



RFID Infrastructure Design

A Case Study of Two Australian RFID Projects

The design, implementation, and deployment of large-scale RFID infrastructures can occur only through innovative solutions that address RFID challenges at all levels, including device, data processing, and data integration. The authors present and discuss several design issues associated with two Australian national RFID projects, implemented based on the EPCglobal Network. Using practical experience gained from these projects, they identify several areas of improvement and research opportunities for future large-scale RFID implementations and deployments.

John P.T. Mo
RMIT University

Quan Z. Sheng
The University of Adelaide

Xue Li
The University of Queensland

Sherali Zeadally
University of
the District of Columbia

The ability to track individual objects is essential to many aspects of modern life, such as manufacturing, logistics, counterterrorism, and anti-counterfeiting. RFID – a wireless communication technology useful for precisely identifying objects – uses radio-frequency (RF) waves to transfer identifying information between tagged objects and readers without line of sight, thus enabling automatic tracking and tracing.^{1,2} Passive RFID can track products in supply chains from the supplier to the distribution center, warehouse, and point of sale. In a typical deployment, RFID tags are affixed to individual products, and readers installed at various checkpoints capture tag-reading events. Companies then transport these

tagged items along the supply chain, leaving trails of tag-reading events behind. This lets those companies precisely record product movement.

Unfortunately, passive RFID technology has limitations, both in terms of physical constraints placed on signal transmissions and how much data a tag can store.¹ To access more data about a product (such as its production date, batch number, or package size), organizations could store most of that information externally in a central data warehouse and use the RFID tag as a key. All organizations would have to agree on a common storage format and continuously upload their data to the central warehouse. This approach has several severe drawbacks. Compet-

Related Work in RFID Applications

Many organizations have recently deployed RFID pilot applications. However, significant challenges remain to developing and deploying large-scale and distributed RFID applications (such as nation-wide supply-chain management across companies).¹ Others have concluded that RFID technology's full potential can develop only through collaboration and data sharing across multiple sites and organizations.² One emerging industry framework that manages RFID-tagged product movement is the EPCglobal Network,³ designed at the Auto-ID Center in the US and further developed by EPCglobal (www.epcglobalinc.org). The International Organization for Standardization and the International Electrotechnical Commission are developing the 18000 RFID Air Interface Standard (www.hightechaid.com/standards/18000.htm) as an alternative to EPCglobal. The basic

idea behind EPCglobal is to realize a so-called *data-on-network* system in which RFID tags contain an unambiguous identity — that is, an electronic product code (EPC) — and other data pertaining to objects stored and accessed over the Internet.

References

1. Q.Z. Sheng, X. Li, and S. Zeadally, "Enabling Next-Generation RFID Applications: Solutions and Challenges," *Computer*, vol. 41, no. 9, 2008, pp. 21–28.
2. C. Bornhovd et al., "Integrating Automatic Data Acquisition with Business Processes Experiences with SAP's Auto-ID Infrastructure," *Proc. 30th Int'l Conf. Very Large Databases (VLDB 04)*, Morgan Kaufmann, 2004, pp. 1182–1188.
3. EPCglobal, "The EPCglobal Network and the Global Data Synchronization Network (GDSN): Understanding the Information and the Information Networks," position paper, Oct. 2004; www.epcglobalinc.org.

ing enterprises are often reluctant to put their data in a shared central database. In addition, such an approach isn't feasible for large-scale, distributed applications because RFID data has unique characteristics such as streaming and large volume.^{3–5}

Here, we examine two Australian RFID deployments to explore the functionality requirements associated with large-scale RFID applications. Led by one of the authors (John Mo), we designed these national projects to integrate with normal business operations between industry partners in two cities, Melbourne and Sydney. The constraints imposed when goods move physically between distant transaction locations have serious implications on the underlying RFID infrastructure's design. Numerous companies took part in these projects, increasing the complexity of business transactions. The practical experiences we obtained while working on them provided insight into how applicable the prevailing theoretical framework is to real-world industrial applications. In this article, we also examine some research opportunities — in particular, a data lineage model (DLM) — for constructing an infrastructure for large-scale, multi-enterprise RFID applications.

