Sense The Physical, Walkthrough The Virtual, Manage The Co (existing) Spaces: A Database Perspective

Beng Chin Ooi   Kian Lee Tan   Anthony Tung
Department of Computer Science
National University of Singapore
Computing 1, 13 Computing Drive, Singapore 117417
{ooibc, tankl, atung}@comp.nus.edu.sg

ABSTRACT

In a co-space environment, the physical space and the virtual space co-exist, and interact simultaneously. While the physical space is virtually enhanced with information, the virtual space is continuously refreshed with real-time, real-world information. To allow users to process and manipulate information seamlessly between the real and digital spaces, novel technologies must be developed. These include smart interfaces, new augmented realities, efficient storage and data management and dissemination techniques. In this paper, we first discuss some promising co-space applications. These applications offer experiences and opportunities that neither of the spaces can realize on its own. We then argue that the database community has much to offer to this field. Finally, we present several challenges that we, as a community, can contribute towards managing the co-space.

1. INTRODUCTION

Traditionally, the physical space and the virtual space are disjoint and distinct. Users in each space operate within the scope of the space, i.e, they may communicate among themselves but do not cross the boundary to the other space. However, technological advances in ubiquitous computing, smart interfaces and new augmented realities have made it possible for these two spaces to co-exist within a single space, the co (existing) space.

In a co-space environment (or cyber-physical system), the physical space and the virtual space interact simultaneously in real-time. Locations and events in the physical world are captured through the use of large number of sensors and mobile devices, and may be materialized within a virtual world. Correspondingly, certain actions or events within the virtual domain can affect the physical world (e.g. shopping or product promotion and experiential computer gaming). Thus, on one hand, the physical space is virtually enhanced with information. On the other hand, the virtual space is continuously refreshed with real-time, real-world information. Figure 1 shows the information flow within a co-space environment - data may flow within a single space, but more importantly, data also flows into the other space. It is this that distinguishes co-space from mixed reality (or augmented reality or augmented virtuality) [28] - while mixed reality integrates the real and virtual worlds (e.g., augmenting live video imagery with computer generated graphics), it is done in a rigid and static manner, and does not capture real-time changes and their effects on either of the spaces.

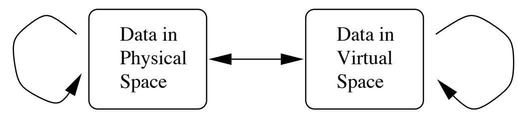

Figure 1: Data flow within co-space as a result of simultaneous interaction.

In co-space, we can design innovative applications that provide experiences and opportunities that neither the physical nor the virtual spaces alone can offer. Some example applications include partnership in shopping among online and physical shoppers, an enhanced digital model that captures physical troop movement, location based games and social networking.

Within such a context, it is easy to see that large amount of data and information must flow to/from co-space in order to ensure that the real and virtual worlds are synchronized. This brings new challenges such as a need to process heterogeneous data streams in order to materialize real world events in the virtual world and more intelligent processing to send interesting events in the co-space to someone in the physical world.

To allow users to process and manipulate information seamlessly between the real and digital spaces, novel technologies must be developed. These include smart interfaces, new augmented realities, efficient storage and data management and dissemination techniques. In this paper, we first present a sample of promising co-space applications. Given that these applications are data-driven, and the potential size of the data that could be generated is enormous, we believe that the database community has much to offer to drive the growth of this field. Finally, we identify and present several challenges that we, as a community, can contribute to manage the large amount of data, the huge number of events and the massive number of concurrent users within co-space. These include the development of efficient storage and indexing methods, processing engines, parallel and distributed architectures.

2. CO-SPACE SCENARIO APPLICATIONS

The co-existence of the physical and digital spaces offers opportunities for novel applications. We shall highlight three of them here.

