

Optimizing Business Processes by Automatic Data Acquisition: RFID Technology and Beyond

Gregor Hackenbroich, Christof Bornhövd, Stephan Haller, and Joachim Schaper

3.1 Introduction

Information technology and enterprise software form the backbone of enterprise operations. Enterprise resource planning (ERP) systems automate and control financial and administrative processes. Manufacturing systems structure complex production lines. Inventory and shipping systems help organize the shipment of products to the customer. The enormous success of enterprise software is the result of both business needs such as increasing internal efficiency and customer responsiveness, as well as of advances in computer and software technology.

The first wave of ERP systems took place in the early 1990s [13]. The implementation of such systems forced companies to understand, document, and streamline their internal processes. The benefit was a higher efficiency of the core financial, accounting, and controller functions. In this generation of systems the data for all business programs was integrated in one massive database. This allowed process integration but at the same time led to large ERP installations whose implementation and customization were difficult and expensive.

The 1990s saw the shift from the mainframe architecture to the three-tier client/server model. New applications emerged that could talk to other applications on a network leading to a new level of complexity and flexibility. The modules of enterprise applications grew and became larger components: sales and distribution became Customer Relationship Management (CRM), materials management became Supply Chain Management (SCM), and product development and manufacturing became Product Lifecycle Management (PLM). The functionality of these applications bundled in business suites coordinated the flow of work across departments, divisions, and even companies, thus leading to a high level of information and process integration.

This chapter focuses on the next generation of enterprise applications [2]: in addition to information and process integration it will allow for the automatic identification and tracking of real-world objects, as well as the automatic

acquisition of additional data about these objects and their physical environments. The first such technology that currently enters large scale deployment is Radio Frequency Identification (RFID) [7], the contact-less electronic identification of objects through radio waves. RFID tags typically combine a modest storage capacity with a means of wirelessly communicating stored information like an electronic product code (EPC) to an RFID reader. In a supply chain management context, an object to be tagged is usually a pallet, a case or even a single sales item. Passive RFID tags require no on-board battery and can be read from a distance ranging from a few centimeters to a few meters. Active tags, on the other hand, come with an on-board battery which provides larger read ranges and memory sizes but also higher unit cost and size and a limited lifespan of typically 3-5 years.

RFID, however, is not the only emerging technology for object data acquisition: sensor nodes go beyond object identification and measure the status or condition of objects such as the object location, temperature, or acceleration [12]. Sensors equipped with processing and communication capabilities will be able to identify themselves, organize themselves into networks, and collaboratively perform simple business processes [1]. Networked embedded systems comprising even larger processing power as well as sensors and actuators will add intelligence to devices like cars, mining, or manufacturing equipment [10]. They will report the device status to backend systems or download from such systems business rules or software updates.

RFID tags, sensor networks, and networked embedded systems offer a range of devices for product integration; collectively we refer to objects equipped with such devices as *smart items*. The defining property of a smart item is that it can provide data about itself and that it has the ability to communicate this information [9]. Integration of smart items into business information systems can provide businesses with more accurate and timely data about their operations and can help streamline and automate the operations themselves. This leads to cost reduction and additional benefits like increased asset visibility, improved responsiveness, and extended business opportunities. However, bridging the gap between the physical and the digital world requires a flexible and scalable system architecture to integrate automatic data acquisition with existing business processes.

In this chapter, we discuss the kinds of changes that have to be made for enterprise systems to take into account smart items information and also the kinds of business applications that can be built if such information is available. As a specific example, we discuss *SAP's Auto-ID Infrastructure (AII)*, which converts object data into business process information by associating it with specified mapping rules and metadata. These mapping rules can feed incoming observation data directly to business processes running on either SAP or non-SAP backend systems, execute predefined business logic, or simply record the data in a persistent store for later analysis.

