracks of individuals are composites of daily journeys to work, weekend outings, seasonal vacations, and other activities. Data models for spatial dynamics needs capabilities to represent the periodicity over space and compute patterns of life or meaningful aggregates of activities, events, or processes from disaggregate space-time data. Database representation of spatial dynamics needs to connect individuals to levels of aggregates in space and time and resolve properties which are common to all or are only associated with individuals or distinct levels of aggregates.

An agglomerate consists of many types of functionally related objects. The ability to represent spatial agglomerates is to capture the component objects and functions that hold these spatial objects to operate synergistically. An airport is an example of spatial agglomerate with flight control tower, runways, terminals, and gates. The concept of spatial agglomeration applies commonly in the study of urban growth and industry clusters. Agglomeration economies take advantages of spatial clustering of complementary business sectors that support each other. Socially, spatial agglomeration of a group is the spatio-social network of its members and their functional connections. The Occupy Wall Street movement can also be considered as a spatial agglomerate in which event organizers and participants work together for coordinated activities. Each type of spatial agglomerates has minimal members and necessary organizational structures. Therefore, representation of spatial agglomerates will serve a foundation to identify additional spatial agglomerates, such as automatic recognition of airports in imagery or detection of emergence of spatial social groups or activities in text analytics. In addition to the examples of spatial agglomerates in human environments, natural processes are also spatial agglomerates. A hurricane, for example, consists of an eye, an eye wall, cyclonic winds, and rainbands, each of which has distinct properties.

Spatial narratives emphasize sequences of events in space and time. The history of a place (the city of Norman history), a person's journey (a vacation in the Yellowstone), the development of a process (hurricane), the workflow of a plan (Mardi Gras parade), and the procedure of an operation (boats entering or leaving a port) are all examples of spatial narratives with distinct narrative components and structures. Narrative databases and computational methods for narrative generation and narrative analysis can help elicit preconditions, precursors, evolution of activities, consequences, and spatial histories as a high-level knowledge synthesis for space-time data from multiple sources.

SPATIAL COMPUTING SERVICES

Digital Cityscapes: Challenges and Opportunities

Ming C. Lin and Dinesh Manocha
University of North Carolina at Chapel Hill

Abstract

In this position paper, we examine the algorithmic and computational challenges in real-time modeling and simulation of *digital cityscapes*. We outline the areas of research challenges in *digital cityscapes* that can benefit from more advanced computational techniques and algorithms. We briefly survey some of our recent progress as part of our early attempt in adopting multi-agent plannning and simulation for digital cityscapes and highlight some of the remaining challenges.

1 Introduction

With industrial revolution, recent economic and social development, increasingly more people are leaving rural areas and migrating to cities, thereby leading to rapid urbanization of the world population in the last century. Today, more than 50% of the global population live in urban areas, with the figure projected to rise to 60% by 2030[1]. Given the ubiquitous urban development across all advanced and developing countries, modeling and simulation of cityscapes is clearly emerging as an important topic for city planning and urban development that require interactive visualization to evaluate various alternatives and options for design and planning. The scale and complexity of the problem demand a new set of algorithms and methodologies for visualizing rich, intricate, and dynamic urban landscapes with constant flows of crowds and traffic.

Numerous efforts have been devoted in acquiring and visualizing "urbanscape". Over the last decade, there has been considerable progress on multiple fronts: acquisition of imagery and 3D models using improved sensing technologies, real-time rendering, and procedural modeling. For example, aerial imagery of most cities is used in Google Earth and Microsoft Virtual Earth. The problem of reconstructing 3D geometric models from videos and scanners has been an active area of research in computer vision and related areas. Similarly, many efficient techniques have been proposed to stream the imagery and geometric data over internet and display them at real-time rates on high end workstations or handheld devices. However, all these efforts are limited to capturing, displaying, or modeling predominantly static models of urbanscapes and do not include dynamic elements, such as crowds or traffic. In many aspects, the realism of models shown in Google Earth or Microsoft Virtual Earth is lacking due to the absence of dynamic behaviors.

Figure 1: *An example of simulated crowds at Shibuya crossing in Japan*

In addition to high-rise buildings and architectural scenes on city landscapes, moving pedestrians and vehicle traffic are an integral part of any metropolitan region, yet they have not received sufficient attention. Aggregates of numerous entities, such as a group of people and fleet of vehicles, form complex systems that exhibit interesting biological, social, cultural, and spatial patterns observed in nature and in society. Modeling of the collective behaviors remains an open research challenge in artificial intelligence, computer vision, architecture, physics, psychology, social sciences, and civil and traffic engineering, as complex systems often exhibit distinct characteristics, such as emergent behaviors, self-organization, and pattern formation, due to multi-scale interactions among individuals and groups of individuals, despite of decades of observation and studies.

2 Research Challenges

The challenges in real-time modeling and simulation of digital cityscape stem from its extremely large scale, i.e. in the range of hundreds of thousands or even millions, crowds and vehicle traffic commonly encountered in metropolitan areas across the globe. We refer to such a physically vast scale of computational challenges as "metropolitan scale." Below we briefly list a few problems in realizing this vision and provide pointers to some recent progress toward this goal:

[1] "World Urbanization Prospects" by *United Nations* Population Division, Department of Economic and Social Affairs, 2005.

Figure 2: *An example of reconstructed traffic in an European cityscape*

- **Modeling of intricate pedestrian dynamics that leads to better understanding of complex crowd pheonomena**: Recently we have developed a new trajectory planning algorithm for virtual humans. Our approach focuses on implicit cooperation between multiple virtual agents in order to share the work of avoiding collisions with each other. Specifically, we extend recent work on multi-robot planning to better model how humans avoid collisions by introducing new parameters that model human traits, such as reaction time and biomechanical limitations. We validate this new model based on data of real humans walking captured by the Locanthrope project. Extending such approach to many thousands or millions of people in a large crowd remains a significant challenge. See:

  http://gamma.cs.unc.edu/RCAP

  http://gamma.cs.unc.edu/PLE

- **Real-time reconstruction metropolitan-scale traffic flows given discrete temporal-spatial sensor data**: We introduce a novel concept, Virtualized Traffic, to visualize reconstructed continuous traffic flows from traffic sensor datas. Given the positions of each car at two recorded locations on a highway and the corresponding time instances, our approach can recreate the traffic flows (i.e. the dynamic motions of multiple cars over time) in between the two locations using a priority-based scheme for multiple agents. Our algorithm is applicable to high-density traffic on highways with a small number of lanes and takes into account the geometric, kinematic, and dynamic constraints on the cars. Although our framework can process a continuous stream of input data in real time by reducing the search space for planning, extending such approaches to a large number of lanes with finer discretization to better approximate contiuous motion makes this approach quickly intractable. More efficient techniques would be needed. See:

  http://gamma.cs.unc.edu/TRAFFIC-RECON

- **Data-driven personality models based on perceptual studies for simulating crowd and driver behaviors**: To generate heterogeneous crowd behaviors using personality trait theory, we adopt results of a user study to derive a mapping from crowd simulation parameters to the perceived behaviors of agents in computer-generated crowd simulations. We establish a linear mapping between simulation parameters and personality descriptors corresponding to the well-established Eysenck Three-factor personality model. Furthermore, we propose a novel two-dimensional factorization of perceived personality in crowds based on a statistical analysis of the user study results. Extension to this approach to establish dynamic mappings and factorizations for generating heterogeneous crowd behaviors in settings with external factors (such as interaction with other agents, environments, and other stress factors) would need to be considered as well. See:

  http://gamma.cs.unc.edu/personality

- **Applications to traffic rerouting and congestion management**: While state-of-the-art systems take into account current trafc conditions or historic trafc data, current approaches ignore the impact of their own plans on the future trafc conditions. We introduce a novel algorithm for self-aware route planning that uses the routes it plans for current vehicle traffic to more accurately predict future traffic conditions for subsequent cars. Our planner uses a roadmap with stochastic, time-varying traffic densities that are defined by a combination of historical data and the densities predicted by the planned routes for the cars ahead of the current trafc. We have applied our algorithm to moderate-scale traffic route planning, and demonstrated that our self-aware route planner can more accurately predict future traffic conditions, which results in a reduction of the travel time for those vehicles that use our algorithm. Extension of such planning and simulation framework to metropolitan-scale traffic that incorporates dynamic sensing and real-time traffic prediction would introduce new challenges to multi-agent simulations. See:

  http://gamma.cs.unc.edu/TROUTE

  Other applications including emergency response and planning, architecture and engineering design evaluation, etc. should also be investigated. In addition, validation of such techniques should be also addressed in the context of applications.

3 Conclusion

We have suggested a list of problems in developing digital cityscapes that can benefit from applications of more advanced multi-agent planning and simulation algorithms and techniques. Addressing these problems can lead to attaining plausible explanations of the behavior and motivation of individual agents (e.g. pedestrians or vehicles) and how they interact with each other under different settings, across varying scales and levels of social organizations, from individuals to groups, with applications ranging from urban planning, civil and traffic engineering, transportation system design, architectural layout, training of first-responders electronic commerce, to education and entertainment.

Data Prospecting Framework for Geoscience

Dr. Rahul Ramachandran (*rahul.ramachandran@uah.edu*)
Information Technology & Systems Center, University of Alabama in Huntsville

Vision

Decade-long, big-science projects, such as the Human Genome project, Large Hadron Collider, Sloan Digital Sky Survey, and Earth Observation System have significantly advanced our understanding and revolutionized domains by collecting and analyzing datasets of unprecedented size. Due to advances in sensors, computation and storage, the cost and effort required to produce datasets of comparable scale is decreasing significantly. As a result, we are seeing a proliferation of large amounts of data being assembled in almost every science field, from the core sciences to physical sciences and engineering, to social sciences. The scientific opportunities inherent in these large datasets are enormous: novel hypotheses become evident by combining and analyzing large amounts of data, and data-intensive science is now considered the fourth paradigm of science [1]. In a data-centric approach, scientists typically first explore data using standard statistical analysis tools and visualization, and then perform sophisticated data mining to extract deeper information and create knowledge [2].

In geosciences, given ever-increasing data volumes, scientists often find themselves conducting more data storage and management tasks rather than entirely focusing on data exploration and data mining. Most providers of large datasets provide simple mechanisms for data exploration, such as search based on metadata. Consequently, scientists first have to download and manage large volumes of data before undertaking any meaningful data exploration and data mining. For many scientists who fall in the "long tail" of science and are resource-constrained [3], focusing significantly on data storage and management tasks affects productivity; such tasks become the rate-limiting step towards making impact-worthy science contributions.

We believe a drastic reduction in scientists' data management tasks can be achieved if providers of large data improve the mechanisms for data exploration, and in particular, allow scientists to conduct statistical analysis-based data exploration on large-scale data as part of the data selection process. Such enhanced exploration will provide immediate interaction with and feel for the data through histograms, similarity searches, and visual analysis. Scientist can then isolate and download portions of some potential datasets, which are often smaller in size, and analyze them deeply using their favorite analytical algorithms and progress towards new hypotheses. This enhanced data exploration process is akin to geophysical prospecting, in which mineral sites of interest are first identified over the vast landscape through appropriate screening methods and then more expensive ore-extraction methods are employed. Based on this metaphor, a ***data prospecting*** approach that enables statistical analysis on the content of the data can lead to improvements in storage, management, and selection of large datasets within scientific disciplines.

This vision and rationale leads to an effective business model. *The proposition of data prospecting creates utility for both the scientists and the big data centers serving geospatial data. By enabling data prospecting as a service, big data centers reduce costs of data transfers and improve their data discovery process. Scientists improve productivity by eliminating time-consuming tasks that arise due to movement, storage and management of data of various types and formats*. In addition, the proposition is of tremendous value to the NSF EarthCube initiative, which is transforming the conduct of geosciences research through community-guided knowledge-management cyberinfrastructure. Towards this goal, EarthCube established the Data Discover, Mining and Access (DDMA) community workgroup [4] to create a roadmap, through community consensus, for the current and future data needs within the geosciences. The roadmap called for integrative approaches that discover geoscience data, enable efficient access, and allow core capabilities for analysis and mining. In addition, it highlighted the universal needs (i) to collocate computing with data, (ii) to develop scalable data access methods, and (iii) for "scientist-friendly" approaches and environments to fully exploit data mining technologies.

Example Prototype

In the DISCOVER project[1], one science objective is to isolate atmospheric phenomena such as a cumulus cloud, a thunderstorm shower, a rogue wave, a tornado, an earthquake, a tsunami, a hurricane, or an El Niño within a large dataset. In general, detecting these phenomena requires scientists to download very large datasets from data providers and conduct data exploration and data mining analyses on their own. However data prospecting can substantially improve data exploration by summarizing events, which are episodes of geoscience phenomena. An event has a finite duration and an associated geo-location as a function of time and can be viewed as an entity in four-dimensional (4D) spatiotemporal space. Summarizations build representations of a given type of event, by studying the characteristics and distribution of a large number of events, such as spatio-temporal distribution, intensity, annual cycle, duration, etc.