The National EPC Network Demonstrator Projects

The EPCglobal Network standard⁶ provides a promising architecture for tracking and tracing objects over the Internet (see the sidebar for more information). However, few real-world, large-scale application examples were reported. Consequently, practitioners lack guidance and have

only their own limited deployment experiences from which to learn. To gain first-hand, practical experience with using the EPCglobal Network standard, we developed two large-scale, national projects involving a total of 16 industry and four nonindustry partners, working between Melbourne and Sydney. To the best of our knowledge, these projects were the first in the world that successfully used the EPCglobal-defined protocol's full stack to enable interorganizational transactions and supply-chain management.⁷

The first project, the National EPC Network Demonstrator Project (NDP), aimed to identify the business benefits of

- sharing information securely using the EPCglobal Network standard,
- providing authentication to interacting parties, and
- enhancing organizations' ability to track and trace product movement within the entire supply chain and for transactions among multiple enterprises.⁸

The EPCglobal model defines one authoritative registry of numbers that companies can query for links to access detailed information from local servers instead of having each company store this information and communicate it to the next partner. Items of interest include *products* (identified by a Serial Global Trade Item Number, or SGTIN), *pallets* (identified by a Global Returnable Asset Identifier, or GRAI), and *unit loads* (identified by a Serial Shipping Container Code, or SSCC). To enable secure information sharing among the partners, the NDP consortium set up

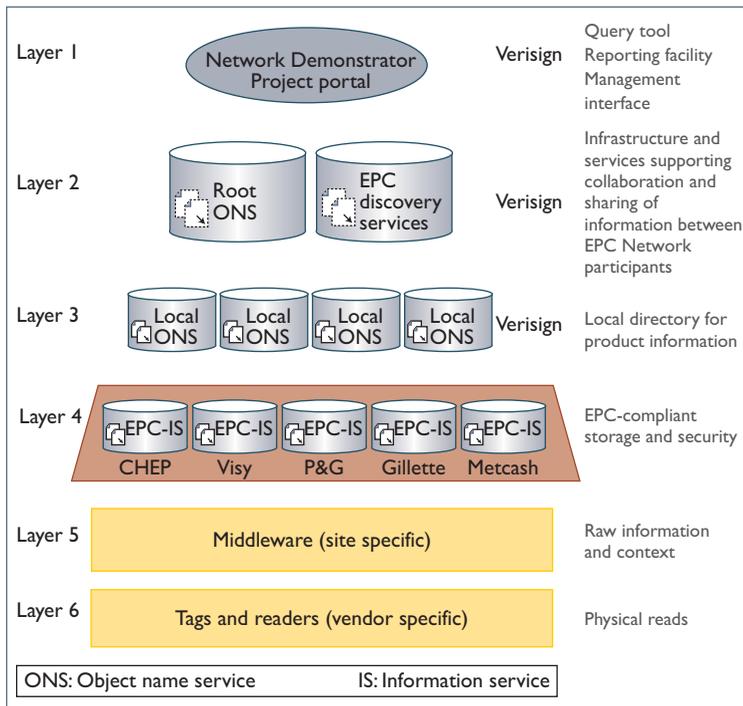


Figure 1. The NDP layered architecture. Verisign implemented the top three layers. The other parts of the EPC Network were implemented by individual partners or vendors separately.

an NDP portal with access control on a global server. Once partners logged in, they could access corresponding product information and its history using the electronic product code (EPC) as the search key. This data-sharing capability, which enabled data transparency in the traded items, was one of NDP's major advantages.

Encouraged by NDP's success, some partners regrouped as a second consortium to work on an extension project called the National EPC Network Demonstration on Business Information Integration (NDP Extension).⁹ This extension aimed to develop a paperless system to reduce human errors (for example, in data entry) and the subsequent corrective costs. Achieving this goal required participating companies to integrate the EPC Network with their business systems. The first NDP project didn't consider implementing such integration.

The NDP Extension particularly focused on asset management – in this case, pallets. Managing assets throughout a supply chain is onerous and surprisingly complex. As assets move among trading partners, the error in quantities sent and received often goes undetected until the invoice becomes available, which could be several weeks later. The partners must then negotiate over any discrepancies, which becomes

tedious given that each partner has its own version of events and supporting paperwork that might or might not prove who's responsible for the error. Because more than 10 million pallets are in circulation throughout Australia, the potential for loss is enormous, and efforts spent on managing assets have proved extremely costly.

The NDP Extension consortium performed a series of paperless test runs on pallet dispatch and delivery transactions among six sites between Melbourne and Sydney. It developed process models for each transaction route and analyzed the models' associated benefits.¹⁰ Efficiency improvements due to eliminating data entry, verification, and reconciliation processes varied among sites, with an average 18.1 percent reduction in labor time.⁹ Furthermore, improvements in inventory accuracy and real-time processing likewise improved performance in other logistic operations, such as planning and forecasting.