The remainder of this chapter is organized as follows: in Section 3.2 we describe enterprise software in more detail, focusing on SAP's supply chain

management solution [3]. We point out the improvements that Auto-ID technology can bring to SCM. We then address the requirements on Auto-ID systems (Section 3.3) and introduce the architecture of SAP's AII (Section 3.4). Section 5 discusses cases studies that were performed with a research prototype of the AII. Section 3.6 summarizes open research issues. And finally, we summarise our findings in Section 3.7.

3.2 Enterprise Software for Supply Chain Management

The adoption of RFID technology by leading retailers has created a large momentum for automatic product identification and ubiquitous technologies in general. To fully unfold the potential of such technologies, object data must be linked to product and asset information kept in business information systems. This involves a number of conceptual and technical challenges that result both from the object identification technology and the business application in question.

This section provides a closer look at the business applications, with a twofold goal: first to impart the reader with basic knowledge of business information systems, and second to introduce a class of business processes that will integrate ubiquitous technologies very soon. An exhaustive discussion of enterprise software applications is beyond the scope of this chapter; therefore we focus on solutions for supply chain management where the business value of Auto-ID technologies is most evident. A discussion of asset management solutions and the benefit that RFID support will bring to them can be found in Chapter 4 by Lampe, Strasser, and Fleisch.

Figure 3.1 depicts the interplay of selected SCM components with ERP systems and ubiquitous devices. For definiteness, we discuss solutions and use the terminology of SAP's enterprise software [8]. The components that reside in backend-systems reside within the top layer; they include the ERP system as well as components for supply chain and data management such as Warehouse Management (WM), Advanced Planner & Optimizer (APO), Supply Chain Event Management (EM), and the Business Information Warehouse (BW) (the functionality of these components will be described below). The business applications exchange data with the Auto-ID Infrastructure that collects and aggregates RFID and sensor data.

The purpose of supply chain support solutions is to provide businesses with all software necessary to analyze, control, operate, and plan its supply chain activities. Within SAP's supply chain management four applications support this task: Supply Chain Networking provides customers, suppliers and partners with a unified view on the activities within the supply chain; Supply Chain Planning synchronizes and optimizes planning activities; Supply Chain Execution deals with all production, procurement and transaction activities; and Supply Chain Coordination helps partners to assess both performance and risks and provides notification upon detection of unexpected

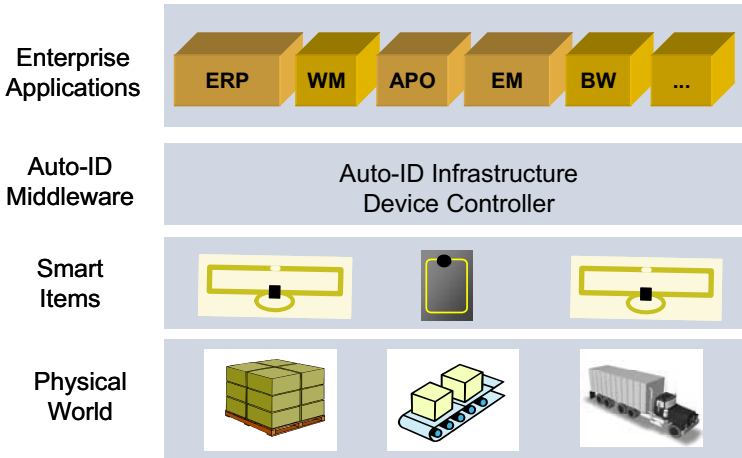


Fig. 3.1. Integration of real-world data into enterprise application. Warehouse management (WM), Advanced Planning & Optimizer (APO), and Event Management (EM) are SCM applications.

events. While offering a broad range of functionalities, these applications can still only address particular aspects of the supply chain. To provide a global view on all processes within the supply chain, SCM solutions must also integrate financial and business data; for example, financial data is required for analyzing the costs of supply chain processes. Likewise, Key Performance Indicators (KPIs) obtained from business intelligence systems are crucial for monitoring and controlling the supply chain. As a consequence, SCM solutions must tightly integrate with ERP and Business Warehouse systems.