Towards improving data exploration, Dr. Ramachandran has developed an initial distributed client-server system, the Visual Data Exploration Environment[2], for demonstrating data prospecting capabilities. The server hosts four data products derived from multiple Special Sensor Microwave Imager instruments, staged on a modest cluster. The client can issue queries of four kinds, viz., 1D and 2D histograms, descriptive statistics and thresholding [5]. An initial demonstration to geoscience researchers has met with very positive feedback. In one interactive session lasting about an hour, a researcher was able to identify several phenomena in the data set such as Mistral, Somali Jet, and Wind patterns around the Hawaiian Islands. Similarly, another researcher was able to identify Tehuantepecer in the data set within minutes. Tehauntepecer is a gap wind that is triggered by a synoptic scale high pressure system over the Great Plains of North America that pushes air through a narrow Sierra Madre mountain range gap. The analysis of such events is important because of their significant regional climate impact. A screen shot showing the results of this analysis is presented in Figure 1.

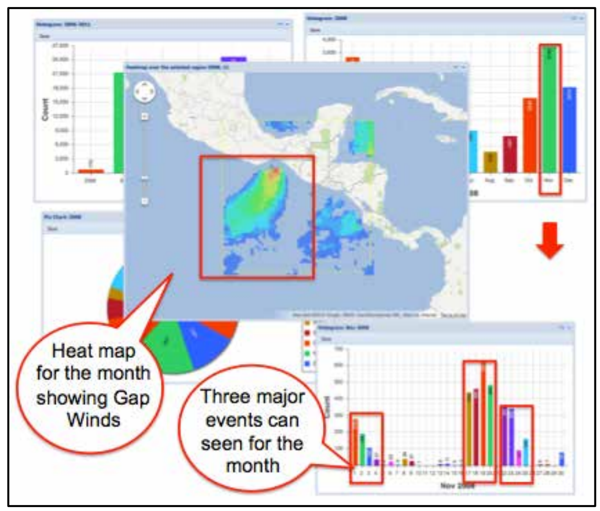

Figure 1. Data prospecting prototype showing visual analysis results.

While the current prototype is very preliminary both in design and capabilities it has clearly demonstrated the value of data prospecting to the geoscience research community. Data prospecting capabilities need to be inherent part of any spatial computing road map to support data intensive science. As such infrastructure requirements to support this capability such as suitable data organization methods, scalable data processing, enhanced data content-based searches, etc. must addressed.

[1] http://discover.itsc.uah.edu

[2] http://jurassic3.itsc.uah.edu/vdee

References

1. Hey, A.J.G., Tansley, S., Tolle, K.M., *The fourth paradigm: data-intensive scientific discovery*, 2009, Microsoft Research Redmond, WA.

2. Shoshani, A., and Rotem, D., *Scientific Data Management: Challenges, Technology, and Deployment*2009: Chapman and Hall/CRC.

3. Heidorn, P.B., *Shedding Light on the Dark Data in the Long Tail of Science.* Library Trends, 2008. **57**(2).

4. The Earthcube Data Discovery Mining and Access Group. *"The DDMA EarthCube RoadMap"*. 2012; Available from: https://docs.google.com/document/d/1RtR8A0h0HFhZP9kIPFlfCChFIYLIf-G-TZMbNSzw58M/edit.

5. Rushing, J., R. Ramachandran, S. Graves, et al. *Visual Data Exploration Environment for Data Intensive Science*. in *IEEE Geoscience and Remote Sensing Society* 2012. Munich.

From the mirror worlds to everyware in the metaverse:
Or what is special about spatial (computing)?

Daniel Sui
Department of Geography
The Ohio State University
Email: sui.10@osu.edu

Eight years ago, along with bio- and nano- technologies, geospatial technologies were touted as one of the three defining technologies in the 21st century (Gewin, 2004). Compared to bio- and nano- technologies, geospatial technologies are obviously still the underdog among three. As a geographer interested GIScience and GeoComputation, I am pleased to notice colleagues in multiple other fields are increasingly drawn to the spatial turn, as evidenced by the growing interests in spatial humanities or spatially integrated social sciences. It is also gratifying to notice a growing numbers of computer scientists are interested in dealing with geo- or spatial aspects of computing.

After reading the workshop proposal, one question instantly came to mind – what is the meaning of the term "spatial computing." Without sounding too pedantic, spatial computing immediately reminds me of multiple other labels/words we have used so far: GeoComputation, context-aware computing, pervasive computing, ubiquitous computing, ambient intelligence, sentient city, and smart planet etc. As most of the critical scholars often argue that new world creation always starts with creation of new words. My position paper outlines some preliminary thoughts on what this new word – spatial computing – means to me, followed by discussions on what kinds of new worlds this new word will create and what kinds of research we should be focusing on in the years ahead.

The question I have been struggling with after I read the spatial computing workshop proposal is, if I can borrow Luc Anselin's (1989) early phrase, what is special about spatial computing? By special I mean to move above and beyond what has been done during the past two decades related to spatial data handling. To me, I believe the discussion on the meaning of the term "spatial computing" should be conducted in the context of ubiquitous computing and the emerging metaverse[1].

Back in 2005, I wrote a short piece in CEUS discussing the implications of ubiquitous computing on GIS (Sui, 2005). It is interesting to notice the great asymmetry of the literature: very few researchers in the geospatial communities have paid due attention to ubicomp whereas more folks in the ubicomp communities have started working on locational and spatial aspects of ubicomp. I am sure that CCC's workshop on "spatial computing" will promote more meaningful interactions between GIScientists and ubicomp researchers.

While speculating about the future of computers in the 21st century, Mark Weiser (1991) argued that ''the most profound technologies are those that disappear. They weave themselves into the fabric of everyday life until they are indistinguishable from it (p. 94).'' Although the computer as we know it has not completely vanished as Weiser (1991) predicted, advances in ubiquitous computing (ubicomp) during the past two decades have accelerated the pace toward embedding more and more computer chips into the devices and our surroundings, e.g. mobile phones, car navigational systems, gas pumps, ATM machines, electronic road/bridge tolls, and smart home products in sentient cities. Additionally, radio chips, led by the radio frequency identification (RFID) technology, are designed to replace barcodes on manufactured objects. RFID, along with existing location technologies, will be able to make the location of every single entity on earth trackable (NRC, 2001, 2003). Hand-held communication media such as smart phones or iPads can easily mutate into wearable remote-control devices for the physical world. Instead of the traditional distinction on hardware and software, we have witnessed the emergence of everyware (Greenfield, 2006) as ubicomp replaces the traditional mainframe and desktop computers to become the dominant paradigm for computing.

Concomitant with the growth of ubicomp and everyware, we are also rapidly entering a new age of metaverse – a hybrid world in which the virtual world based upon bits is increasingly linked to the atom-based physical world (www.metaverseroadmap.org). Nowadays when I think about GIS in general and spatial computing particular, I cannot separate it from the emerging metaverse. I believe that metaverse roadmap developed by the cross-industry foresight group could serve as a possible roadmap for us to think about the challenging issues related to spatial computing.

The rapidly evolving metaverse is a result of several converging technologies. According to the metaverse road map report, the browser for engaging this metaverse will be based upon a 3-D Web that brings together the following four technologies:

● Mirror worlds – digital representations of the atom-based physical world, such as Google Earth, Microsoft Virtual Earth, NASA's World Winds, ESRIs ArcGlobe, USGS's National Map, and the massive georeferenced GIS data bases developed during the past fifty years.
● Virtual worlds - digital representations of the imagined worlds, such as Second Life, World of War Craft, computer games, various cellular automata models, and agent-based models;
● Lifelogging - the digital capture of information about people and objects in the real or digital worlds, such as twitter, blogs, flickr, YouTube, social networking sites such FaceBook or MySpace.
● Augmented reality – sensory overlays of digital information on the real and virtual worlds using head-up displays (HUDs) or other mobile/wearable devices such as cell phones or sensors via participatory sensing.

Viewed from a metaverse perspective, the workshop proposal for spatial computing seems to have focused almost exclusively on the components of mirror worlds. For this workshop, I hope we can shift the discussion on spatial computing from the mirror worlds to the broader field of ubicomp/everyware in the emerging metaverse. By making this shift, we will inevitably enter into a world where space or spatial may have a different meaning. For example, Kitchin and Dodge (2011) made a distinction between coded space (where information is inscribed digitally that enhances the functioning of a particular environment) and code/space (where information and space are so fused that the space cannot function without the information and there is no uncoded, manual alternative). Will spatial computing be capable of analyzing and model the ontogenetic process Kitchin and Dodge (2011) outlined for the ubicomp (everyware) environment we are in right now? I think so.

Endnote:

1. First coined by Neal Stephenson in his 1992 science fiction novel *Snow Crash*, metaverse refers to a fictional virtual world where humans, as avatars, interact with each other and software agents, in a three-dimensional space that uses the metaphor of the real world.

References cited

Gewin, V. (2004). Mapping opportunities. Nature, 427(6972), 376–377. Anselin, L. 1989. What is special about spatial data? Alternative perspectives on spatial data anlaysis. NCGIA Technical Report. Santa Barbara, CA.

Greenfield, A., 2006. Everyware: The Dawning Age of Ubiquitous Computing. New Riders Publishing.

Kitchin, R. and M. Dodge (2011). Code/Space: Software and Everyday Life. Cambridge,MA.: MIT Press.

NRC, 2001. Embedded, Everywhere: A Research Agenda for Networked Systems of Embedded Computers. Washington, D.C.: National Academies of Press.

NRC, 2003. IT roadmap to a geospatial future. Washington, DC: National Academy Press. Weiser, M. (1991). The computer for the twenty-first century. Scientific American, September, pp. 94–10.

Sui, D.Z., 2005. Will ubicomp make GIS disappear? *Computers, Environment, and Urban Systems* 29(1): 361-367.

Big Spatiotemporal Data Analytics: Recent Advances and Future Challenges

Ranga Raju Vatsavai ({vatsavairr}@ornl.gov)
Computational Sciences and Engineering Division
Oak Ridge National Laboratory, Oak Ridge, TN 37831, USA.

Keywords: Big Data, Remote Sensing, Data Mining

1 Introduction

We are living in the era of 'Big Data.' Spatiotemporal data, whether captured through remote sensors (e.g., remote sensing imagery, Atmospheric Radiation Measurement (ARM) data) or large scale simulations (e.g., climate data) has always been 'Big.' However, recent advances in instrumentation and computation making the spatiotemporal data even bigger, putting several constraints on data analytics capabilities. Spatial computation needs to be transformed to meet the challenges posed by the big spatiotemporal data. The Geographic Information Science and Technology (GIST) group has been engaged in developing novel spatiotemporal data mining and machine learning approaches to efficiently process big spatiotemporal databases to extract knowledge that is highly useful for various stakeholders. Through this white paper we share some of the recent advances made in big spatiotemporal data analytics, specifically in the area of remote sensing, and point the community to some of future research challenges.

2 Data Challenges

Recent advances in remote sensing instrumentation, and commercialization of remote sensing technology has lead to the unprecedented growth in the acquisition and archival of high resolution imagery. Figure 1 shows progression of remote sensing instruments along three important sensor characteristics: spatial, spectral, and temporal resolutions. Though these improvements are leading to increase in volume, velocity, and variety of remote sensing data products making it hard to manage and process, they are also enabling new applications. For example, improvements in temporal resolution allows monitoring biomass on a daily basis.

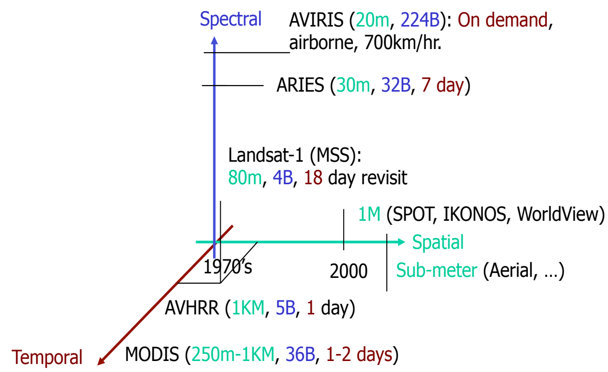

Figure 1. Advances in remote sensing data products (1970's through present)

Improvements in spatial resolution allows fine-grained classification (settlement types), damage assessments, and critical infrastructure (e.g., nuclear proliferation) monitoring. Now let us consider biomass monitoring application requirements more closely from big data perspective.

Monitoring biomass over large geographic regions for identifying changes is an important task in many applications. With recent emphasis on biofuel development for reducing dependency on fossil fuels and reducing carbon emissions from energy production and consumption, the landscape of many countries is going to change dramatically in coming years. Already there are several preliminary reports that address both economic and environmental impacts of growing energy crops. These changes are not limited to the United States alone. Developing countries like India, the rural areas are facing increasing demand for energy. Monitoring biomass using daily remote sensing imagery for identifying changes is a critical application in achieving food and energy security. Satellite data products collected from the MODerate Resolution Imaging Spectroradiometer (MODIS) sensor can provide global coverage at a spatial footprint of 250 meters and a fine temporal resolution of one day. Analyzing MODIS data at regional and global scales poses several computational and I/O challenges. Since data at global scale is difficult to handle, MODIS data is organized into tiles of 10 x 10 degrees (4800 x 4800 pixels). Though there are 460 daily MODIS tile products available, we need to process 326 products, which contain land pixels, making it very hard to process the data in a day before new data comes in. On the other hand, monitoring critical infrastructure requires high-resolution (both spectral and spatial) satellite imagery. Just to give an understanding of data complexity, we have collected more than 10 TB of data for post Katrina damage assessments which included both imagery products (Digital Otho Images, ASTER, IKONOS, LANDSAT-5/7, Quickbird, and SPOT) and vector (digital elevation, flood contours, etc.) data. Table 1 summarizes computational and storage requirements for major applications being pursued by our research group.