Military Mission Exercises

Traditionally, military exercises are either carried out in the physical realm or the virtual domain. In the physical realm, soldiers and military vehicles are mobilized for operations in some physical terrain. In the virtual domain, commanders "sweat out" in air-conditioned rooms in a simulated warfare over 3D virtual world models of the physical entities. While the former is realistic, it is limited by scale (both in the number of personnel and physical space); the latter, however, handles large scale warfare at the expense of actual ground happenings (e.g., it may take much longer time to cross a river physically than estimated in a model because of ground constraints and fitness of the soldiers; moreover, in a model, soldiers can walk through a building destroyed by artillery, but given that this may not actually happen in the physical space during an exercise, the time to bypass the building may be much longer.).

With co-space technology, we can now conduct a more realistic military exercise that takes on a completely new experience and flavor. Consider an exercise that involves both a small scale military exercise (the physical space) and a virtual model of a large scale military exercise. The physical space essentially forms a small part of the entire military exercise (e.g., a physical exercise over a physical space of 5 km by 5 km compared to a virtual model that simulates a war over 100 km by 100 km space). Now, on the physical ground itself, soldiers and vehicles are equipped with location tracking devices to monitor their movement as well as other information such as fire-power, casualties, etc. At the command center, based on real-time feed of the sensed data, the virtual model more accurately reflects the ground situations. Moreover, actions taken within the virtual world (e.g., simulating reinforcement, enemy counter-attack, etc) will be relayed to the ground troops that may then further influence the ground decisions. For example, if a region in the ground occupied by troops were air-raided, then the troops must "die".

Co-Space Marketplace

In today's marketplace, you either shop in a mall or you can buy your products online. If a physical shop exists that also has a web page (a very primitive form of virtual model), there is a very limited real-time interaction between the two spaces - when a customer purchased an item, the quantity-on-hand may be updated immediately.

In the near-future co-space marketplace, a physical mall will be "expanded" into a mall (virtually) that houses many more shops than the physical mall. In addition to the virtual correspondence of the physical shops, the virtual mall can rent out virtual space for virtual shop owners. At the physical mall, screens (large displays) can be set up within each physical shop for cyber shoppers to communicate with physical shoppers within the same shop (e.g., through text messages). While a physical shopper is restricted to the shops that are physically located in the mall, the online shopper has a wider selection of shops (and products). The virtual mall need to be kept up-to-date with real time information from the physical mall, e.g., live programs that are happening in the physical mall, on-going lucky draws, updates on availability of products, etc. In addition, the cyber and physical shoppers can interact with one another. When both are in the same shop (one in the physical space, the other in the virtual space), they can communicate and benefit from discounts (e.g., for a "buy two for the price of one" offer, each can buy one while sharing the cost) or complain to one another over poor services.

This concept can be easily extended to build and "expand" a stadium sitting capacity to target global audience, "expand" the space for exhibits in a museum, and so on.

Co-Space Gaming and Social Networking

One class of gaming in a co-space environment is location based gaming (LBG). LBG is gaining popularity and is believed to be the future of video-gaming where a player's everyday experience (e.g., of moving around the city) is interleaved with the extraordinary experience of a game. These games deliver an experience that changes according to the player's locations and actions.

In LBG, a user equipped with a GPS-enabled handset (e.g., a mobile phone) can play a video game that combines a player's real world (aka. his physical location) with a virtual world on the handset. The physical location becomes part of the game board, and the player's movement directly influences the gaming progress (may affect the game character and/or environment). BotFighters and Swordfish are examples of LBG.

Another form of co-space games integrates a physical environment with a corresponding virtual model. Here, RFIDs and sensors are used to capture information about the players' current context, which are transmitted to a server. The server (which may be controlled by another player) follows the game rules and relays back to the physical players information that help them to proceed (e.g., locations of enemies in the vicinity). Examples of this category of games include Wanderer, PAC-LAN, MobHunt, GoogleTron, and Tourality.

It is also not hard to visualize that social networking can also be conducted in co-space. A person in a certain location in the physical space will be able detect a friend at the same location in the virtual space and together fight some monsters that are in the virtual space or do some shopping together in the co-space. They may form interest groups to share information and trade user-created contents and virtual valuables. Similarly, two "comrades" who fight together in the virtual space will be able to detect each other when they are near to each other in the physical space giving opportunity for more interaction.

It would be interesting to see how the multi-billion dollar industry of games and social networking will grow as advances in technologies to support co-space become mature.