To select the software components and business processes that gain most value upon integration with ubiquitous technologies, SAP performed an extensive market analysis. That analysis involved customer surveys, workshops, and customer councils. Key business problems that customers need to solve include poor data quality, limited supply chain and inventory visibility, and the substantial amount of manual labor involved in checking and managing incoming and outgoing goods. The components WM, APO, EM, and BW depicted in 3.1 address these problems but can be improved by automatic data acquisition as we show below.

WM deals with all processes relevant within a warehouse, including picking and packing, moving and storing goods, and in/outbound delivery. Warehouse management is facilitated by the integration of mobile devices that workers use to enter and retrieve reliable information about goods stored and moved through warehouse locations. BW may be thought of as a metadatabase. It integrates all databases that an enterprise has in place, guarantees a consistent view on the respective data, and generates from this data Key Performance Indicators. These, in turn, are exploited by the SCM planning system. The

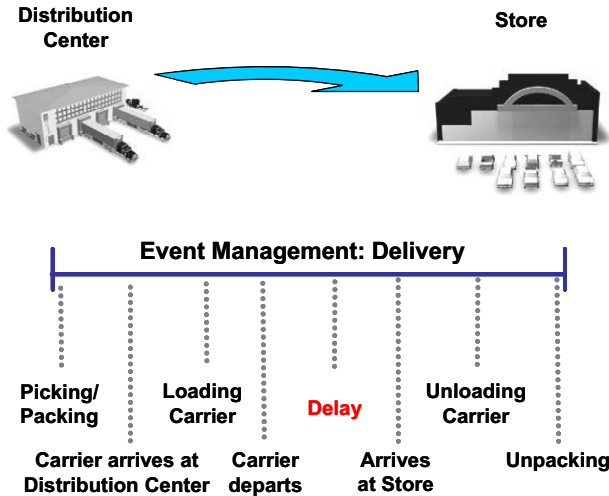


Fig. 3.2. Sequence of events in a delivery process. In traditional enterprise systems events must be reported manually. An RFID-enabled solution automatically generates events, thus allowing for automatic real-time supply chain planning and execution. SAP's Auto-ID Infrastructure supports the processes packing/unpacking and loading/unloading.

focus of BW is on the quick integration of huge amounts of data (in the Terabyte-range) from all possible sources within the enterprise. Finally, EM is an application for monitoring, managing, and tracking the progress of supply chain activities. Activities are represented in terms of events like start production/end production, packing/unpacking etc. For each event one can define assessment and execution rules (for example follow-up activities to be executed upon late deliveries). In addition, notification rules specify the persons and systems to be notified, for example in case of alerts, together with the respective communication channels.

To allow for automatic decision making, EM makes use of complex rule engines that take into account correlations along the full supply network including production planning. EM simulates the logistic processes and the correlation of events such as order compliance and replenishment of materials. The system continuously monitors events and detects when they happen outside of predefined tolerances. This allows quick identification of supply chain disturbances so the system can proceed according to the respective exception rules. If, for example, the quality of preliminary products is reported to be outside the allowed tolerance, then planning and scheduling systems get this information and calculate a new plan. A sequence of typical events for a delivery from a distribution center to a store is shown in Figure 3.2.

The sequence of events highlights some of the improvements that smart items technologies can bring to enterprise systems: in traditional enterprise

systems events must be reported manually; this is time-consuming, event reporting is not guaranteed, and the process is error-prone. By contrast, RFID-enabled event management ensures accurate, real-time event data that, in turn, improves supply chain execution and optimization. In addition to improving existing processes, smart items technology may generate new processes, e.g., automatic gates can monitor the movement of goods within the store. For example, the system can determine if goods are still in the store back room or already on the shop floor. We discuss this process in one of the case studies below. We note that some of the events shown in Figure 3.2 can be automatically reported using RFID technology and passive tags. In particular this holds for packing/unpacking and loading/unloading, which are already supported by SAP's Auto-ID Infrastructure. Automatic reporting of delays will probably require the use of active tags or sensor nodes.