Application	Computational (TFlops)	Storage (TBs)
LandScan Global	50	250
LandScan HD + Settlement Mapping	50	2000
Global Infrastructure, Population Mobility, Evacuation Modeling	50	5
Biomass and Nuclear Proliferation Monitoring, Damage Assessments	50	150
Mobile/Trailer Park Mapping	10	50

Table 1. Computational and Storage Requirements of Few Big Spatiotemporal Data Applications

3 Algorithmic Challenges

Most of the pattern recognition and machine learning algorithms are per-pixel based (or single instance). These methods worked well for thematic classification of moderate and high-resolution (5 meters and above) images. Very high-resolution (VHR) images (submeter) are offering new opportunities beyond thematic mapping, they allow recognition of structures in the images. For example, consider the problem of settlement mapping [2]. The high rate of urbanization, political conflicts and ensuing internal displacement of population, and increased poverty in the 20th century has resulted in rapid increase of informal settlements. These unplanned, unauthorized, and/or unstructured homes, known as informal settlements, shantytowns, barrios, or slums, pose several challenges to the nations as these settlements are often located in most hazardous regions and lack basic services. Though several World Bank and United Nations sponsored studies stress the importance of poverty maps in designing better policies and interventions, mapping slums of the world is a daunting and challenging task. VHR images provides the ability to distinguish informal settlements from formal settlements. However, per-pixel based methods do not work well for very high-resolution (VHR) images (submeter). The main problem being that the pixel size (less than meter) is too small as compared to the object size (10s of meters) and contains too little contextual information to accurately distinguish between given set of pixels. As shown in Figure 2 often do not provide sufficient discrimination power between classes. One way to alleviate this problem is to consider a bigger window or patch consisting a group of adjacent pixels which offers better spatial context than a single pixel. Unfortunately, this makes all well known per-pixel based classification schemes ineffective. Multi-instance learning approaches might be useful in moving from pixel-based or object-based structure recognition in VHR images, but computational complexity is too high to be practically applied for global settlement mapping.

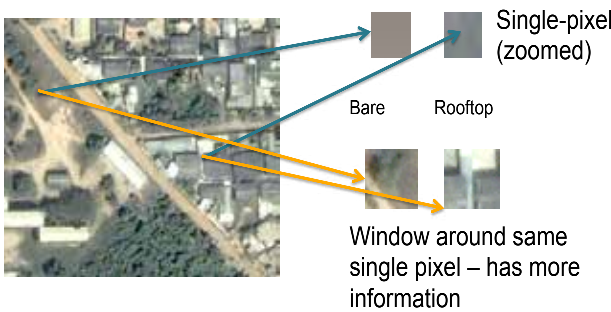

Figure 2. Problems with pixel-based pattern recognition methods

Now, let us consider the problem of identifying complex facilities (e.g., nuclear facilities, thermal power plants) [3] in VHR images. As can be seen from Figure 3, thematic classification is designed to learn and predict thematic classes such as forest (F), crops (C), buildings (B), etc., at pixel level. However, such thematic labels are not enough to capture the fact that the

(a) FCC Image with Thematic class labels (B-Buildings, C-Crop, F-Forest)

(b) Thematic Classified Image (B-Buildings, C-Crop, F-Forest)

(c) FCC Image with Semantic Labels (S-Switch Yard, C-Containment Building, T-Turbine Generator, CT-Cooling Tower

Figure 3. Thematic vs. Semantic Classes

given image contains a nuclear power plant. What is missing is the fact that the objects, such as switch yard (S), containment building (C), turbine building (T), and cooling towers (CT) have distinguishing shapes, sizes, and spatial relationships (arrangements or configurations) as shown in Figure 3(c). These semantics are not captured in the traditional pixel- and object-based classification schemes. In addition, traditional image analysis approaches mainly exploit low-level image features (such as, color and texture and, to some extent, size and shape) and are oblivious to higher level descriptors and important spatial (topological) relationships without which we can not accurately discover these complex objects or higher level semantic concepts. Figure 5 shows four different images (baseball and football fields, two residential neighborhoods) where they share common objects, for example, grass and soil across baseball and football fields, and two (economically) different neighborhoods where in one neighborhood buildings are colocated with cars (parked on the road) while in the other builds are collocated with swimming pools. Both pixel- and object-based methods often fails to capture these complex relationships. Future research requires models that explicitly learn complex spatial relationships among the objects to accurately predict semantic classes and scale to big VHR image collections.

4 Computational Challenges

In general, increasing resolutions (spatial, spectral, and temporal) and as well modeling of spatial and temporal constraints leads to increased computational complexity. For example, the Gaussian Process (GP) based change detection technique [1] developed to monitor biomass using MODIS NDVI time series signals is computationally expensive (time complexity is $O(n^3)$) and memory requirements are of $O(n^2)$, making this infeasible for large study regions. At daily temporal resolution, MODIS time series contains about 3600 data points (at each pixel location). We need to process about 7,511,040,000 time series, where each time series contains 3600 data points. We need to process this data in a day before new set of MODIS products arrive. Using GP-based change detection, sequential code requires about two days to process this data. We developed efficient and parallel techniques using shared and distributed memory models which make it possible to apply this technique for continuous monitoring of biomass at regional scales. However, further research is needed to scale these algorithms for daily temporal resolution at continental scales by utilizing the modern computing infrastructure consisting of data staging and intelligent I/O techniques. Modern computing architectures can be leveraged for spatial computation, especially GPU and cloud computing. However, there are several challenges that need to be addressed before these technologies can be widely adopted. For example, in our experiments, widely used Gaussian Mixture Model (GMM) based clustering scaled upto 70x on GTX-285 GPU (240 CUDA cores, 1 GB), but the main bottleneck is I/O (data transfer between host and device memory). On the other hand, many spatial computation tasks are not amenable for simple data parallelization. For example, consider the clustering of a large image (or set of images) stored in a distributed file system (e.g., HDFS). GMM clustering requires costly exchange of data samples between the nodes. We developed a novel distributed clustering algorithm (Figure 4) which estimates local models from the data available at each node (map task) and computes global model (reduce task) by just exchanging the model parameters. This clustering algorithm is shown to scale linearly in terms of nodes. Further research is needed for efficient data partitioning.

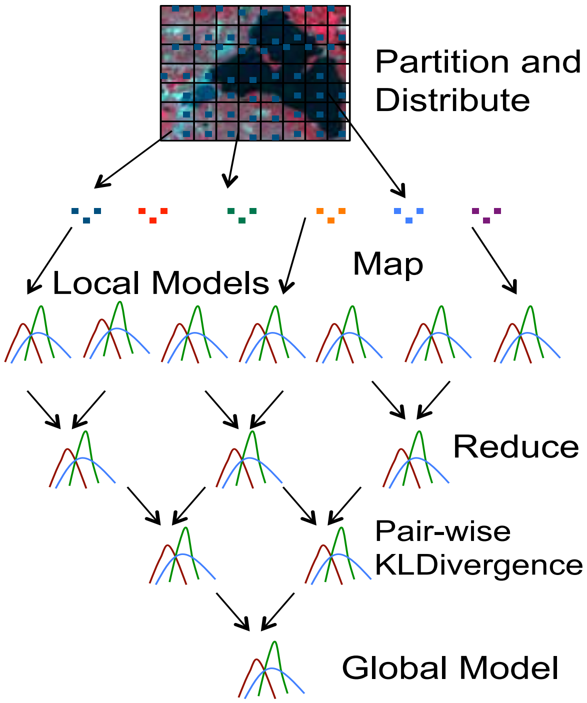

Figure 4. Distributed GMM Clustering Algorithm

5 Conclusions

Big spatiotemporal data, though opening up new applications, posing several challenges. New approaches are required to overcome both computational and I/O challenges, and new models that explicitly model spatial and temporal constraints efficiently.

6 Acknowledgments

I would like to thank my collaborators Prof. S. Shekhar, Prof. T. Burk, Dr. B. Bhaduri, Mr. E. Bright, Dr. A. Cheriyadat, Dr. V. Chandola, and Mr. J. Graesser. Prepared by Oak Ridge National Laboratory, P.O. Box 2008, Oak Ridge, Tennessee 37831-6285, managed by UT-Battelle, LLC for the U. S. Department of Energy under contract no. DEAC05-00OR22725.

(a) Baseball Field (b) Football Field

(c) Building + Car (d) Building + Swimming Pool

Figure 5. Example images sharing similar objects (e.g., grass, buildings, roads, cars, water) but entirely different global labels

References

[1] V. Chandola and R. R. Vatsavai. A scalable gaussian process analysis algorithm for biomass monitoring. *Statistical Analysis and Data Mining*, 4(4):430–445, 2011.

[2] J. Graesser, A. Cheriyadat, R. R. Vatsavai, V. Chandola, J. Long, and E. Bright. Image based characterization of formal and informal neighborhoods in an urban landscape. *Selected Topics in Applied Earth Observations and Remote Sensing, IEEE Journal of*, 5(4):1164 – 1176, August 2012.

[3] R. R. Vatsavai, A. Cheriyadat, and S. S. Gleason. Unsupervised semantic labeling framework for identification of complex facilities in high-resolution remote sensing images. In *ICDM Workshops*, pages 273–280, 2010.

SPATIAL COMPUTING SYSTEMS

Integrating Spatial-features in Data-centric Applications

Divyakant Agrawal
Department of Computer Science
University of California at Santa Barbara
Santa Barbara, CA 93106

August 15, 2012

1 Introduction

Contemporary information and data rich applications often have to be integrated with spatial features. In this position paper, we illustrate the research challenges that arise in integrating spatial features in data-centric applications. Our intent in presenting these case-studies is to underscore the importance of spatial feature integration as well as demonstrate that this integration gives rise to significant complexity. Given that location and time information will be an integral feature of all end-user devices, it is imperative that there is a systematic exploration of research and development challenges in the area of spatial computing.

2 Location-based Services

The last few years have witnessed a significant increase in hand-held devices becoming location aware with the potential to continuously report up-to-date location information of their users. This has led to a large number of location based services (**LBS**) which customize a user's experience based on location. Some applications—such as customized recommendations and advertisements based on a user's current location and history—have immediate economic incentives, while some other applications—such as location based social networking or location aware gaming—enrich the user's experience in general. With major wireless providers serving hundreds of millions of subscribers [wir10], millions of devices registering their location updates continuously is quite common. Database management systems (DBMS) driving these location based services must therefore handle millions of location updates per minute while answering near real time analysis and statistical queries that drive the different recommendation and personalization services.

Location data is inherently multi-dimensional, minimally including a user id, a latitude, a longitude, and a time stamp. A rich literature of multi-dimensional indexing techniques—for instance, K-d trees [Ben75], Quad trees [FB74] and R-trees [Gut84]—have empowered relational databases (RDBMS) to efficiently process multi-dimensional data. However, the major challenge posed by these location based services is in scaling the systems to sustain the high throughput of location updates and analyzing huge volumes of data to glean intelligence. For instance, if we consider only the insert throughput, a MySQL installation running on a commodity server becomes a bottleneck at loads of tens of thousands of inserts per second; performance is further impacted adversely when answering queries concurrently.

On the other hand, Key-value stores, both in-house systems such as Bigtable [CDG$^+$06] and their open source counterparts like HBase [hba10], have proven to scale to millions of updates while being fault-tolerant and highly available. However, Key-value stores do not natively support efficient multi-attribute access, a key requirement for the rich functionality needed to support LBSs. In the absence of any filtering mechanism for secondary attribute accesses, such queries resort to full scan of the entire data. MapReduce [DG04] style processing is therefore a commonly used approach for analysis on Key-value stores. Even though the MapReduce framework provides abundant parallelism, a full scan is wasteful, especially when the selectivity of the queries is high. Moreover, many applications require near real-time query processing based on a user's current location. Therefore, query results based on a user's stale location is often useless. As a result, a design for batched query processing on data periodically imported into a data warehouse is inappropriate for the real-time analysis requirement.