NDP Infrastructure Design

The EPCglobal standard gives guidelines about the network's necessary component characteristics. However, the standard doesn't contain implementation details. More work is thus needed to clarify system components' functionality and purposes. This section clarifies some of the implemented functionalities.

Layered Architecture

The distance between transaction locations was a significant factor, causing undesirable delays in the system's ability to complete transactions or provide visibility for the items in transition. The various NDP partners jointly developed the system to support these transactions, and it must thus maintain consistency across physical and organizational boundaries. The partners adopted a six-layer model (Figure 1), in part from the EPCglobal architecture, which specifies the top three layers; the partners developed the operational elements that span layers 4, 5 and 6:

- *NDP portal.* The consortium developed the portal to display the data collected and stored in the system and to allow companies to query the movement of goods through the supply chain.
- *Root object name service (ONS) and discovery service.* The ONS identifies unique numbers

for manufacturers, whereas the discovery service points to a particular EPC information service (EPC-IS) in which companies can obtain detailed information on a specific item.

- **Local ONS.** Each industry partner maintains its own repository of product-specific data. The local ONS provides a pointer to the local database.
- **EPC-IS.** At each site, the industry partners maintain details about site-specific product data that other participants can query.
- **Middleware.** An onsite software component converts multiple reads to one read, adds contextual information (such as the reading location and timestamp), and formats data for storage in an EPC-IS.
- **Tags and readers.** Physical reads capture RF signals in the form of hexadecimal numbers and transmit them to the server for the middleware layer to filter.

Adopting a common architecture solved half the visibility problems. However, partners still had their own internal systems and processes that weren't compatible for information exchange. Next, we'll look at three key challenges that the partners overcame during the NDPs – namely, *security*, *process synchronization*, and *event tracing*.

Security

Many people from various backgrounds participated in the projects, including managers, engineers, business analysts, assembly-line workers, researchers, and IT professionals. The system needed to accommodate different practices from different companies and overcome serious challenges in security during transactions and information exchange. To accomplish this, the partners implemented the EPC Network system's functionalities on secure platforms from VeriSign (for NDP) and Telstra (for the NDP Extension). Authorized users could then use the NDP portal querying process to query EPC data. To retrieve information they had to specify an EPC via the form

- urn:epc:id:grai:<manager number>:<class number>:<asset number>,

- urn:epc:id:sgtin:<manager number>:<class number>:<item number>,

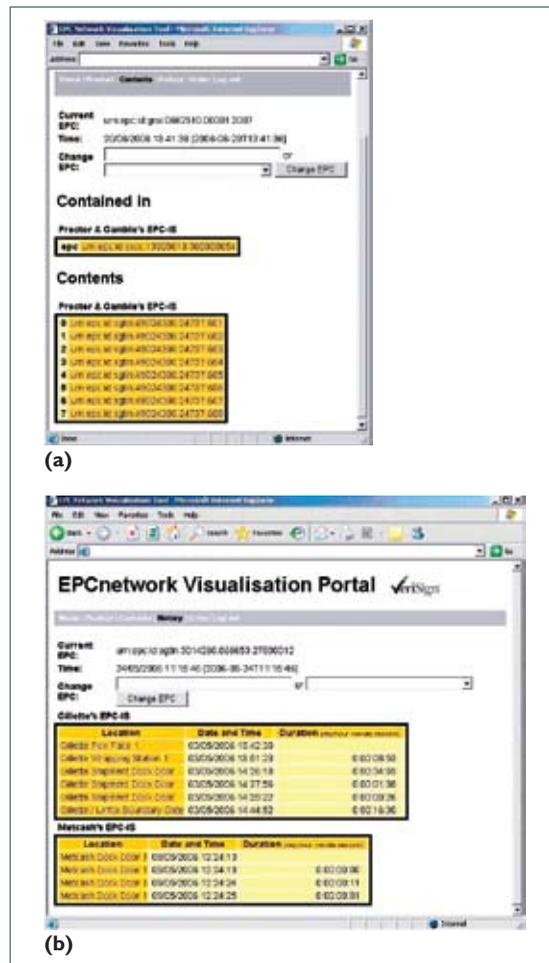


Figure 2. Containment and change of ownership. (a) Containment specifies the packing relationship when one or more RFID items are embedded into another RFID item and must be recorded in an RFID-enabled tracking environment. (b) An ownership change directly affects insurance, costs, and responsibilities.