3.3 Auto-ID System Requirements

The existing architecture for our Auto-ID Infrastructure has been developed with the following system requirements in mind:

Scalability. Companies like large retailers are assumed to require throughput rates of about 60 billion items per annum [11]. Assuming 100 distribution centers, each with an average of five checking points per item, the system needs to guarantee an average throughput of at least 100 messages per second per distribution center. The size of an observation message can be assumed to be around 200 bytes, and the processing of an incoming observation message usually requires multiple database updates and the execution of business procedures at the backend system.

Open System Architecture. In addition to being hardware-agnostic, the architecture should be based on existing communication protocols like TCP/IP and HTTP, as well as syntax and semantics standards like XML, PML [8] and EPC [4]. This will allow the use of sensors from a wide array of hardware providers and will support the deployment of Auto-ID solutions across institutional or even country boundaries.

Efficient Event Filtering. The infrastructure needs to provide efficient means to filter out false or redundant readings from RFID or sensor devices. Also, it needs to provide flexible and configurable filtering of events to pass on only relevant information to the appropriate backend processes.

Event Aggregation. The infrastructure needs to support the composition of multiple related events to more complex events for further processing. For example, the system must allow the composition of individual object identification events for multiple individual cases and the corresponding pallet to only one complete-pallet-detected event.

Flexibility. The infrastructure needs to be adaptable to different business scenarios. Furthermore, the infrastructure needs to provide flexible means at the business logic layer to respond to abnormal situations, like the missing of




Fig. 3.3. System architecture of the SAP Auto-ID Infrastructure.

expected goods or company-internal re-routing of goods. To avoid redundant implementations of the same business rules in different enterprise applications, the infrastructure needs to offer means to deploy and execute them within the Auto-ID Infrastructure.

Distribution of System Functionality. A real deployment of an Auto-ID solution can be distributed across sites, across companies, or even across countries. This naturally requires a distributed system architecture. As a first step, we require that the Auto-ID Infrastructure supports the distribution of message pre-processing functionality (for example, filtering and aggregation) and, to some degree, business logic across multiple nodes to better map to existing company and cross-company structures.

System Administration and Test Support. The infrastructure must provide support for the testing of individual custom components used in the filtering and aggregation of events, as well as the end-to-end processing of RFID and sensor data. Good administration and testing support is a prerequisite for the deployment of a distributed Auto-ID solution in large-scale applications.

3.4 System Overview

The architecture of our Auto-ID Infrastructure (AII) is shown in Figure 3.3. Conceptually, it can be divided into the following four system layers.

At the Device Layer, different types of sensor devices can be supported via a hardware-independent abstraction layer. It consists of the basic operations for reading and writing data and a publish/subscribe interface to report observation events. By implementing this API, different kinds of smart item devices can be deployed within the Auto-ID infrastructure. Besides RFID readers, these devices can include environmental sensors, or PLC devices. The Device Operation Layer coordinates multiple devices. It also provides functionality to filter, condense, aggregate, and adjust received sensor data before passing

it on to the next layer. This layer is formed by one or more Device Controllers (DC). The Business Process Bridging Layer associates incoming observation messages with existing business processes. At this layer status and history information of tracked objects is maintained. This information includes object location, aggregation information, and information about the environment of a tagged object. A so-called Auto-ID Node (AIN) realizes this functionality. Finally, the Enterprise Application Layer supports business processes of enterprise applications, such as Supply Chain Management or Asset Management, running on SAP or non-SAP backend systems.

Our Auto-ID Infrastructure provides an infrastructure for realizing a complete Auto-ID solution. Most existing solutions only focus on a portion of such a complete solution. For example, the Auto-ID Savant middleware [5] provide only Device Controller functionality within our reference model. Since Auto-ID solutions can span organizations or even countries, standards for the interfaces between the components are essential. Therefore, the AII is compliant with the standards proposed by the EPCglobal consortium.

As part of the infrastructure, a test and workload generator tool is provided that can simulate messages coming from one or more Device Controllers or backend systems to an Auto-ID Node. Also, a scriptable simulator is available that can simulate multiple RFID readers. These tools allow the testing of an Auto-ID deployment without the installation of physical devices.