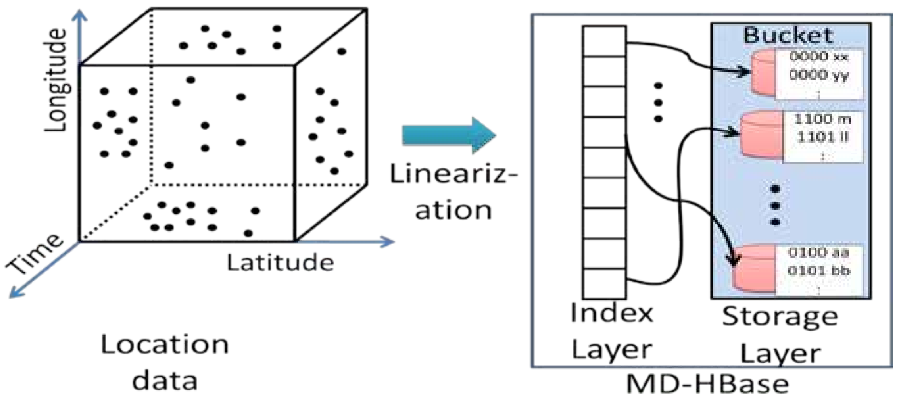

Figure 1: Architecture of $\mathcal{MD}$-HBase.

RDBMSs provide rich querying support for multi-dimensional data but are not scalable, while Key-value stores can scale but cannot handle multi-dimensional data efficiently. Our solution, called $\mathcal{MD}$**-HBase**, bridges this gap by layering a multi-dimensional index over a Key-value store to leverage the best of both worlds.[1] We use linearization techniques such as Z-ordering [Mor66] to transform multi-dimensional location information into a one dimensional space and use a range partitioned Key-value store (HBase [hba10] in our implementation) as the storage back end. Figure 1 illustrates $\mathcal{MD}$-HBase's architecture showing the index layer and the data storage layer. We show how this design allows standard and proven multi-dimensional index structures, such as K-d trees and Quad trees, to be layered on top of the Key-value stores with minimal changes to the underlying store and negligible effect on the operation of the Key-value store. The underlying Key-value store provides the ability to sustain a high insert throughput and large data volumes, while ensuring fault-tolerance and high availability. The overlaid index layer allows efficient real-time processing of multi-dimensional range and nearest neighbor queries that comprise the basic data analysis primitives for location based applications.

We evaluated different implementations of the data storage layer in the Key-value store and evaluate the trade-offs associated with these different implementations. In our experiments, $\mathcal{MD}$-HBase achieved more than 200K inserts per second on a modest cluster spanning 16 nodes, while supporting real-time range and nearest neighbor queries with response times less than one second. Assuming devices reporting one location update per minute, this small cluster can handle updates from $10 - 15$ million devices while providing between one to two orders of magnitude improvement over a MapReduce or Z-ordering based implementation for query processing. Moreover, our design does not introduce any scalability bottlenecks, thus allowing the implementation to scale with the underlying Key-value data store.

3 Spatio-temporal Trends in Online Social Networks

Information that is shared in a social network may have certain semantic properties such as the location and time. For instance, one might be interested to know the trends in California alone or short/long term trends . Such queries cannot be answered using trends analysis at the scale of the entire network. Therefore we believe there is a need for trend definitions that explore such dimensions. Our belief is also supported by the growing body of research in this field [TLP$^+$08, SST$^+$09].

Spatial trends can be defined in various ways. For instance, the goal can simply be to detect heavy hitters for each location. However, such a technique fails at identifying topics of true geographical nature since a topic of global importance incidentally also has a high frequency of occurrence in various localities without really being related to such locations. Distinguishing such a topic from ones that are trending in only certain localities is not possible without considering the *correlations* between places and topics. Therefore, we need to focus on the problem of identifying the correlation of information items with different geographical places. We propose *GeoWatch*: an algorithmic tool for detecting geotrends in online social networks by reporting *trending* and *correlated* location-topic pairs. *GeoWatch* also captures the temporality of trends by detecting geo-trends along a sliding window. With the use of different window

[1]The name $\mathcal{MD}$-HBase signifies adding multi-dimensional data processing capabilities to HBase, a range partitioned Key-value store.

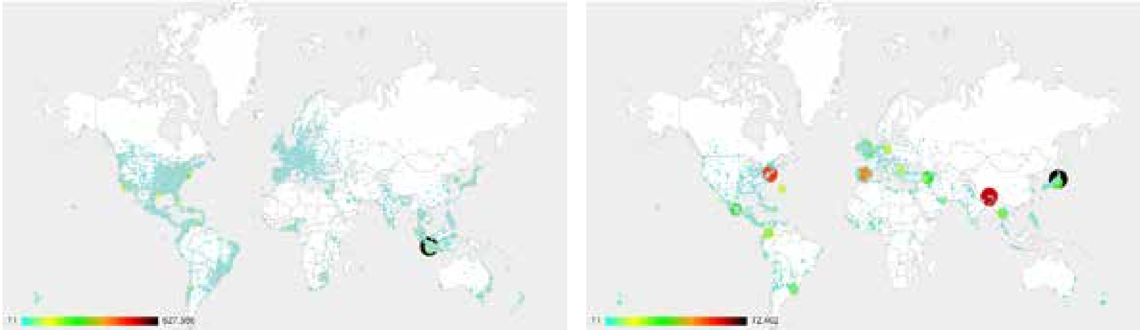

(a) Tweets *in* Cities

(b) Tweets *about* Cities

Figure 2: Heat Map for # of tweets in/about cities of the world

sizes, trends of different time granularity can be detected. Our analysis on a Twitter data set shows that such geo-trend detection can be very important in detecting significant events ranging from emergency situations such as earthquakes to locally popular flash crowd events such as political demonstrations or simply local events such as concerts or sports events. The fast detection of emergency events such as the March 11 Japan earthquake indicates the possible value of *GeoWatch* in crisis management. In Figure 2, we present a heat map of tweets for a period of approximately 2 months of tweets (March 9 to May 8, 2011). More particularly, we capture the volume of tweets originating from various cities in Figure 2(a) and tweets about cities in Figure 2(b). In these plots, every city associated with more than 10 tweets is marked– color and size is proportional to the number of tweets. Our approach helps identify various characteristics of the social network usage. The two figures resemble each other but there are certain interesting distinctions. It is worthwhile to note that the part of the map corresponding to Japan is denser in Figure 2(b). This is mostly due to the Japan Earthquakes that took place within the time period captured in our data set. This important event spanned a long time period due to the after effects and was an important headliner, making it a trending topic in Twitter. On the contrary, a drop in significance can be observed for countries such as Indonesia when comparing the tweets *in* cities to tweets *about* cities. This big difference originates from the fact that Indonesia is a highly active country for Twitter [ind], while there are no important events taking place in its cities that would result in people mentioning them.

Formally, given a stream S of location-topic pairs of the form (l_i, t_j), a window size of N, and three user defined frequency thresholds θ, ϕ, and ψ in the interval [0, 1]; our goal is to keep track of all locations l_i s.t. $F(l_i) > \lceil \theta N \rceil$ alongside their frequencies as well as all topics t_x and their frequencies $F(t_x)$. In addition, in order to detect the correlations, we aim to find all pairs (l_i, t_x) s.t. $F(l_i) > \lceil \theta N \rceil$, $F(l_i, t_x) > \lceil \phi F(l_i) \rceil$, and $F(l_i, t_x) > \lceil \psi F(t_x) \rceil$; where $F(l_i, t_x)$ is the number of information items on topic t_x from location l_i in the most recent N items in S; $F(l_i)$ is the aggregate number of occurrences of all the items from l_i in the current time window; and $F(t_x)$ is the aggregate number of items on t_x. The window size can be set in terms of maximum number of elements or an actual time window such as an hour or a day. In the latter case, the number of elements N in the current window is variable.

We now explore a sketch-based structure for *GeoWatch* to detect correlations between locations and topics. As can be seen from Figure 3, *GeoWatch* consists of two main components. *Location-StreamSummary-Table* contains a $StreamSummary_{l_i}$ structure for each location l_i that has a current estimated relative-frequency of at least θ. In order to provide a solution in a sliding window where deletions as well as insertions of elements need to be supported, *Location-StreamSummary-Table* also needs to include a sketch structure. This sketch structure is maintained to keep track of frequencies of locations in a sliding window by allowing both insertion and deletion operations [JQS$^+$03]. In general *GeoWatch* uses sketches to keep track of the frequencies of tracked elements. The second component is the *Topic-StreamSummary-Table*, a hash table that monitors the topics that are potentially correlated with at least one location and a sketch structure to keep track of the topic frequencies. For each tracked topic this structure also keeps track of the number of locations the topic is trendy for. Once this value reaches 0, the topic is removed from *Topic-StreamSummary-Table*.

Spatial properties of trends are crucial to capture many trend detection problems. One of the challenges associated with trend detection in general is the large number of possible topics. This problem intensifies in the case where we introduce yet another dimension of space. In the same vein, detecting spatial trends at different temporal granularities poses another challenge in the context of this type of analysis. Indentifying topics that suddenly become popular, i.e., a topic that is not necessarily a heavy-hitter in the traditional sense but exhibits a sharp increase in frequency over a short period of time poses another challenge. Solutions to discover such trends, it is necessary to consider bith the frequency and temporal order of elements in a data stream.

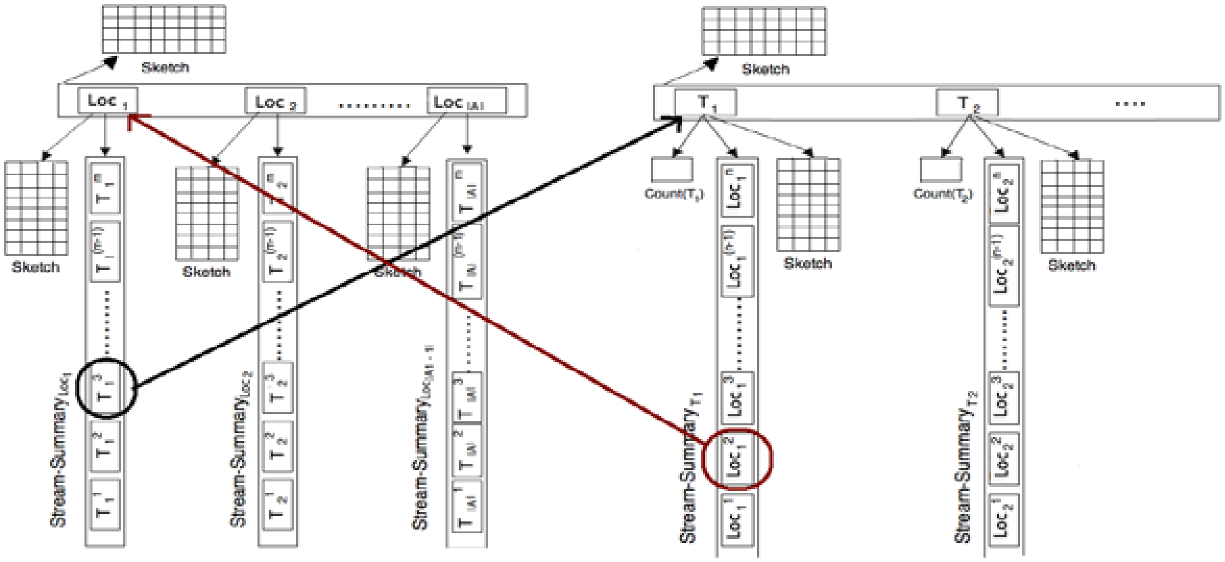

Figure 3: Overview of Data Structure: The two main sub-components are *Location-StreamSummary-Table* (on the left) and *Topic-StreamSummary-Table* (on the right). *Location-StreamSummary-Table* keeps track of ϕ-frequent topics for each of the θ-frequent locations. *Topic-StreamSummary-Table* keeps track of ψ-frequent locations for each topic that is ϕ-frequent for at least one location. Here the third most important topic for Loc_1 is T_2 and the second most important location for T_2 is Loc_1

4 Discussion

Spatial computing is likely to pay an important role in a variety of applications. With our reliance on Navigation systems and maps, next-generation transportation system will greatly benefit from a tighter integration of spatio-temporal data and computing. Similarly, as latency-sensitive applications such as distributed gaming, many-to-many interactive video applications, and synchronous multi-user sessions in online social networks become prevalent, the traditional cloud computing platforms will likely integrate computing, storage, and networking resources at the network edge. In this context, the integration has to be not only be based on network metrics such as latency and bandwidth but also on the basis of spatial metrics such location of the resource. Our preliminary investigation indicates that this poses a formidable research and development challenge.

References

[Ben75] Jon Louis Bentley. Multidimensional binary search trees used for associative searching. *Commun. ACM*, 18(9):509–517, 1975.

[CDG+06] Fay Chang, Jeffrey Dean, Sanjay Ghemawat, Wilson C. Hsieh, Deborah A. Wallach, Mike Burrows, Tushar Chandra, Andrew Fikes, and Robert E. Gruber. Bigtable: A Distributed Storage System for Structured Data. In *OSDI*, pages 205–218, 2006.

[DG04] Jeffrey Dean and Sanjay Ghemawat. MapReduce: simplified data processing on large clusters. In *OSDI*, pages 137–150, 2004.

[FB74] Raphael A. Finkel and Jon Louis Bentley. Quad trees: A data structure for retrieval on composite keys. *Acta Inf.*, 4:1–9, 1974.