- urn:epc:id:sgcc:<manager number>:<item number>.

Process Synchronization

As mentioned earlier, the projects included 16 industry partners who each brought their own business practices. Synchronizing and agreeing on processes among these interacting companies was difficult. Two key business processes were essential for accomplishing this within the NDPs. First, we used the concept of *containment* to relate items that were packed into a larger unit for transportation and storage. Containment specifies the packing relationship when one or more RFID items are embedded into another RFID item and must be recorded in an

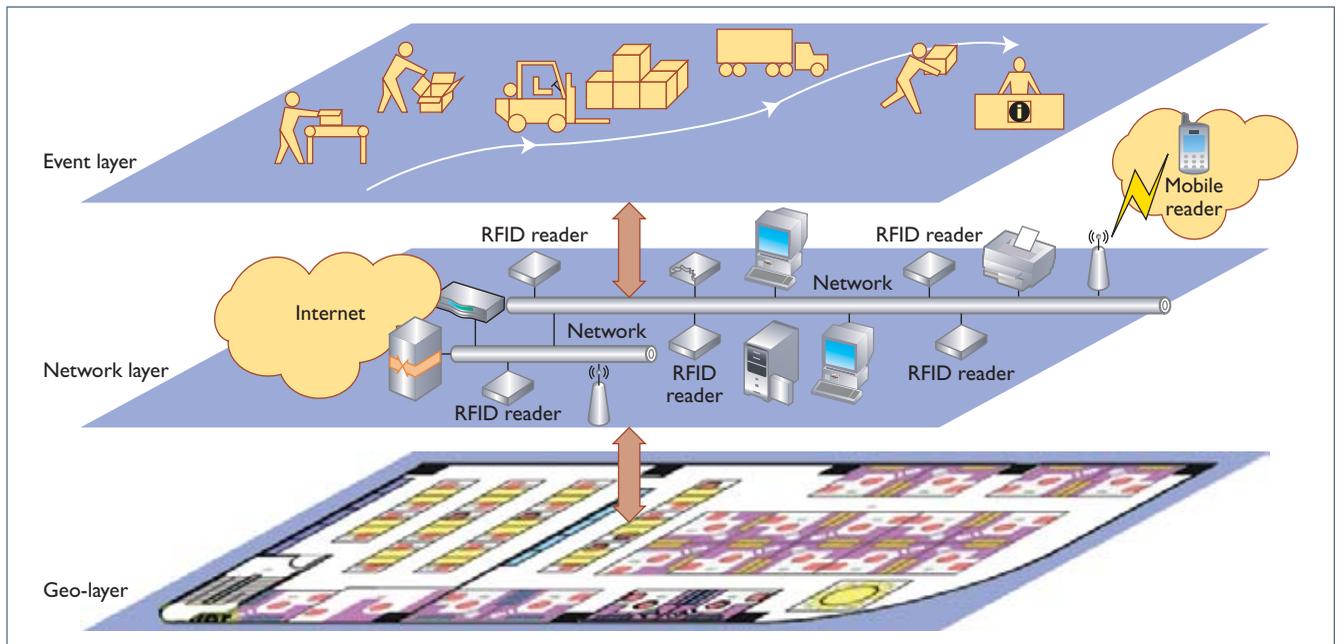


Figure 4. A data lineage model for tracking and tracing goods in supply chains. These three layers are logically independent from each other with changes. The system maintains a mapping mechanism.

where the truck had gone or what happened to the pallet).

Model for Large-Scale RFID Deployment

Enabling fast, large-scale RFID deployment requires a novel model – the system model the partners specially developed for both NDPs required a huge implementation effort. In particular, our experience within the NDPs demonstrated that the data in an RFID network is constantly evolving due to RFID tag mobility, change of item ownership, and change of workflows and event patterns. For any piece of RFID data, we need to see its physical values (for example, location and time), as well as its logical ones (such as its event pattern, ownership, and business transaction). RFID data has different representation levels. At the physical level, data represents floor plans and geographical locations. At the network level, it represents connections between mobile readers and networks. Finally, at the business level, data represents business events. Subsequently, the NDP partners must specially develop every component to match their different business practices and system requirements.