The following two subsections will explain the two main building blocks of the AII: the Device Controller and the Auto-ID Node.

3.4.1 Device Controller

A Device Controller (DC) is responsible for coordinating multiple smart item devices and reporting incoming observation messages to one or more Auto-ID Nodes. A DC supports two operation modes. In the synchronous mode, the Device Controller receives messages from an Auto-ID Node for direct device operations, such as to read or write a specific data field from/to a tag currently in the range of an RFID reader, or to read the value from a temperature sensor at a given point in time.

In the asynchronous listening mode, the DC waits for incoming messages from the sensor devices. Upon receiving such a message, additional data can be read and event messages can be filtered or aggregated according to the configuration of the DC. Note that when a DC is configured for asynchronous operations, it is still capable of synchronously receiving and executing commands.

Message processing in the DC is based on so-called Data Processors. We distinguish six different types of data processors. (1) Filters filter out certain messages according to specified criteria. For example, they can be used to filter out all event messages coming from case tags, or clean out false reads (“data smoothing”). (2) Enrichers read additional data from a tag’s memory or other devices and add this data to the observation message. (3) Aggregators




Fig. 3.4. Typical Data Processor chain.

can be used to compose multiple incoming events into one higher-level event (for example, mapping data from a temperature sensor to a temperature-increased event), or for batching purposes. (4) Writers are used to write to or change data on a tag or control an actuator. (5) Buffers buffer event messages for later processing and/or keep an inventory of tags currently in the reading scope of an RFID reader. (6) Senders transform the internal data structure of the messages to some output format and send them to registered recipients. Our implementation of the DC currently uses PML Core [8] as the output format. As new standards are developed, they can be incorporated by simply implementing appropriate new Senders.

The core functions of the Device Controller, in particular the message processing described above, are independent of the hardware used. For reading and writing the data on the tags, we use logical field names to abstract from concrete tag implementations. A field map provides the mapping between memory addresses on the tag and logical data fields.

Since all Data Processors implement the same publish/subscribe interface, they can be arranged into processing chains. Powerful message processing and filtering operations can be achieved by chaining together the right, possibly customized, set of simple data processors. This results in a very flexible framework which allows for the distribution of message processing functionality close to the actual sensor devices to reduce message traffic and improve system scalability.

Figure 3.4 shows an example of a typical processor chain used for dock doors in a supply chain scenario. For full coverage dock doors commonly use more than one reader. The main reason for this is national and international regulations that set upper bounds for the emission power in respect to the field strength of RFID systems. Such restrictions limit the range of RFID readers, so frequently more than one devices are required to provide full door coverage.

RFID readers sometimes generate false event messages. For example, because of physical reasons a tag is not seen during a particular read cycle. To

filter out these false tag-disappeared messages, a `LowPassFilter` is applied. Also, every tag that passes the radio field will issue two event messages: a tag-appeared and a tag-disappeared message. Since in the dock door scenario we are only interested in the fact that an item has passed the door, we can safely filter out tag-disappeared messages by using an `EventTypeFilter`. The `EPCEnricher` in the example is only needed if non-EPC tags (which are still common today) are used. These tags have a unique ID set by the manufacturer, and the EPC is actually stored in the user memory of the tag. In this case, the `EPCEnricher` reads the EPC and adds it to the event message. At a dock door, we want to collect all tags that are seen during a certain time window and report them in a single message to the backend system. The `TimeFixedSizeAggregator` and the `Send` processor in our example do this. In addition, a `StateBuffer` keeps track of all tags currently in the reader's scope for auditing and reporting purposes.

3.4.2 Auto-ID Node

An Auto-ID Infrastructure can contain multiple Auto-ID Nodes. An AIN is responsible for integrating incoming observation messages from the Device Controllers with the business processes running at the backend systems.

For an AIN, we distinguish between the interactions with Device Controllers (reader events from and control commands to Device Controllers), and interactions with backend enterprise systems (such as receiving master data from a logistics system and returning a confirmation). These interactions with the AIN are treated as either incoming or outgoing messages.