[Gut84] Antonin Guttman. R-trees: a dynamic index structure for spatial searching. In *SIGMOD*, pages 47–57, 1984.

[hba10] HBase: Bigtable-like structured storage for Hadoop HDFS, 2010. `http://hadoop.apache.org/hbase/`.

[ind] Indonesia, brazil and venezuela lead global surge in twitter usage. `http://www.comscore.com/Press_Events/Press_Releases/2010/8/Indonesia_Brazil_and_Venezuela_Lead_Global_Surge_in_Twitter_Usage`.

[JQS+03] Cheqing Jin, Weining Qian, Chaofeng Sha, Jeffrey X. Yu, and Aoying Zhou. Dynamically maintaining frequent items over a data stream. In *CIKM '03*, pages 287–294. ACM, 2003.

[Mor66] G. M. Morton. A computer oriented geodetic data base and a new technique in file sequencing. Technical report, IBM Ottawa, Canada, 1966.

[SST+09] Jagan Sankaranarayanan, Hanan Samet, Benjamin E. Teitler, Michael D. Lieberman, and Jon Sperling. Twitterstand: news in tweets. In *GIS '09: Proceedings of the 17th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, pages 42–51, New York, NY, USA, 2009. ACM.

[TLP+08] Benjamin E. Teitler, Michael D. Lieberman, Daniele Panozzo, Jagan Sankaranarayanan, Hanan Samet, and Jon Sperling. Newsstand: a new view on news. In *GIS '08: Proceedings of the 16th ACM SIGSPATIAL international conference on Advances in geographic information systems*, pages 1–10, New York, NY, USA, 2008. ACM.

[wir10] `http://en.wikipedia.org/wiki/List_of_mobile_network_operators`, 2010.

The Austin Project and its Ingress Data Layer

Mohamed Ali
Microsoft Corporation
mali@microsoft.com

There has been a tremendous investment in spatial computing from both academia and industry, yet, from two different angles. While they trust their fellow researchers in academia and rest assured that industry will benefit from the topnotch research, industry professionals are held accountable for helping set up the right stage for research. Industry provides real world challenges for research. Industry builds systems, services and tools to help researchers scale their effort to the appropriate level. Finally, industry shoots for opportunities to productize research.

Spatial computing has been advanced by the state of the art technologies in GPS devices and wireless communications. These technologies enable the real time streaming of location data. Hence, spatial computing would be advanced even more by the investment in software systems that monitor, process and analyze streams of location data. This paper addresses the ongoing effort at Microsoft to build such a system. This paper raises awareness of the Austin project and its Ingress Data Layer within the spatial community. The Austin project provides the system, the services and the tools for researchers so that their effort in advancing the field is magnified.

Background

Microsoft StreamInsight [1, 2] is a platform that supports building Complex Event Processing (CEP) applications. It allows analyzing streaming event data at high throughputs from multiple sources with low latency. Its computations are event-driven, such that results are delivered instantaneously in near-real time.

Ali et al [3] present two approaches (the extensibility approach and the native support approach) to enable spatiotemporal query processing in DSMSs and, more specifically, in StreamInsight. The extensibility approach has been further investigated [4, 5, 6] to extend StreamInsight with the capabilities of the SQL Server Spatial Library [8]. Using the extensibility framework proposed by Ali et al [7], incremental streaming-oriented versions of spatial operators are developed and integrated with the query execution pipeline.

Project Austin

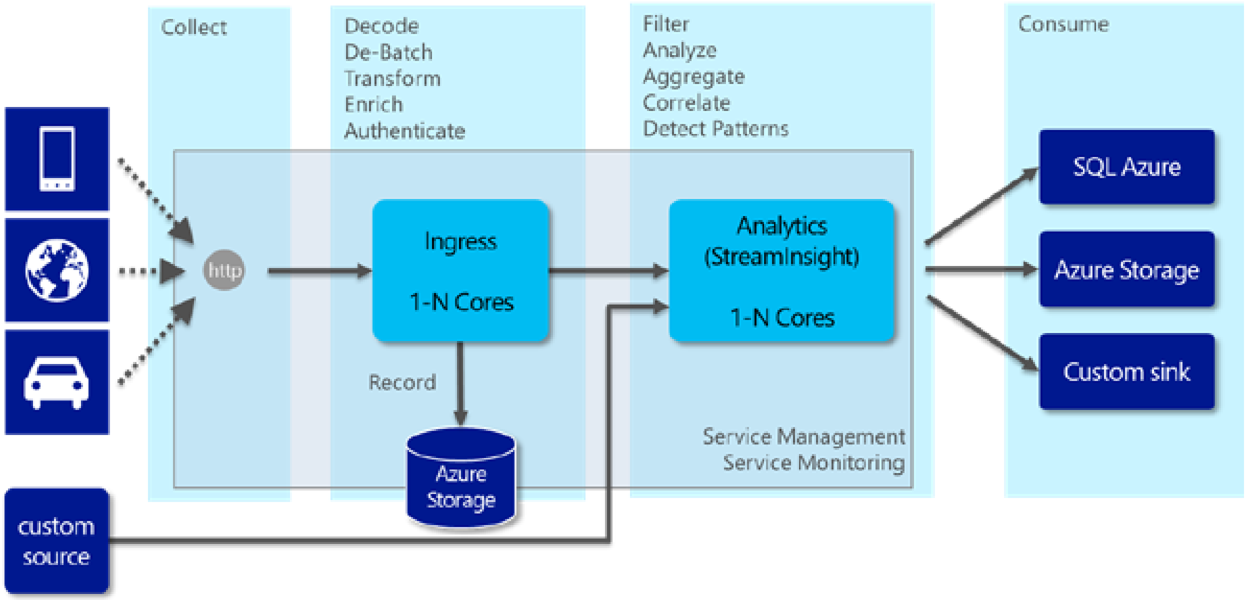

Figure 1. High level architecture of Project Austin and Ingress Data Layer.

Project Codename "Austin" makes Microsoft StreamInsight's complex event processing capabilities available as a service on the Windows Azure Platform (Figure 1). This allows Microsoft's customers and partners to build event-driven applications where the analysis of the events is performed in the Cloud. Such a deployment becomes relevant in scenarios where event data needs to be collected from globally distributed assets or equipment such as connected cars and cell phones. The output of the event processing is also consolidated and made globally available.

Project Austin also addresses the needs of customers to scale out and to deal with distributed systems with occasional loss of connectivity. Customers will be able to monitor, collect telemetry and gain real-time insights from highly distributed assets such as mobile devices, process and control devices & manufacturing robots, etc. Instead of pulling data into an on-premise analytics environment and then possibly distributing it again, it can be processed in the Cloud using StreamInsight's event-driven computation framework, providing cloud computing benefits for many application scenarios. Beyond the core StreamInsight functionality, Austin will provide built-in connectivity to common data sources and sinks (Azure Storage, SQL Azure) with a consistent developer experience, deployment experience, manageability and monitoring experience.

The Ingress Data Layer gives the ability to collect data from a wide variety of client devices and is a key value proposition of Austin. Data ingress functionality in Austin leverages the ubiquity of HTTP connections and can accept any HTTP request (GET, PUT, POST). Data ingress endpoints can also accept and de-serialize OData encoded payloads, based on a .NET type description. The variability of endpoints would enable a wide range of communication protocols (REST, XML, JSON) to co-exist and to interface with Austin. In parallel to the event-processing pipeline, Austin allows for the recording of the raw data stream into Azure blob storage for off-line processing, replay, and data mining.

Project Goals

The high level project goals can be summarized as follows:

Connectivity. Real time data collection from connected devices (Servers, Tablets, Phones) through standards compliant endpoints (REST, XML, JSON) with the ability to enrich and transform data (e.g., geo-tagging).

Analytics. Rich temporal analytical models through dynamic, flexible query and data management experience.

Scale. Multi-tenant cloud service with flexible, elastic capacity and federated scale out collection and analytics.

Security. Secure data transfer through secure endpoints and secure storage of data.

Monitoring. Automatic restarting of Austin instances and analytics upon node failures with notifications, alerting and explorative Troubleshooting.

Tooling. Logging and tracing infrastructure for user code through diagnostic tools and frameworks.

Uptime. Higher availability through redundant instances.

Spooling. Enabling historical data analysis and replay through archives of raw received data in Azure Storage.

References

1. Roger S. Barga, Jonathan Goldstein, Mohamed Ali, and Mingsheng Hong. Consistent Streaming Through Time: A Vision for Event Stream Processing. In Proceedings of CIDR, 412-422, 2007.

2. Jonathan Goldstein, Mingsheng Hong, Mohamed Ali, and Roger Barga. Consistency Sensitive Streaming Operators in CEDR. Tech. Report, MSR-TR-2007-158, Microsoft Research, 2007.

3. Mohamed Ali, Badrish Chandramouli, Balan Sethu Raman, Ed Katibah. Spatio-Temporal Stream Processing in Microsoft StreamInsight. IEEE Data Eng. Bull. 33(2): 69-74 (2010).

4. Jalal Kazemitabar, Ugur Demiryurek, Mohamed Ali, Afsin Akdogan, Cyrus Shahabi, Geospatial Stream Query Processing using Microsoft SQL Server StreamInsight, In VLDB, 2010.

5. Mohamed H. Ali, Badrish Chandramouli, Balan Sethu Raman, Ed Katibah: Real-time spatio-temporal analytics using Microsoft StreamInsight. In ACM GIS 2010.

6. Jeremiah Miller, Miles Raymond, et al, An Extensibility Approach for Spatio-temporal Stream Processing using Microsoft StreamInsight. In SSTD, 2011.

7. Mohamed Ali, Badrish Chandramouli, Jonathan Goldstein, Roman Schindlauer. The Extensibility Framework in Microsoft StreamInsight. In ICDE, 2011.

8. SQL Server Spatial Library, http://www.microsoft.com/sqlserver/2008/en/us/spatial-data.aspx, last accessed in March 2011.

Challenges for Very Large Graphs

Dr. Siva Ravada, Oracle Spatial

Graphs are rapidly becoming the next big challenge in the spatial DB industry. Graphs can be found in simple road navigation applications to finding connectivity patterns in social networks. Graphs have a very simple representation to model very complex relationships. And the general graph tracing and path computation algorithms are very well understood. But the size of today's graphs makes it hard to scale the traditional solutions to new and emerging applications. A graph that is familiar to most Spatial database experts is the road network. Traditionally, the graphs are used in road network navigation are built on a static road network. With the proliferation of mobile devices, various classes of dynamic content like traffic can now be added to the road network. The complexity of graph tracing algorithms increases with the addition of this dynamic content to the road network. In addition to this dynamic content, some road networks can also be augmented with historical traffic patterns which can be used in the navigation systems. This adds the temporal element to the graph tracing algorithms. In the social networks world, the graphs tend to be much larger in size, and with more complexity in terms of connectivity. In this paper, I describe 3 graph problems which are very challenging due to the size and complexity of the graphs.

Multi-Modal Network Tracing: Spatial Graphs are very common in today's mobile and navigation systems. In addition to the complexities described above, there are new requirements from the industry for multi-modal networks where a single logical network can be used to navigate across multiple transportation networks. For example, a user might want to find the best possible way to go from an address in Boston to an address in San Francisco. Such a trip would span at least 2 different types of transportation modes (Car to and from the airport and the flight). If the person wants to take the public transportation to and from the airport, there will be another mode of transportation added to the mix. With mobile devices, sometimes the users want to include foot paths along with some type of public transportation in finding the shortest path. And when you include public transportation or commercial flights, you also need to consider the temporal aspect of these systems. Thus finding a shorted path would need to consider different modes of transportation in addition to the published time tables for these transportation systems. And this gets even more complex if you add the historical time delays and traffic patterns to the network. A route finding system should be able to consume all of these types of input to find the best possible route between two points. With the current systems, there is no standard way to represent all of these types of transportation systems and the corresponding temporal data associated with them. There is a strong requirement to define a data model that can capture all the data in such a system. There is also the requirement to define standards so that different transportation systems can exchange the temporal information like time tables and delays with a route guidance system.

Network Analysis with Map-Reduce: Data is growing very rapidly to increase the sizes of these graphs. Network tracing algorithms combining most of these different types of attribute (as described above) data are available today. But the biggest drawback of most of these existing solutions is that they are all memory based solutions. That is, all the network data has to be loaded into memory before the network tracing algorithms can be executed. There are some solutions where this memory management is handled by the system (Oracle Spatial Network Model) so that only portions of the network are loaded into memory while doing the network analysis. This still poses problems when the networks are large and the latency of loading different parts of the network into memory is much more expensive than the actual network tracing cost. With clustered systems where many machines can be used to solve a problem, the total memory of a system is usually sufficient to hold most of the network in memory. With several nodes in a cluster, different nodes can hold different pieces of the network in memory. And if there are enough nodes, the whole network can be kept in memory. With the advent of Hadoop and other map reduce technologies, these clustered systems can be easily programmed to solve different tasks. But this architecture is not suitable for solving network tracing type problems as many of these algorithms assume the whole network can be accessed in one place. So there is real need to develop map reduce based algorithms for network tracing applications. For example, a network tracing algorithm like A* can be done in such a way that each node in the cluster solves part of the problem and the reduce phase can be used to combine these partial traces to build a complete solution. Some problem like finding all pairs shortest paths, network clusters are more suitable for these map reduce environments. While some others like single source single destination shortest path are not harder to do in map reduce environments. Thus this provides a very challenging opportunity to build new algorithms for solving these complex problems on a very commonly used platform.