3. WHAT THE DATABASE COMMUNITY HAS TO OFFER

From the above discussions, we have the following observations of a co-space environment.

- There is a large amount of data/information generated within co-space. Some of these are static (e.g., maps, quantity-on-hand), while others are dynamic (e.g., locations, sensor data) and frequently changing. Moreover, large amount of data may have to be streamed from one space to another, particularly from the physical to the virtual to ensure real-time tracking of the environment.

- There is a large number of sensors that are used to capture the data from the physical environment. In-network processing may be needed to aggregate data before transmission.
- There is a large number of events generated within co-space. These have to be monitored, and may trigger further actions/events both in the physical and virtual worlds.
- There is a large number of users (and queries). Each user device basically contributes a distributed node into a highly distributed environment.

Clearly, our community has been dealing with the above-mentioned (perhaps, not at the scale that co-space entails). We pride ourselves for managing large datasets. We have addressed and are addressing a wide range of research problems that are relevant - sensor networks, data streams, distributed databases, update-intensive operations, search and data retrieval. As such, our experience will enable us to contribute to this new field and to chart the research directions ahead.

4. CO-SPACE CHALLENGES

Being an integration of the physical and virtual spaces, it is certain that co-space brings with it the research issues within each space. In the physical domain, we need to design efficient and effective methods to sense the physical environment (through extensive use of RFIDs or sensing devices), to transform these data into a form that users will appreciate (through data cleansing, data mining, aggregation or interpolation), and to process queries in-network, and so on. While some work has been done (e.g., [12, 21, 26, 36]), we are only scratching the surface to realize practical deployment.

In the virtual space, with the popularity of Massively Multiplayer Online Games, there has been tremendous amount of interest in recent years to design techniques to support interactive virtual environments for users to communicate with each other in real-time [14, 37, 38, 41]. As pointed out in [38], there are a number of research challenges that need attention, including designing database engines for games workloads and methods to guarantee consistency across multiple virtual views. Techniques for caching and indexing virtual environments (e.g., [33, 34]) need further study to scale to the large number of users.

For the rest of this section, we shall focus on challenges that arise as a result of the integration between the two spaces that may be of interest to the database community. Some non-database related issues include (a) novel interface technologies that can seamlessly link the physical and cyber spaces to support real-time interaction between users within the two spaces; (b) innovative visualization and presentation of output (events and data) within the co-space on a wide range of devices and platforms (small vs large displays, fixed vs mobile); (c) techniques, tools and devices for capturing data from the physical environment, and for creating content (high quality digital images, animation and effects) for the virtual environment; (d) language translation, transcription and mediation methods to support social networking and learning, and many others (e.g,. security and networking infrastructure).

Figure 2: The Co-space of a Library

4.1 Data Fusion over Heterogenous Data Sources

Data fusion is generally defined as the use of techniques that combine data from multiple sources through inference in order to produce data that is potentially more accurate than if they were obtained from a single source [15]. While data fusion has been studied in the context of sensor networks, data fusion in co-space is more challenging as the inputs may come from a wide variety of sources including blogs, video/audio clips, photographs about events that took place in the digital and physical world.

As an example, consider the co-space of a library in Figure 2, information from both video camera and RFID readers will be needed to ensure that the location of books are represented accurately in the digital space. Furthermore, reviews and opinion on the book can also be drawn from both the Web and the social network of the user to enhance the browsing experience. Such fusion of information on a single entity requires a substantial amount of inference over semantics that are extracted from multiple data sources.

From the above discussion, we note that co-space data management is related to the well studied fields of data stream processing [42], sensors network [26] and data integration [25, 24]. However, it also differs in at least two ways. First, unlike the relatively simple aggregation that is being done over data streams and sensors presently, co-space data management requires more complex logic inference over these data sources. Second, unlike data integration which aims to derive a common schema for a set of heterogeneous databases, co-space data management need not attempt to do so but will instead try to detect events that had taken place based on these data sources and try to depict these events accurately and efficiently in the co-space.

There is a clear need to develop data fusion mechanisms that can deal with these two issues effectively.