Incoming observation messages are routed to a rule engine which, based on the message type, evaluates a specified list of conditions. The result of the evaluation step is a set of qualifying rules for which one or more actions are executed in a specified order. Such an action can, for example, update the system status of an object in the local repository, communicate with the backend system, or generate and write EPC data to a tag.

Actions of a rule can pass on parameters and can trigger other rules at the Auto-ID Node. Based on the message type, messages can be assigned different processing priorities and can be specified as being persistent in the Auto-ID Node.

An Auto-ID Node provides a local repository which contains information about the current status and history of the objects being processed. This information includes data about the operations that have been applied to an object (e.g. move, pack, or unpack), its movement and current location, and its structure (e.g. packing information). Also, the repository replicates master data from the backend system about products and business partners, or the physical location and type of the RFID readers. The Auto-ID Repository provides the basis for the execution of business logic in the Auto-ID Node.

The use of customizable rules provides a flexible mechanism to specify and execute business logic at the Auto-ID Node. This allows the pre-processing

of incoming observation data and the handling of abnormal situations within the Auto-ID Infrastructure, such as discrepancies between a received advanced shipping notification (ASN) and a detected pallet. This, in turn, allows the system to offload processing from the backend systems.

Our work to date has focused mainly on Supply Chain Management scenarios, for which a standard set of rules is in place. The deployment of our Auto-ID Infrastructure in a different context, for example, in asset management or product lifecycle scenarios, simply requires the adoption or extension of the existing rules.

The Auto-ID Administrator provides a graphical tool that supports the reconfiguration of existing or the definition of new rules in an AIN at runtime. In addition, it allows the central configuration and control of the Device Controllers and smart item devices in the system.

3.5 Case Studies

The following sub-sections discuss two Auto-ID pilot installations based on a research prototype of the AII: a retail application, and an adaptive planning application. Finally, we summarize the lessons learned from these and other real-world experiences.

3.5.1 A Retail Application

The first pilot was conducted at a large retailer in Europe. Here, the Auto-ID Node was used as a kind of "Auto-ID data hub" to feed business event information from several processes to two backend systems: a data warehouse (SAP BW) for analytical purposes, and a tracking system (SAP EM) to track the status of deliveries. On top of this, the SAP Enterprise Portal was used as the user interface to provide both employees and project partners with a unified view of the entire system.

From the perspective of the retailer, the goal of this project was mainly to evaluate if and how RFID technology can be used in practice. While there is a lot of hype about the technology, only by putting it in a real environment can one learn what works and what does not, and possibly how to get around technical difficulties. In addition to technical issues, another question was how customers would accept the technology.

The main process covered by RFID technology was the tracking of deliveries from the distribution center to one dedicated store, as well as the movement of goods from the store's back room to the shop floor. Tagging was done on case and pallet level. There were four read points in this business process:

1. **Packing Station:** At the distribution center, all cases needed to be tagged and assembled into deliveries. An association between the pallet

and the cases loaded onto it was recorded. Once the packing was finished, a message was sent to the Auto-ID Node with information about the pallet and its associated cases.

2. **Goods Issue Gate:** After the deliveries were loaded onto a truck at the distribution center, they passed through a reader that registered what had passed. The reader was mounted in the dock door. The data from the reader was filtered, aggregated and then sent to the Auto-ID Node, which updated its inventory of goods.
3. **Goods Receiving Gate:** Similar to the previous read point, incoming goods were read and recognized as they arrived at the store.
4. **Back Room/Shop Floor Gate:** This represents another automatic gate where goods were scanned when they passed through to determine if goods were in the back room or already on the shop floor. Previously, the retailer was not able to make this distinction.