The constant changes inherent in RFID systems call for a novel RFID data model that lets us easily track and trace RFID data among these different representation levels. Inspired by the

uncertainty lineage databases (ULDBs) that Omar Benjelloun and his colleagues proposed,¹¹ we believe a DLM can support tracking and tracing RFID data efficiently. The DLM has three layers, as Figure 4 shows.

The *event layer* is responsible for modeling business transactions. In supply-chain management, we can define events by *pre-* or *post-conditions* and *triggers* that describe under which (spatial or temporal) conditions such events will occur and the actions the system should take. For example, we might model an ownership change as both a pre- and post-condition for different owner IDs when a distributor has an order on the product from the manufacturer and the retailer has a payment record in the distributor's accounts-receivable. The event could be a reading at a warehouse dock door, and the action would be registering the shipment and checking the product quantity on hand.

At the *network layer*, the system maps all events from the event layer to network configurations to make them traceable. For each business event, the system registers the associated network devices. This layer also protects the data lineage from illegal data injection and readings.

Finally, the *geo-layer* ties all network devices to their geo-properties (for example, in a floor plan with 3D locations) and their business ownerships. The system can further map

business events from the network devices' configurations to their corresponding physical locations. Given that we can use a mobile network to track moving objects, the system should dynamically map network devices' logical connections to their new locations. In this case, the mapping between the network layer and the geo-layer would represent a tight coupling for tracing business transactions at the event layer.

A *logic independence* DLM guarantees that changes occurring in one layer (such as a change to a business process) will be independent from other layers within the RFID system. Hence, changes in the lower layers (for example, a change in the network configuration) won't affect the business logic. Note that although the RFID Application Level Event standard⁶ that EPCglobal proposed emphasizes the importance of RFID data stream processing, it relies heavily on system developers to implement application-specific processes into the infrastructure.

Real-Time Query Language

RFID data is time-dependent, dynamically changing in large volumes and carrying implicit semantics.^{3,5} So, we need a general RFID data processing framework based on the lineage databases to process high-volume RFID data streams in real time and automatically transform physical RFID observations into virtual counterparts linked to business applications. RFID applications need a real-time query language that specifies how individual events are filtered and how multiple events are correlated via time- and value-based conditions. Based on our experiences with the NDPs, we believe that a real-time RFID query language should be able to query

- a sequence of complex events involving multiple readings from different local storage spaces and time ranges;
- event patterns in given occurrence frequencies; and
- network data streams along with classical functions based on the static data stored in a central database.

Unfortunately, no such language currently exists that satisfies all these properties. To enable temporal event tracing, we're currently working on a new time notion, *time-to-live*

(TTL), which represents the length of time an RFID event can legally live in an RFID system. TTL covers various complex temporal event patterns, including those that existing systems can process. In particular, we classify TTL into four categories to denote different RFID events:

- absolute TTL, which specifies the time that an RFID tag can live in the physical world;
- relative TTL, which specifies the time that an RFID tag can be used for a particular application (after which we can reassign the tag to other applications);
- periodic TTL, which specifies the time between two successive events with the same event type – in other words, it controls the period in which similar events occur periodically; and
- sequential TTL, which specifies the time between any two successive events.

A TTL-enabled query language can capture a wide range of interesting RFID applications with temporal restrictions. A typical supply chain, such as the one the NDPs implemented, requires companies to pack and re-pack items due to redistribution requirements. Having full traceability at the DLM geo-layer would generate global location data that is unified with temporal data and would provide a data-dependence-free model for globally tracing objects. Using the TTL concept to maintain temporal information about items, especially returnable assets, is also critical in tracking shipment items globally. Details about our work on TTL are available elsewhere.¹²

The practical experience gained working within the NDPs demonstrates areas in which we can improve RFID infrastructure for large-scale deployments. Adopting the EPCglobal RFID infrastructure for the NDPs enabled global identification, and standardizing system and data architectures among the partners allowed EPC tracking and event tracing across multiple business partners and workflows. However, the NDPs didn't have all the elements required for meeting current supply-chain RFID applications' needs. An RFID system aims to make decisions based on the identities of commodity items passing through different links in the supply chain. Subse-

quently, such information should efficiently support settlement (acknowledging delivery/receipt of goods), operation (inventory levels), and planning (resource scheduling).