4.2 Distributed/Parallel Architecture

With a large number of cyber users, and physical users with handheld devices, the co-space environment naturally forms a distributed (peer-to-peer) system. The system is highly complex because of the heterogeneity of the devices. Moreover, there is an enormity of static and dynamic data that flow within each space and across spaces.

For queries that access static data that are stored locally, techniques that can facilitate search/discovery of relevant information are critical. P2P search methods may be applicable here [17, 20, 39]. However, for dynamic data that need to be streamed from one space to the other, these methods may not be suitable. While there has been considerable work on distributed stream processing [1, 5, 18], these are restricted to query processing and typically assume a smaller number of sites and do not address the heterogeneity across the sites. Here, it seems that publish/subscribe architecture [9, 13, 42, 43] may be more effective. Novel architectures that can support streaming data and search efficiently are needed. For example, we envision a publish/subscribe system over peer-to-peer networks where each peer may be a highly parallel cluster that can support large number of mobile clients.

The need for supporting a large number of concurrent and both data and computational intensive activities, requires new system architectures to be autonomic and adaptive and scalable, in which loads are adaptively balanced and new nodes can be easily added without substantial reconfiguration effort. Recently, the processing paradigm of MapReduce [7] and other similar applicative programming frameworks have revolutionaized the extreme data analysis on clusters, and systems such as Clustera [8] exploit modern software building blocks for efficiency and scalability. These and some other recent efforts in exploiting multi core architectures and commodity hardware may provide a basis for development of new database engines. We shall examine some of the related issues below.

4.3 Database Engines for Co-Space

Managing co-space calls for a re-examination of the database engines as we understand today. This is because we are dealing with (a) a large amount of diverse types of data, ranging from structured to unstructured, textual to video, static and dynamic; (b) data that exist in two different spaces.

Storage Manager

While it is clear that data of different types need to be managed separately, it is not immediately clear that data of the same type from the two spaces should be treated separately. In other words, should the location of a shopper in the physical mall be stored together with the location of an online shopper; or should the real-live images of exhibits in a museum be handled in the same way as the corresponding pictures available in the virtual space. On one hand, we can simply tag data to reflect the space it belongs to. This offers a unified view of the co-space and simplifies the management of data. However, for operations that involve only data from a particular space, the performance may be penalized. On the other hand, we can organize the data from the two spaces separately. But, this may end up duplicating resources. Moreover, it may be possible to have a hybrid strategy - for certain data types, integrating them may be the best; for others, keeping them distinct may be optimal. It would also be interesting to study how recent storage designs such as row- or column- oriented stores [2] and self-organizing storage [19] can be exploited for co-space applications.

In the context of a distributed architecture, we need to design techniques that partition the data across the sites for efficient processing.

Query Processing and Optimization

Query processing and optimization in a co-space environment will require novel mechanisms. First, new operators may have to be introduced. As an example, sensor data may have to be interpolated (or combined using some user-defined functions) for them to be consumed by the virtual space. In fact, data can be processed and transformed as soon as it is received; alternatively, it can be transformed at runtime. As another example, data in the virtual space may be interpreted in a different way from those in the physical space. These will inevitably lead to changes to the optimizer so that it can be aware of these operators in order to generate an optimal plan. Hellerstein's earlier work on optimizing queries with expensive predicates may offer a good starting point [16].

Second, the performance requirements for the two spaces may not be necessarily the same. For example, it is reasonable to prioritize sales for a shopper in a physical mall than for an online shopper (when they both wanted the last available item). As another example, in the case of a cyber user, while real-time information is highly desirable, approximate data may be tolerated (e.g., instead of a high resolution video stream, a low resolution stream or animation may be acceptable). This calls for query processing or optimization techniques to be "space" aware. Moreover, efficient approximation techniques in the virtual space that do not sacrifice the quality of the output significantly are highly desirable.

Third, besides I/O, CPU and bandwidth consideration, the optimizer may have to be device-aware so that a feasible (and optimal for the device) plan can be generated. Some works on processing in portable devices [23, 27] and energy-efficient optimization [3] can potentially be extended for co-space.