Bulk reading took place at all read points except at point 1. Depending on the nature of the products (metal cans, bottles with soda water, and so on) in the delivery, it was impossible to always read all tags of inner cases. However, since the pallets were not unpacked until they were on the shop floor, it proved to be unnecessary to achieve a read accuracy of 100%. Since the system provided the information about what was packed together from read point 1 (where the operational process guarantees a 100% read accuracy) in principle it was sufficient to detect only a single tag of the whole delivery at the other read points. The Auto-ID Node was able to deduce the other tags from the packing information.

The Auto-ID Node received observation messages from all read points to update its inventory information in its repository. Information was also sent to the Business Information Warehouse to allow for later analysis. The AIN stored information about the deliveries (actual and expected), so whenever it received a message from one of the first three read points, it inferred the delivery number from the EPCs detected. It could then tell the Event Manager, which was responsible for tracking deliveries overall, about the status change of the delivery, for example, that it had arrived at the shop.

RFID technology was also used on the item level for some distinct goods. For example, processed cheese was tagged in order to track expiration dates using a Smart Shelf. Because of limited space we will not describe these processes in more detail in this paper. This pilot implementation showed that item level tagging is technically feasible, but that the cost of tags themselves, of applying the tags to products, and of the required infrastructure (readers and so on) is currently still too high to make sense economically. Other reasons against tagging at the item level are public concerns regarding privacy.

At the time of the project, EPC tag data standards as now defined by EPCglobal [6] had not been developed. However, user requirements required standard identifiers encoded on the tag. Cases needed to have a GTIN (Global Trade Identification Number) of the product, and pallets either a SSCC (Serial

Shipping Container Code) or a GRAI (Global Returnable Asset Identifier). We therefore had to define our own mapping of these standard identifiers to EPCs, adding a serial number in the case of the GTIN mapping. Our mapping was based on [4].

The main benefit provided by RFID technology in this pilot was increased visibility of the goods, which could be used to make better decisions on when to reorder goods, leading to cost reductions because of lower inventory levels and increased sales because of increased on-shelf availability.

The software worked reliably. Because of the size of the pilot, scalability was not an issue and a single Auto-ID Node was sufficient. More daunting were the challenges regarding the hardware, like positioning and tuning the reader antennas to achieve good read accuracy while conforming to regulatory requirements in Europe, tag placement, cabling, and availability of tags, just to name a few. Workplace safety regulations added additional constraints.

3.5.2 A Real-Time Adaptive Planning Application

The second pilot involved a large retailer and a manufacturer in North America. In this pilot, SAP provided the same components as in the pilot described in Section 3.5.1, plus the supply chain planning component SAP APO. This pilot included three sites: a distribution center of the manufacturer, a distribution center of the retailer, and a retail shop. The main operational process consisted of three steps.

First, in the distribution center of the manufacturer, items were packed into cases and shipped to the distribution center of the retailer based on shipment orders. In the second step, the distribution center of the retailer verified the shipment on the case level and then sent a case to the retail shop. Finally, in the retail shop, the case was first placed in the backroom and then moved to the shop floor. The items contained in the case were put on a smart shelf in the shop. The following read points were defined:

1. **Pack Station at the Manufacturer:** After packing a case, a message with all the EPCs of the case and the contained items was sent to the Auto-ID Node. The Auto-ID Node forwarded the EPC of the case and associated shipment order to the tracking system (SAP EM), where an Event Handler was created with the expected shipment time, a tolerance for the shipping time and rules for exception handling — that is, what to do when a shipment did not arrive in time.
2. **Goods Receiving Gates at the Retailer:** There were similar gates both at the distribution center and at the shop. When these gates detected a case tag, messages with the detected case tags were sent to the Auto-ID Node. After updating the status of the associated locations and the physical objects, the Auto-ID Node sent a message to the tracking system to update the status of the corresponding Event Handler.

3. **Back Room / Shop Floor Gate:** These read points were similar to the read points at the receiving gates of the retailer.
4. **Smart Shelves in the Shop:** When items were added or removed from the smart shelf, messages containing the EPC of the moved objects were sent to the Auto-ID Node with the logical reader ID and the timestamp of when the objects were scanned. The AIN then forwarded the observation message for the first item from a case that appeared on the smart shelf to the tracking system to indicate that the contents of a case had been put onto the shelf and that the tracking process for that case was completed.