Visualization: Network visualization is one area that has been neglected by the industry. All the existing network visualization technologies work very well for small networks (with thousands of nodes and links). But once the networks grow to be millions of nodes and links, there are no solutions available for visualizing these networks. Consider the well known facebook social network map.

Source: http://motherboard.vice.com/2010/12/14/facebook-world-a-map-of-the-social-graph

This is a network with 10 million nodes and many more edges. But this is a static map. Imagine the possibilities if this map can be made interactive so that when you click on a link, you can get other attribute data about each of those links. For an interactive visualization application, this amount of data is not practical. So we first need a way to abstract (aggregate) the representation for this large network. This is similar to the aggregation concepts used in normal maps where the small scale maps show less detail and large scale maps show more detail. We need the ability to create aggregate representations for large (spatial and non-spatial) networks. The typical requirement is to be able to drill down a network when users click on certain links and get a more detailed view of the network. This can apply to both the links and nodes. That is, both links and nodes have to be aggregated to provide a concise representation of the large network. Since these networks tend to be very dynamic, these aggregation methods should be aware of such changes as well. For a given Spatial network, such aggregation methods may be able to use existing Spatial clustering techniques. But for a logical network, this problem is much harder to solve.

Managing Competition in Spatial Computing

Ouri Wolfson,

University of Illinois, Chicago

A lot of work has been devoted to collaboration among distributed computing devices. Sensor networks and crowdsourcing are prominent examples of distributed collaborative computing. Much less attention has been paid to competition among computing devices, particularly competition for space. An example of such competition is drivers, guided by their smartphones, attempting to park; they are competing for a limited number of parking spaces. Similarly, car navigation systems compete for road space, commuters compete for seats on a bus, taxi cabs compete for customers, and customers compete for taxi cabs. In other words, spatial competition abound.

In this paper we use parking to illustrate the challenges of information systems that manage spatial competition. Cruising for parking by driving around an urban area looking for available parking slots has been shown to be a major cause of congestion. For example, studies conducted in 11 major cities revealed that the average time to search for curbside parking was 8.1 minutes and cruising for these parking slots accounted for 30% of the traffic congestion in those cities (see [1]). This means that each parking slot would generate 4,927 vehicle miles traveled (VMT) per year [2]. That number would of course be multiplied by the number of parking slots in the city. For example, in a big urban city like Chicago, with over 35,000 curbside parking slots [5], the total number of VMT becomes 172 million per year due to cruising for parking. Furthermore, this would account for waste of 8.37 million gallons of gasoline and over 129,000 tons of CO_2 emissions.

The proliferation of mobile devices, location-based services and embedded wireless sensors has given rise to applications that seek to improve the efficiency of the transportation system. In particular, new applications to help drivers find parking in urban settings are becoming available. For example, wireless sensors embedded in parking slots are used to detect the availability of slots in some area, and the locations of currently available parking slots are disseminated to the mobile devices of drivers that are looking for parking in the area. A municipality that uses sensors embedded in the streets is San Francisco (see SFPark [3]). When a user is looking for parking in some area of the city, the application shows a map with the marked locations of the open parking slots in the area.

We propose that smartphone apps that navigate drivers to parking slots be developed. Furthermore, they have a focus on a conceptual gap between two notions of parking, which is a spatio-temporal matching between mobile agents (drivers) and spatial resources (parking slots). The two matching notions are *optimality* and *equilibrium.* Ideally, we would like the matching to be optimal, i.e. minimize the total time driven to park by all vehicles. However, achieving this optimality requires a central authority that can eliminate competition by dictating the slot in which a driver should park, even if that driver can do better. Figure 1 below illustrates this point by an example.

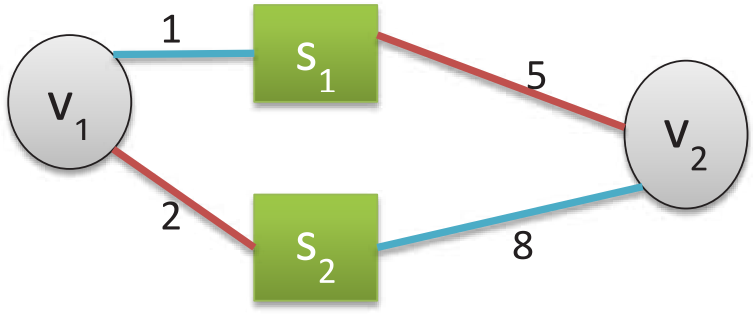

Figure 1: A parking example with two vehicles and two parking slots

Suppose that the edge labels represent travel times in minutes, i.e., vehicle v_1 is 1 minute away from slot s_1, 2 minutes away from s_2 and so forth. To achieve minimum total driving time, v_1 will have to park in s_2, and v_2 will have to park in s_1. This parking configuration, called minimum, has a total time of 7 minutes. However, this requires v_1 to drive to a farther slot, s_2, i.e. an inferior slot from her point of view because s_1 is closer. There is no central authority that can dictate this parking configuration to v_1. If v_1 drives to s_1 (and captures it since she is closer than v_2), then v_2 must settle for s_2. The total driving time of this configuration, called equilibrium, is 9 minutes, i.e., higher than the minimum. However, in an equilibrium configuration no driver *d* can unilaterally deviate and improve *d*'s cost. In other words, if there is no competition then v_2's parking app should navigate her to s_1, the closest slot; whereas if there is competition from v_1, the app should navigate her to s_2. This means that the parking app P should navigate differently in the face of competition for spatial resources. Furthermore, P usually does not know whether there is competition, and where the competing devices are located. In [4] we introduced a method of competitive parking, Gravitational Parking, that is provably superior to the one that simply goes to the closest slot, i.e. ignores competition. The superiority reaches 30%. This means that the price of using a noncompetitive algorithm in competitive situations in Chicago would be 25 million Vehicle/miles traveled per year, over 1.2 million gallons of wasted gasoline, and over 19,000 tons of CO_2 emissions.

Gravitational Parking works by having a vehicle be “attracted” to available parking slots, i.e., slots applying a “gravitational force” on vehicles. Each force is represented by a vector, and the vehicle moves in the direction of the vector-sum of the forces. This means that the vehicle *v* does not always pursue an available parking slot because the slot may become unavailable by the time *v* reaches it. Instead, it moves in a general direction that is promising in the sense that contains multiple available slots. Figure 2 shows the gravitational force field generated by five available slots. The arrows represent the direction in which a vehicle will move when it is located at the start point of the arrow, and the small dots represent the slots. This diagram indicates how vehicles move across the map when using Gravitational Parking, and it shows that they will eventually converge to a slot.

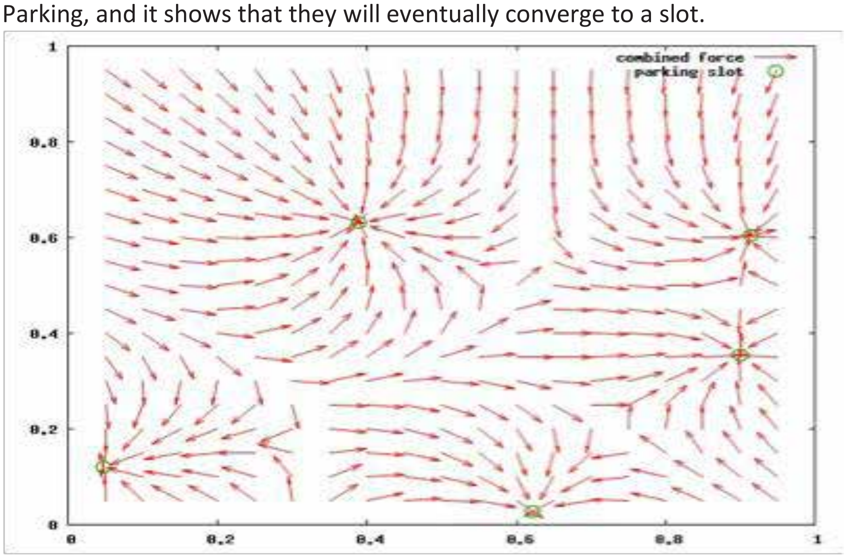

Figure 2: Field force generated by 5 slots

How should Gravitational Parking be extended to account for competition in routing? In other words, the shortest-distance path to a destination is inferior if it is congested. And a shortest-time path will also be inferior if many drivers are led to it by their car navigation systems. This situation becomes increasingly likely as car navigation systems with traffic information proliferate, leading to the phenomenon called herding. Similarly, what is the equivalent of Gravitational Parking in competing for other spatial resources? To address these questions we propose competition management as research direction for the emerging spatial computing community.

References:

[1] D. Shoup. The High Cost of Free Parking. American Planning Association, Chicago, IL, 2005.

[2] D. Shoup. Cruising for parking. Transport Policy, 13: 479-486, 2006.

[3] http://sfpark.org/

[4] D. Ayala, O. Wolfson, B. Xu, B. Dasgupta, and J. Lin. Parking slot assignment games. In Proc. of the 19th Intl. Conf. on Advances in Geographic Information Systems (ACM SIGSPATIAL GIS 2011), Chicago, IL, November 2011

[5] Transportation Alternatives (www.transalt.org), New York, NY. Pricing the Curb: How San Francisco, Chicago and Washington D.C. are reducing traffic with innovative curbside parking policy, July 2008.

CROSS-CUTTING SPATIAL COMPUTING

Challenges of Spatiotemporal Data Fusion

Dr. Sara J. Graves
Director of the Information Technology and Systems Center and Professor of Computer Science
University of Alabama in Huntsville

Introduction

Fusion of spatiotemporal data is a difficult process that is often required for many of today's data-intensive spatial computing applications [1]. While much work has been done in this area, the problem has not yet been fully addressed. In fact, given the continuing proliferation of geospatial data, data fusion is both more complex and more urgent. For example, data collection and analysis workflows can feed into each other, with each stage of the process generating data products which must be managed, converted to different formats, and fused together for the next stage. Figure 1 shows an example of a cyclic, data-intensive workflow framework for collecting, processing, exploiting, and disseminating spatiotemporal data. Opportunities for data fusion should be explored at each of these stages.

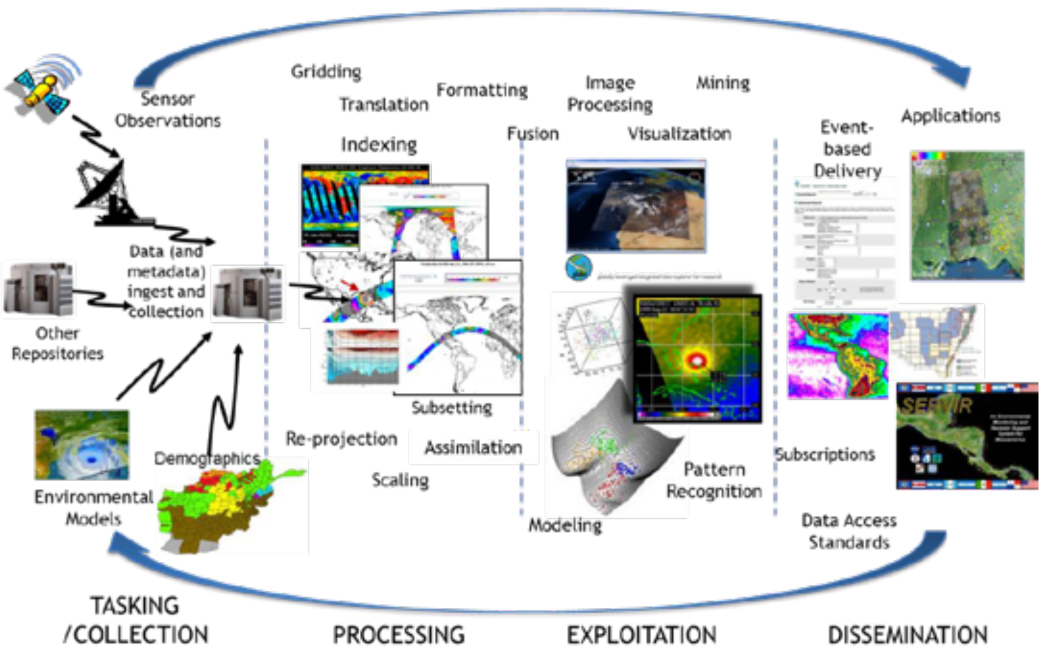

Figure 1. Example of a cyclic, data-intensive spatiotemporal application workflow framework

Data Fusion Approaches and Challenges

There is no single, best approach to spatiotemporal data fusion that will solve all spatial computing problems. A key preliminary step is specifying and gathering appropriate data to be fused. Fusion can be accomplished directly by combining data products or indirectly by using them as a part of a knowledge discovery process. In order to illustrate the varied nature of data fusion and the inherent difficulties and requirements, we discuss several data fusion problems, along with approaches implemented at the University of Alabama in Huntsville's Information Technology and Systems Center (ITSC), for projects funded by NSF, NASA, NOAA, DoD, DoE, and other agencies.