Fourth, we are dealing not only with moving objects (some moving in the physical space), we are also dealing with moving queries (user moving in the virtual environment may need to track all users within his or her view - as the user moves, his or her view of the space changes). There are very few works on moving queries over moving objects [11, 10], and this area is certainly worth further exploration.

Finally, one key challenge in designing a distributed architecture is to ensure that meta-data that are required for optimization can be estimated locally at each site/cluster to minimize information exchange, while at the same time the quality of the generated plan may not be significantly compromised. Designing such a co-operative system is difficult. Techniques from distributed databases may be relevant here [35].

Indexing

As mentioned, co-space offers a wide diversity of data. To manage this, we may need novel indexing methods. For example, in [34], a HDoV tree is proposed to index content at different degrees of visibility in a virtual walkthrough environment. This structure is obtained statically, and requires high computational overhead. In co-space, we may need a more robust and dynamic structure to cater to the frequent updates of information. While some work has been done for location data [6, 22], no such indexing methods have been designed for the virtual domain. We need more flexible schemes to be able to handle update intensive applications and frequently changing scenes.

Buffer Management/Caching
The two categories of data (coming from the physical and virtual domains) call for novel buffer management and caching schemes. In particular, we expect an effective scheme to be conscious of the semantics. For example, data from the real space may be given higher priority over data from the virtual space. However, we need to develop criteria to compare the priorities across the two domains.

4.4 Data Consistency

In networked virtual environments, it is important that users have a consistent view of the virtual world. This requires transmitting data within the virtual world. Unfortunately, there is to-date no solutions that can scale well. Now, in co-space, the requirement of consistency becomes even more challenging - the virtual world must also reflect what is happening in the real world. Given the constraints in bandwidth and the large amount of data to be transmitted, we do not expect to see a truly consistent view in both worlds. However, we can try to keep the virtual world as close to the real world as possible. One solution is to tolerate some degree of discrepancies - for numerical data, this may be within certain coherency requirement; for multimedia data, a low resolution image/video may be used instead. Some recent works have looked at how to disseminate streaming data to a large number of clients while preserving data coherency [4, 30, 31, 43]. These techniques assume a small number of distinct objects, and so do not scale to large number of objects.

A closely related approach is to study how the data to be transmitted should be prioritized. For example, more critical data can be transmitted first before less critical data. We can learn from methods developed for intermittently-connected and disruptive networks [40]. We believe there is much needed avenue to be explored in this aspect, e.g., to study different scheduling schemes. Besides prioritizing data, it may also be necessary to develop techniques to schedule multiple (continuous) queries that meet different Quality of Service (QoS) metrics. While techniques developed in [32, 29] provided some insights on how this can be effectively handled, we believe this direction deserves further investigation.

5. CONCLUSIONS

The advancement in technologies has changed the way we live. In the real world, we can participate in virtual games. In the world of the virtual, we can shop, engage in strategic games that thrill us and receive real-time information and acquire knowledge. The merging of these two spaces will further enhance user experience. This paper has argued for the co-existence of the two spaces, not as independent entities but as an integrated world where the two spaces interact simultaneously, and users experiencing an augmented world (either reality or virtuality) seamlessly. We have presented several promising applications of co-space, and discussed some research issues that the database community can contribute.

In our discussion, we have focused primarily on the present; with virtual space technology, time no longer "bounds" us - we can, for example, be physically at a historical site experiencing virtually an event that transpired in history on the exact spot that we are standing; likewise, we can have a virtual futuristic view of the current location.

As researchers, we look forward to the exciting challenges in this field, and encourage members of our community to join us. Perhaps, by 2015, we will experience the world of co-space as end-users and be brought "back to the future"!

6. REFERENCES

[1] D. J. Abadi, Y. Ahmad, M. Balazinska, U. Cetintemel, M. Cherniack, J.-H. Hwang, W. Lindner, A. Maskey, A. Rasin, E. Ryvkina, N. Tatbul, Y. Xing, and S. B. Zdonik. The design of the borealis stream processing engine. In *CIDR*, pages 277–289, 2005.