Our experience with the NDPs identifies the need to segregate functionalities in large-scale RFID applications (such as a nation-wide supply chain) into three layers. We're investigating the data lineage approach to address this need. Mapping the DLM layers to the EPC Network requires support in several areas that we must address in the future. We'll need new functionality over the NDP infrastructure to cater to crucial business requirements in business event processing, continuous tracking, and identification regeneration. □

Acknowledgments

The Australian government has supported the projects discussed in this article through the Information Technology Online (ITOL) program in the Department of Communications, Information Technology, and the Arts. Quan Z. Sheng's work has been partially supported by Australian Research Council Discovery Grant DP0878917. Xue Li's work has been partially supported by Australian Research Council Discovery Grant DP0558879. We thank the anonymous reviewers for their valuable feedback.

References

1. B. Nath, R. Reynolds, and R. Want, "RFID Technology and Applications," *IEEE Pervasive Computing*, vol. 7, no. 1, 2006, pp. 22–24.
2. G. Roussos, *Networked RFID: Systems, Software and Services*, Springer, 2008.
3. S.S. Chawathe et al., "Managing RFID Data," *Proc. 30th Int'l Conf. Very Large Databases (VLDB 04)*, Morgan Kaufmann, 2004, pp. 1189–1195.
4. H. Gonzalez et al., "Warehousing and Analyzing Massive RFID Data Sets," *Proc. 22nd IEEE Int'l Conf. Data Eng. (ICDE 06)*, IEEE CS Press, 2006, p. 83.
5. Q.Z. Sheng, X. Li, and S. Zeadally, "Enabling Next-Generation RFID Applications: Solutions and Challenges," *Computer*, vol. 41, no. 9, 2008, pp. 21–28.
6. EPCglobal, "The EPCglobal Network and the Global Data Synchronization Network (GDSN): Understanding the Information and the Information Networks," position paper, Oct. 2004, www.epcglobalinc.org.
7. "EPC Network Australia Demonstrator Project Report," GS1 Australia, CSIRO, 2006; www.gs1au.org.
8. J.P.T. Mo, "Development of a National EPC Network for the Tracking of Fast Moving Consumer Goods," *Int'l J. Enterprise Network Management*, vol. 2, no. 1, 2008, pp. 25–46.
9. J.P.T. Mo et al., "Achieving Full Business Process Integration using EPC Compliance Technology," *Proc. 5th Int'l Conf. Supply Chain Management and Information Systems (SCMIS 2007)*, Monash Univ., 2007, p. 15.
10. S. Gajzer, *Process Modelling and Development of the Key Performance Indicators for the National Electronic Product Code Network Demonstrator*, MEng dissertation, School of Aerospace, Mechanical and Manufacturing Eng., RMIT Univ., 2007.
11. O. Benjelloun et al., "ULDBs: Databases with Uncertainty and Lineage," *Proc. 32nd Int'l Conf. Very Large Databases (VLDB 06)*, ACM Press, 2006, pp. 953–964.
12. X. Li et al., *Time To Live: Temporal Management of Large-Scale RFID Applications*, tech. report SSE-2008-02, School of Information Technology and Electrical Eng., Univ. of Queensland, Oct. 2008.

John P.T. Mo is professor and head of the discipline of manufacturing and materials engineering at RMIT University. His research interests include system integration, risk analysis, infrastructure modeling, signal diagnostics, and manufacturing optimization. Mo has a PhD in manufacturing engineering from Loughborough University, UK. He is a fellow of Engineers Australia. Contact him at john.mo@rmit.edu.au.

Quan Z. Sheng is a lecturer in the School of Computer Science at the University of Adelaide. His research interests include Web services, business process integration, data integration, and Internet computing. Sheng has a PhD in computer science from the University of New South Wales. He is a member of the IEEE and the ACM. Contact him at qsheng@cs.adelaide.edu.au.

Xue Li is an associate professor in the School of Information Technology and Electrical Engineering at the University of Queensland. His research interests include data mining, multimedia data security, database systems, and intelligent Web information systems. Li has a PhD in information systems from the Queensland University of Technology. He is a member of the IEEE and the ACM. Contact him at xueli@itee.uq.edu.au.

Sherali Zeadally is an associate professor in the Department of Computer Science and Information Technology at the University of the District of Columbia. His research interests include wired and wireless computer networks, mobile computing, RFID, network and system security, and ubiquitous computing. Zeadally has a PhD in computer science from the University of Buckingham, England. He is a fellow of the British Computer Society and a fellow of the Institution of Engineering and Technology, UK. Contact him at szeadally@udc.edu.