Preliminary Data Fusion

A preliminary aspect of data fusion is the process of gathering multiple observations or models of the same geophysical area or spatiotemporal phenomenon. To streamline the timely delivery of targeted data to users and processes, ITSC has prototyped an Event-Driven Data Delivery (ED3) system that provides users the ability to subscribe for the delivery of specific data products, triggered by the occurrence of a specific event. For instance, a weather researcher might want to subscribe for the automatic delivery of several observational datasets based on the occurrence of a tropical storm in the Gulf of Mexico. ED3 is able to "listen" for occurrences of matching events based on advisories issued by the National Hurricane Center and then initiate the requested data retrievals, processing, packaging and notification of the user that the data is ready for pickup. This approach is applicable to any events with spatial and temporal characteristics and can greatly shorten the time involved with accessing and utilizing data for analysis and decision making. ED3 performs the first stage in spatiotemporal data fusion, acquiring and staging temporal and spatially matched data that can be used for further analysis.

Challenges in this area include:

- Further automating the acquisition of relevant data,
- More intelligent and precise data selection,
- More widely available preprocessing capabilities to subset, subsample, regrid, and/or reformat data in preparation for subsequent fusion and analysis.

Direct Data Fusion

Data can be combined or fused at a variety of levels, from raw observations to features within the data, to decisions inferred from the data. Raw sensor data level fusion requires the sensors to be observing the same phenomena. These fusion techniques involve statistical methods such as principal component analysis, band ratios within spectral data, etc. Feature level fusion involves extraction of representative features from the sensor data. These features, extracted from multiple sensor observations, are combined into a single feature vector, which is then input into pattern recognition tools such as classifiers. Decision level fusion involves fusion of sensor information after each sensor has made a preliminary determination of a phenomenon's identity, location and attributes. The best approach to take depends on the properties of the data sources and the application. A set of broad capabilities is required in order to accommodate the varied nature of the data and applications.

GLIDER [2] was developed as an integrated geospatial image display and analysis framework with data mining [3] and workflow capabilities, and has been used to support research for NASA, DoD, and DoE. As part of GLIDER, ITSC has developed data fusion approaches for geospatial imagery that allow co-location of imagery of different projections or geographic grids. Imagery from different physical sources can be fused together as either a common grid or a common native sensor view. This approach allows for one image to be transformed, matching the other without transforming both images into a standard projection model, minimizing resolution degradation and preserving the original swath view of the data. The software also supports gridding imagery to a common projection. Combined with ED3, GLIDER allows an analyst to import, manage, process, and analyze data from a variety of sources.

Challenges to direct data fusion include:

- Proliferation of data formats often requires custom software interfaces and makes data import and fusion difficult
- Varying sensor properties (orbital characteristics, spatial resolution, projections, etc.) require sophisticated software for spatial reprojection of data
- Reprojection is computationally expensive and can result in loss of spectral information

Indirect Data Fusion

In some cases, there may be observations of a particular phenomenon that are fundamentally incompatible with one another, making it impossible to directly combine the observations. In other cases combining the observations may simply be undesirable due to large differences in spatial or temporal resolution. A model-based approach to sensor fusion may be appropriate in these cases. In this approach, a model is constructed for the objects or phenomena of interest that describe the possible states of those phenomena. Models are also constructed for each type of sensor to reflect the expected range of values for each sensor for given object states. For example, the maximum likelihood model works by exploring the space of possible object states and finding the one that is the most likely given the set of observations. The observations from the highest fidelity sensors are naturally weighted more highly than those of low fidelity sensors. General purpose search techniques such as stochastic hill climbing, simulated annealing and genetic algorithms are used to generate candidate object states. Each state generated by the search procedure is then evaluated using the state and sensor models to see how likely it is given the actual observations.

An example of the model based maximum likelihood approach has been used in the Vantage tool [4] for tracking of aircraft. Vantage contains models for various classes of sensors. Each of the sensor classes has different properties, which makes direct fusion difficult. All sensor models have probabilistic detection based on the range to the target and its properties such as radar cross-section or infrared signature, as appropriate to the sensor. Using the maximum likelihood model approach to data fusion, the Vantage system is able to track aircraft in real time.

Challenges to indirect data fusion include:

- Indirect fusion requires integration of powerful visualization and analysis software
- Algorithmic analysis is problem specific so a method for encapsulating analysis workflows is needed
- Visual analysis tools are needed to evaluate and refine results

Progress toward a Data Fusion Framework

In an approach that combines direct data fusion with decision level or indirect fusion, ITSC is investigating techniques for correlating and fusing data from a number of disparate data sources in order to provide improved information on characteristics of phenomena. Examples of the types of data that could be used include are infrasound data, overhead remotely sensed measurements such as environmental data, weather satellite imagery, infrared observations, and telemetry. The technical objective of this research is to demonstrate an analysis framework and techniques for integrating these data sources, resulting in an improved understanding of events of interest. Spatiotemporal imager data of different spatial resolutions must be fused together and analyzed in conjunction with vector-based data in a 3D geospatial domain. This analysis requires not only fusion, but also 3D visualization and analysis capabilities. The Globally Leveraged Integrated Data Explorer for Research (GLIDER) is being used as the visualization, fusion, and analysis framework for this study.

Conclusions

Data fusion is a complex and difficult problem that has not been fully addressed. Both direct and indirect data fusion approaches are needed to solve specific problems, depending upon the properties of the data sources and the application. GLIDER and other tools developed at ITSC represent progress toward an integrated framework for fusion, visualization, and analysis of data products from a variety of remote sensing sources, but much work remains. ITSC's involvement in the NSF EarthCube initiative informs our work in data fusion and other technologies for data-intensive spatial computing. For example, ITSC researchers played a lead roles in the EarthCube Data Discover, Mining and Access (DDMA) community workgroup [5] to create a roadmap, through community consensus, for current and future data needs within the geosciences. Data fusion will likewise be a key set of technologies for NSF's new DataWay initiative.

References

[1] Graves, S. and T. Berendes, "Meeting the Challenges of Data-Intensive Science," *IGARSS 2012: IEEE International Geoscience and Remote Sensing Symposium*, 22-27 July 2012.

[2] Ramachandran, R., S. Graves, T. Berendes, M. Maskey, C. Chidambaram, S. Christopher, P. Hogan, T. Gaskins, "GLIDER: A comprehensive software tool to visualize, analyze and mine satellite imagery," *IGARSS 2009: IEEE International Geoscience and Remote Sensing Symposium*, vol.3, no., pp.III-781-III-784, 12-17 July 2009.

[3] Rushing, J., R. Ramachandran, U. Nair, S. Graves, R. Welch, H. Lin, "ADaM: A Data Mining Toolkit for Scientists and Engineers," *Journal of Computers and Geosciences*, Vol. 31, Issue 5, pages 607-618 (June 2005).

[4] Lin, H., J. Rushing, S. Graves, E. Criswell, "A data fusion algorithm for large heterogeneous sensor networks," *International Conference on Wireless Algorithms, Systems and Applications*, Chicago, 1-3 August, 2007.

[5] The Earthcube Data Discovery Mining and Access Group, "The DDMA EarthCube RoadMap," 2012. Available from: https://docs.google.com/document/d/1RtR8A0h0HFhZP9kIPFlfCChFIYLIf-G-TZMbNSzw58M/edit.

A Study Program in Urban Computing that Leverages Machine Learning and Social Media

N. Sadeh, J. Cranshaw, J. Hong, R. Schwartz

School of Computer Science

Carnegie Mellon University

sadeh@cs.cmu.edu

Abstract

With the rapid proliferation of smart-phones in recent years and the subsequent rise of location-based applications producing and consuming geospatial data about the world around us, there is great potential for the development of computational methods that can harness this data to solve a plethora of problems, potentially opening a whole new range of uses and applications. Of particular interest is the application of this data towards the goal of better understanding the intricacies of the urban landscape. In this work, we layout a program of study in *urban computing*, which uses new modalities of ubiquitous and social computing systems and machine learning to help us better understand the urban environment, improve the efficiency of the complex systems that operate it, and enhance the well-being of those that live there. To illustrate the potential opportunities of this field, we discuss the Livehoods project, which uses geospatial check-in data to study the neighborhood structure of a city.

Introduction

Large and dense cities are often characterized by their intricacy and by the vast array of interdependent complex systems that enable them. Municipal systems for transportation, safety, public health, housing, waste removal, and energy all interact with each other and with the equally complex private sector economy to meet the everyday demands of the millions of inhabitants of a large city. At the same time, the forces that shape the movements, preferences, and demands of the populous are equally complex; cultural perceptions, economic factors, municipal borders, demography, geography, and resources—all shape and constrain how people make use of the resources of the city. Understanding how these opposing dynamic systems—the needs of the citizens, and the resources of the city—relate and interact with one another is critical to improving the efficiency of the city and the overall well being of those that live in it. To that end, although methodologies for accurately studying the city are essential to improving it, it can be extremely difficult to convey the intricate social realities of a city to an outsider. When outsiders, such as researchers, journalists, or city planners, want to learn about a city, it can require hundreds of hours of observation and interviews. Although existing methods offer a way to gather deep insights about certain aspects of city life, they can never offer large-scale insights or produce real-time analyses of the inner workings of the city.

To illustrate this challenge we offer two examples demonstrating the difficulty of collecting and analyzing such data. Our first example is from the perspective of transportation engineers, who study how and when people travel [1] to optimize traffic flows and public transportation. This information is typically gathered through a Travel Behavior Inventory (TBI) of diaries from thousands of people, and GPS from a few hundred over a few days. However, given the cost and effort of data collection, TBIs are only done once every ten or twenty years [3]. Our second example is quality of life surveys undertaken in various disciplines such as sociology, urban studies, architecture, and by city governments. These surveys are often deployed at the neighborhood level and require research assistants to go door-to-door interviewing residents about issues related to their daily life, e.g. access to fresh produce. Although this methodology can offer deep insight into the socio-economic demands on a region, because of the costs associated, they can only ever be conducted on a select few neighborhoods and over infrequent time spans.

The widespread adoption of smartphones and social networks opens the door for a new era of **urban computing**, where large and dynamic collections of public data makes it possible to ask a wider range of questions and develop a deeper understanding of urban activity than ever before. It also calls for the development of new algorithms and tools to help diverse groups of individuals and organizations better understand the urban ecosystem, and find new ways to anticipate and respond to the needs of complex and ever-changing urban communities. More specifically, **we argue that there is an exciting opportunity for creating new ways to conceptualize and visualize the dynamics, structure, and character of a city by analyzing the social media its residents generate.** Leveraging such rich sources of data with new and expressive techniques in spatial modeling and machine learning will yield new kinds of analytics tools that will let urban planners, policy analysts, social scientists, and computer scientists explore how people actually use a city, in a manner that is relatively cheap, highly scalable, and insightful. These tools would shed light onto the factors that come together to shape the urban landscape and the social texture of city life, including municipal borders, demographics, economic development, resources, geography, and planning.

Figure 1. The Livehoods project (www.livehoods.org) allows users to visualize neighborhoods identified by clustering public Foursquare check-ins based on measures of social and geographic distance. The product of this analysis has been shown to often coincide with the mental maps of local residents and to offer a highly scalable way of organically identifying a city's neighborhoods.

The Livehoods Project

The Livehoods Project is an urban computing tool that we have built and deployed that exemplifies the case for how social and ubiquitous computing technologies can be used to better understand the urban environment [2] (see Figure 1). The problem that we've focused on with Livehoods is studying *the idea of a neighborhood*. There are at least two perspectives one can take in thinking of city neighborhoods. First, you can think of neighborhoods as the area enclosed by fixed boundaries set by city and local governments. These fixed, municipal, historical boundaries are essential to providing order to the chaos of the city, serving as a common frame of reference for the various people and organizations that operate within a city. However, the city is always changing, and neighborhoods evolve much more rapidly than the boundaries that delineate them do. Often these fixed city boundaries do not necessarily reflect the cultural perceptions people have of the area. Our research hypothesis is that the character of an urban area is defined not just by the types of places found there, but also by the people that make it part of their daily life. To explore this idea, we crawled 18m check-ins from the location-based social network foursquare, and applied clustering algorithms that grouped nearby venues into areas (which we call Livehoods) based on geographic distance and the particular mix of the people who check-in to them. To evaluate our work, we conducted numerous interviews with locals, including city planners, business owners, and local residents. We asked them to draw parts of the city that they were most familiar with (before they saw our maps), describe characteristics of the areas they were most familiar with, and examine the Livehoods that our system generated and offer feedback. In many cases, our Livehoods matched the mental models of locals, and also provided new insights about how neighborhoods were organized and were changing over time. Livehoods demonstrated to us the strong demand for urban analytics, as well as the power of combining data mining with user-centered design to develop tools for diverse stakeholders. Our long term goal is to extent the various spatial computing algorithms currently being used in the construction of Livehoods, so that we can build a suite of tools that might be useful to the various stakeholders of the city.