[2] D. J. Abadi, S. Madden, and N. Hachem. Column-stores vs. row-stores: how different are they really? In *SIGMOD Conference*, pages 967–980, 2008.

[3] R. Alonso and S. Ganguly. Query optimization for energy efficiency in mobile environments. In *FMLDO*, pages 1–17, 1993.

[4] M. Bhide, K. Ramamritham, and M. Agrawal. Efficient execution of continuous incoherency bounded queries over multi-source streaming data. In *ICDCS*, page 11, 2007.

[5] S. Chandrasekaran, O. Cooper, A. Deshpande, M. J. Franklin, J. M. Hellerstein, W. Hong, S. Krishnamurthy, S. Madden, V. Raman, F. Reiss, and M. A. Shah. Telegraphcq: Continuous dataflow processing for an uncertain world. In *CIDR*, 2003.

[6] S. Chen, B. C. Ooi, K. L. Tan, and M. A. Nascimento. St2 b-tree: a self-tunable spatio-temporal b+-tree index for moving objects. In *SIGMOD Conference*, pages 29–42, 2008.

[7] J. Dean and S. Ghemawat. Mapreduce: Simplified data processing on large clusters. In *6th OSDI Conference*, 2004.

[8] D. DeWitt, E. Robinson, S. Shankar, E. Paulson, J. Naughton, A. Krioukov, and J. Royalty. Clustera: An integrated computation and data management system. In *VLDB*, 2008.

[9] P. T. Eugster, P. Felber, R. Guerraoui, and A.-M. Kermarrec. The many faces of publish/subscribe. *ACM Comput. Surv.*, 35(2):114–131, 2003.

[10] B. Gedik and L. Liu. Mobieyes: A distributed location monitoring service using moving location queries. *IEEE Trans. Mob. Comput.*, 5(10):1384–1402, 2006.

[11] B. Gedik, K.-L. Wu, P. S. Yu, and L. Liu. Processing moving queries over moving objects using motion-adaptive indexes. *IEEE Trans. Knowl. Data Eng.*, 18(5):651–668, 2006.

[12] H. Gonzalez, J. Han, and X. Shen. Cost-conscious cleaning of massive rfid data sets. In *ICDE*, pages 1268–1272, 2007.

[13] A. Gupta, O. D. Sahin, D. Agrawal, and A. Abbadi. Meghdoot: Content-based publish/subscribe over p2p networks. In *Middleware*, pages 254–273, 2004.

[14] N. Gupta, A. J. Demers, and J. Gehrke. Semmo: a scalable engine for massively multiplayer online games. In *SIGMOD Conference*, pages 1235–1238, 2008.

[15] D. Hall and S. McMullen. *Mathematical Techniques in Multisensor Data Fusion*. Artech House, 2004.

[16] J. M. Hellerstein. Optimization techniques for queries with expensive methods. *ACM Trans. Database Syst.*, 23(2):113–157, 1998.

[17] R. Huebsch, B. N. Chun, J. M. Hellerstein, B. T. Loo, P. Maniatis, T. Roscoe, S. Shenker, I. Stoica, and A. R. Yumerefendi. The architecture of pier: an internet-scale query processor. In *CIDR*, pages 28–43, 2005.

[18] J.-H. Hwang, U. Cetintemel, and S. B. Zdonik. Fast and highly-available stream processing over wide area networks. In *ICDE*, pages 804–813, 2008.

[19] S. Idreos, M. L. Kersten, and S. Manegold. Database cracking. In *CIDR*, pages 68–78, 2007.

[20] H. V. Jagadish, B. C. Ooi, K. L. Tan, Q. H. Vu, and R. Zhang. Speeding up search in peer-to-peer networks with a multi-way tree structure. In *SIGMOD Conference*, pages 1–12, 2006.

[21] S. R. Jeffery, M. J. Franklin, and M. N. Garofalakis. An adaptive rfid middleware for supporting metaphysical data independence. *VLDB J.*, 17(2):265–289, 2008.

[22] C. S. Jensen, D. Lin, and B. C. Ooi. Query and update efficient b+-tree based indexing of moving objects. In *VLDB*, pages 768–779, 2004.