Conclusions

In this short article, we made the case for a new form of spatial computing that leverages social media data and machine learning techniques to overlay socio-economic data on top of traditional geographical data. This additional layer of information opens the door to a much richer view of spatial information. In particular, in an urban context, social media data allows for the dynamic generation of models that combine both geographic information and social information and provide dynamic snapshots of activities in a city, ways in which resources are being used at different points in time, and more. Combined with historical data, this type of information can be used to spot new trends (e.g. changes in patterns of activities, spread of contagious diseases [4], and much more) as well as identify one-off events that would otherwise elude identification. Work in this area will lead to new algorithms and metrics for processing and refining large sets of social media data that will give insights into various characteristics of a city and its inhabitants. It will also require new algorithms and architectures for managing privacy issues, to minimize the potential for re-identifying individuals based on the filtered and aggregated data that is made available by our tools, and also new methodologies for evaluating and interpreting algorithmic results.

References

[1] F. Calabrese, G.D. Lorenzo, and C. Ratti Human Mobility Prediction based on Indvidual and Collective Geographical Preferences. IEEE International Conference on ITS, (2010).

[2] J. Cranshaw, R. Schwartz, J. I. Hong, and N. Sadeh. The Livehoods Project: Utilizing Social Media to Understand the Dynamics of a City. The Sixth International AAAI Conference on Weblogs an Social Media (ICWSM-12). Dublin, Ireland, June 2012.

[3] Metropolitan Council. 2000 Travel Behavior Inventory: Home Interview Survey: Data and Methodology. 2000.

[4] A. Sadilek, H. Kautz, and V. Silenzio. Modeling Spread of Disease from Social Interactions. The Sixth International AAAI Conference on Weblogs an Social Media (ICWSM-12). Dublin, Ireland, June 2012.

Spatial Computing – Challenges and Opportunities

Johannes Schöning [1,2]
[1] Hasselt University, Belgium
[2] University College London, UK

As stated in the workshop proposal spatial computing is a set of ideas and technologies that will transform our lives by understanding the physical world, knowing and communicating our relation to places in that world, and navigating through those places.

One important "kind of places" are cities. Cities are places where people, meet, exchange and interact. They bring people with different interests, experiences and knowledge close together. They are the centres of culture, economic development and social change. They offer many opportunities to continually innovate with technologies, from the infrastructures that underlie the sewers to computing in the cloud [1]. According to a United Nation report [2] every second the global urban population grows by 2 people. Therefore the urban population is expected to increase from 3.6 billion people in 2011 to 6.3 billion in 2050. In 2020 more than 700 cities will exists with populations of +1million; today we have just 500 cities with populations of +1million. The exploding urban population growth creates unprecedented challenges, among which provision for water and sanitation are the most pressing and painfully felt when lacking [2]. This opens a wide space for spatial computing within future cities. In this abstract I focus on the research questions that arise for the area of spatial computing within future mega cities.

From my point of view one overarching goal of my spatial computing in the near future is to integrate the technological, economic and social needs of cities in ways that are sustainable and human-centered. Thereby it provides fundamental enablers for all technologies that will be targeted towards increasing the quality of living and lowering the barriers for mobility in our future cities.

My own main research interests lie at the intersection of the foundations of human-computer interaction (HCI) [3], geographic information science (GI-Science), and novel interface technologies (e.g. hardware such as stereoscopic displays and depth cameras). To address the challenges in this space, I aim to contribute novel HCI- and geography-oriented adaptations of methods from ubiquitous computing [4,5] to improve the state of the art in the area of intelligent user interfaces (IUI).

In the short position paper I would like to highlight four upcoming research topics the from a human-centred perspective that have growing importance within my research area:

- **New User Interfaces for Spatial Information** [6]: Assuming highly developed technical urban infrastructures there is a need of multimodal user assistance [cite multimodal] and support within future mega cities. This includes research on multimodal navigation systems, as well as on novel augmented reality technologies [cite AR paper] to improve human's daily experiences and activities in future cities.

- **Big Data:** Also this topic is covered within other computer science disciplines it will be very important for future cities. In them tons of data will be produced every second. To handle this data stream, highly developed infrastructure will be needed. In addition efficient and fast algorithms and data processing are needed to generate valuable information out of this data [7].

- **User generated geo data and content:** In addition to the data, which is coming from the technical infrastructure from a city user generated geo data and content will have a growing importance. New ways are needed how to incorporate this information into current systems and how to prove and constantly check the quality and consistence of this information [8,9,10].

- **Supporting Sustainable behaviour through IT systems:** Overall all approaches should help citizens adapt their behaviour in order make cities more sustainable involves increasing their awareness of how they live and then encourage changing habits, at an individual, family, local community and city level. This requires adopting a broad unit (ranging from the individual to communities at large) of analysis: considering the needs of the individual city dweller, families, whole neighbourhoods, councils, and communities at large [11,12,13].

References

1. Schöning, Rogers et al.: Intel Collaborative Research Institute - Sustainable Connected Cities, AmI 2012, to be appear.
2. Department of Economic: The Millennium Development Goals Report 2010. United Nations Publications (2010)
3. Preece, J., Rogers, Y., Sharp, H.: Interaction design: beyond human-computer interaction. Wiley (2002)
4. Weiser, M.: The computer for the 21st century. Scientific American 265(3) (1991) 94–104
5. Bell, G., Dourish, P.: Yesterdays tomorrows: notes on ubiquitous computings dominant vision. Personal and Ubiquitous Computing 11(2) (2007) 133–143
6. Johannes Schöning: Advanced User Interfaces for Spatial Information. PhD Thesis Saarland University advised by Prof. Dr. Antonio Krüger, Dr. Michael Rohs and Prof. Dr. Dr. h.c. mult. Wahlster (2010)
7. Juha Laurila et al. The Mobile Data Challenge: Big Data for Mobile Computing Research. Proc. Mobile Data Challenge by Nokia Workshop, in conjunction with Int. Conf. on Pervasive Computing, Newcastle, 2012
8. Matthias Böhmer, Brent Hecht, Johannes Schöning, Antonio Krüger & Gernot Bauer: Falling Asleep with Angry Birds, Facebook and Kindle - A Large Scale Study on Mobile Application Usage. Mobile HCI 2011: International Conference on Human-Computer Interaction with Mobile Devices and Service, (2011)
9. Hecht, B. and Moxley, E. (2009). Terabytes of Tobler: Evaluating the First Law of Geography in a Massive, Domain-Neutral Representation of World Knowledge. *Proceedings of the International Conference on Spatial Information Theory* (COSIT 2009), pp. 88-105. Berlin: Springer-Verlag
10. Afra Mashhadi et al.: On the Sustainability of Urban Crowd-Sourcing for Maintaining Large-Scale Geospatial Databases, Wikisym 2012.

11. Rogers, Y., Hazlewood, W., Marshall, P., Dalton, N. S. and Hertrich, S. (2010) Ambient Influence: Can Twinkly Lights Lure and Abstract Representations Trigger Behavioral Change? Ubicomp 2010 Proceedings. ACM, 261-270.
12. Arroyo, E., Bonanni, L., Selker, T.: Waterbot: exploring feedback and persuasive techniques at the sink. In: Proceedings of the SIGCHI conference on Human factors in computing systems, ACM (2005) 631–639
13. Bird, J., Rogers, Y.: The pulse of tidy street: Measuring and publicly display- ing domestic electricity consumption. In: Workshop on Energy Awareness and Conservation through Pervasive Applications (Pervasive 2010). (2010)

About the author: Prof. Johanns Schöning and the Expertise Centre for Digital Media (EDM), Hasselt Belgium

The Expertise Centre for Digital Media performs research in computer science in two main areas: Visual Computing and Human-Computer Interaction. The Human-Computer Interaction (HCI) research group was established in 2000, primarily building upon HCI research in virtual environments and gradually extending its research subjects over the past years. The group performs research in three sub-domains: multimodal interaction, emerging interactive systems and user-cantered software engineering. Results of this research are applied in diverse domains including (but not limited to): rehabilitation & robotic systems, culture and education, gaming, e-health and user interface design.
Prof. Dr. Johanns Schöning works within the EDM (HCI group) and he is an adjunct consultant at UCL London within the Intel Collaborative Research Institute for Sustainable Cities [cite AMI]. His research interests relate to new methods and interfaces to navigate through spatial information, such as mobile augmented reality applications and home grown multi-touch surfaces. He is involved in research on mobile apps usage (i.e. understanding how context influences people's interactions with mobile apps and how people interact with specific types of mobile apps such as shopping support apps), his most recent publications in this area including:

More information can be found under:
www.johannesschoening.de
www.uhasselt.be/edm

Qualitative representation and reasoning in the context of big data
(Position Paper for the Spatial Computing Visioning Workshop, Washington DC, 2012)

Mike Worboys
School of Computing and Information Science,
University of Maine, Orono, ME 04473

This position paper proposes qualitative spatio-temporal representation and reasoning as a solution to some of the problems posed by the volume, diversity, and dynamism of big data sets. It has long been recognized that for many human activities, qualitative representations provide both concise and intuitive means of moving up the food chain from data to knowledge to action. For example, Beck's 1933 map of the London underground transport system was a topological (and thus qualitative) distillation of salient information on connectivity that travellers need, abstracted away from voluminous and irrelevant metric and Euclidean information.

To be clear, by a *qualitative representation* we mean a representation or model of a domain that is primarily non-numerical. This debars representations based on Euclidean space, reals, rationals, and large integers. Included are representations based upon set theory, finite topologies, graph theory, and natural language. In general terms, a qualitative representation can provide a concise, relevant, and intuitive representation of big data that can be used as a basis for decisions and maybe more precise quantitative analysis. To be useful, there needs to be an alignment s between a domain's qualitative representation and our cognitive models of it. Furthermore, qualitative representations need to be complemented with appropriate reasoning and computational tools. The topic of *qualitative reasoning* is a well-established branch of artificial intelligence, and its subdomain of qualitative spatial reasoning is part of the GIScience research agenda. Approaches to qualitative representation and reasoning are usually based upon formal languages that may be founded on ontologies.

It is now widely accepted that the domain of spatial computing embraces information related to dynamic geospatial phenomena, and thus requires treatment of temporality and process. Furthermore, such phenomena are situated in fully three (and not just 2.5) spatial dimensions of space. With integration of diverse data sets being one of the preoccupations of big data research, we should also include a principled treatment of uncertainty. Thus, the vision of qualitative representation and reasoning needs expansion to cater to current and emerging requirements.

When these extensions of qualitative representation and reasoning are considered in the context of big data, they lead to a rich research agenda. The overarching item for research is the development of expressive and scalable languages and reasoning mechanisms that target spatio-temporal phenomena. A prerequisite here is the principled treatment of spatial change, including 3-D spatial change, in the presence of uncertainty.

While the development of such scalable approaches provides the core of this research area, there will be many more specific issues. Examples of such more specific items for the agenda are:

1. Investigations of similarities and differences between qualitative representations of indoor and outdoor spaces, vehicular and pedestrian spaces, urban and rural spaces.
2. Ensuring that qualitative approaches to representation and reasoning within specific domains accord with our cognitive models of said domains.
3. Surveying the great body of existing qualitative representation and reasoning (including QSR) and assessing how it can be extended to 3-D and dynamic spaces.
4. Application of process-algebraic methods, used in computer science, to the phenomena ranges above.
5. Investigation and application of qualitative uncertainty, including treatments of vagueness and levels of granularity, in this context.
6. Qualitative representations of dynamic spatial fields (two and three dimensional).

Some background

A synopsis of the principal elements of qualitative reasoning may be found in [2]. Spatial aspects of qualitative representation and reasoning are surveyed in [3]. There have been some attempts to integrate time with space (see, e.g. [4]). Issues that arise with extremely large and diverse data sets are surveyed in [1]. Early work on a qualitative representation of dynamic fields is reported in [5].

References

1. Agrawal, D., Bernstein, P., Bertino, E., Davisdon, S. Dayal, U. *et al.* Challenges and opportunities with big data 2011-1. Purdue e-Pubs, Cyber Center Technical Reports. Paper 1. http://docs.lib.purdue.edu/cctech/1 2011.
2. Bredeweg, B., Struss, P. Current topics in qualitative reasoning. *AI Magazine* 24 (4), 13-16, 2003.
3. Cohn, A., Renz, J. Qualitative spatial representation and reasoning. In F. van Harmelen, V. Lifschitz, (editors) *Handbook of Knowledge Representation*, Springer, 581-584, 2008.

4. Gerevini, A., Nebel, B.: Qualitative Spatio-Temporal Reasoning with RCC-8 and Allen's Interval Calculus: Computational Complexity. *ECAI 2002*. IOS Press, 312-316, 2002.
5. Duckham, M., Stell, J., Vasardani , M., and Worboys, M. Qualitative change to 3-valued regions. *International Conference on Geographic Information Science*, Zurich, Switzerland. In S. Fabrikant et al., editors, *Lecture Notes in Computer Science* 6296, pp. 249-263, Springer, Berlin, 2010.