[23] P. Kalnis, N. Mamoulis, S. Bakiras, and X. Li. Ad-hoc distributed spatial joins on mobile devices. In *IPDPS*, 2006.

[24] M. Lenzerini. Data integration: a theoretical perspective, 2002.

[25] A. Levy. Logic-based techniques in data integration. *Kluwer International Series In Engineering And Computer Science*, pages 575–595, 2000.

[26] S. Madden, M. J. Franklin, J. M. Hellerstein, and W. Hong. Tinydb: an acquisitional query processing system for sensor networks. *ACM Trans. Database Syst.*, 30(1):122–173, 2005.

[27] N. Mamoulis, P. Kalnis, S. Bakiras, and X. Li. Optimization of spatial joins on mobile devices. In *SSTD*, pages 233–251, 2003.

[28] P. Milgram and A. F. Kishino. Taxonomy of mixed reality visual displays. *IEICE Transactions on Information and Systems*, E77-D(12):1321–1329, 1994.

[29] L. A. Moakar, T. N. Pham, P. Neophytou, P. K. Chrysanthis, A. Labrinidis, and M. A. Sharaf. Class-based continuous query scheduling for data streams. In *DMSN*, August 2009.

[30] S. Shah, S. Dharmarajan, and K. Ramamritham. An efficient and resilient approach to filtering and disseminating streaming data. In *VLDB*, pages 57–68, 2003.

[31] M. A. Sharaf, J. Beaver, A. Labrinidis, and P. K. Chrysanthis. Balancing energy efficiency and quality of aggregate data in sensor networks. *VLDB J*, pages 13(4): 384 – 403, December 2004.

[32] M. A. Sharaf, P. K. Chrysanthis, A. Labrinidis, and K. Pruhs. Algorithms and metrics for processing multiple heterogeneous continuous queries. *ACM Trans. Database Syst.*, 33(1):5.1–5.44, March 2008.

[33] L. Shou, J. Chionh, Z. Huang, Y. Ruan, and K.-L. Tan. Walking through a very large virtual environment in real-time. In *VLDB*, pages 401–410, 2001.

[34] L. Shou, Z. Huang, and K. L. Tan. The hierarchical degree-of-visibility tree. *IEEE Trans. Knowl. Data Eng.*, 16(11):1357–1369, 2004.

[35] M. Stonebraker, P. M. Aoki, W. Litwin, A. Pfeffer, A. Sah, J. Sidell, C. Staelin, and A. Yu. Mariposa: A wide-area distributed database system. *VLDB J.*, 5(1):48–63, 1996.

[36] E. Welbourne, M. Balazinska, G. Borriello, and W. Brunette. Challenges for pervasive rfid-based infrastructures. In *PerCom Workshops*, pages 388–394, 2007.

[37] W. M. White, A. J. Demers, C. Koch, J. Gehrke, and R. Rajagopalan. Scaling games to epic proportion. In *SIGMOD Conference*, pages 31–42, 2007.

[38] W. M. White, C. Koch, N. G. 0003, J. Gehrke, and A. J. Demers. Database research opportunities in computer games. *SIGMOD Record*, 36(3):7–13, 2007.

[39] S. Wu, J. Li, B. C. Ooi, and K. L. Tan. Just-in-time query retrieval over partially indexed data on structured p2p overlays. In *SIGMOD Conference*, pages 279–290, 2008.

[40] Y. Zhang, B. Hull, H. Balakrishnan, and S. Madden. Icedb: Intermittently-connected continuous query processing. In *ICDE*, pages 166–175, 2007.

[41] S. Zhou, W. Cai, B. S. Lee, and S. Turner. Time-space consistency in large-scale distributed virtual environments. *ACM Transactions on Modeling and Computer Simulation*, 14(1):31–47, 2004.

[42] Y. Zhou, K. Aberer, A. Salehi, and K. L. Tan. Rethinking the design of distributed stream processing systems. In *ICDE Workshops*, pages 182–187, 2008.

[43] Y. Zhou, B. C. Ooi, and K. L. Tan. Disseminating streaming data in a dynamic environment: an adaptive and cost-based approach. *VLDB J.*, 2008.