In this pilot, a shipment was associated with a single case as only one case was sent from the manufacturer to the retailer at a time. The Auto-ID Node maintained the status and also the history of the objects, including cases, items, and shipments. The tracking system was used to track all shipments. Therefore, only messages on the case level were sent to the tracking system, which monitored the delivery of shipments and handled possible exceptions in almost real-time.

Through the Auto-ID Node, the manufacturer could get inventory information about its products in the retail shop. Based on the history of sale records, the Auto-ID Node maintained a local prediction model. This model could be used to trigger a request to the SAP Advanced Planning and Optimization to adjust the shipment planning.

SAP Business Information Warehouse was used for analytical operations and reporting, in a similar way as in the pilot discussed in the previous section.

3.5.3 Evaluation

Our experiences with the pilots described in the previous sections as well as with the first months of Auto-ID ramp-up phase can be summarized by the following lessons learned.

End-to-end Solutions. Many customers request end-to-end solutions from the tags to the business applications. To offer such solutions, hardware providers, business software providers, and system integrators must join competencies. Solutions on the device and business process layers must support interoperability with different providers to facilitate partnerships.

Cross-Organizational Collaboration. The pilots contained multiple sites, and in the case of the second pilot even multiple companies. The full potential of smart items technology can only be unlocked through collaboration and data sharing across sites and organizations. A basis for such collaborations could be flexible risk sharing strategies that help increase profits for all partners along the supply chain. The hope is that the potential business improvements offered by Auto-ID technology can bring companies to overcome their current reluctance to collaborate in the near future. This reluctance, as well as technical integration challenges, are the main reasons why EDI has not been implemented to the extent initially expected.

Standards. One of the key issues is the use of common standards. To avoid integration nightmares, standards on the hardware layer (readers, tags), the communication layer (HTTP, XML), and also on the syntax and semantics layer (PML, EPC) should be used or must be developed. Deployment of components from different providers becomes feasible at a reasonable cost of ownership only with the right standards in place. We are actively involved in ongoing standardization efforts at EPCglobal and the W3C. Standardization has been largely focused on RFID, and is much less advanced for other smart items technologies such as wireless sensor networks and networked embedded systems.

Automatic Identification Is Not Just RFID. The main use case of smart items today is the universal unique identification of items. RFID is not the only technology that allows this; for example, barcodes can be used as well. Different technologies have different advantages and use cases. Thus, all of these must be easily integratable into one system. Furthermore, in a real environment RFID readers sometimes need to work with other devices such as traffic lights and light beam sensors. These heterogeneities are the rule, not the exception.

100% RFID Reading Accuracy Cannot Be Expected. For physical reasons, one cannot expect to have a 100% tag reading accuracy. One way to work around this problem is the proper selection of RFID frequency and hardware. A second strategy is to keep information about how objects are assembled and have the Auto-ID Infrastructure infer the missing information. For example, detecting the movement of a pallet known in the system will allow the system to infer the movement of all associated cases. We expect the reading accuracy problem to diminish when the hardware technology matures (but not to vanish for the physical reasons mentioned earlier).

Need to Support Out of Sequence Messages. To an Auto-ID Node, the connected Device Controllers form a distributed environment. In a real-world installation, network latency, different system clocks at the readers, and message batch processing all can cause the order in which observation messages arrive to be different from the order in which the corresponding events took place in the physical world. Therefore, the Auto-ID Infrastructure needs to be able to reorder incoming event messages based on knowledge about the physical structure and the business processes of a given site.

Business Logic at the Point of Action. Distributing data and functionality among a hierarchy of device controller(s) and AIN(s) allows for efficient event filtering and event aggregation. However, distribution poses new challenges for error detection: it happens that errors can only be detected at higher layers and that error handling commands must be passed downwards to take proper action. As an example, consider a wrong pallet that arrives at a retailer's receiving gate. The discrepancy between this pallet and the expected pallets according to the advanced shipping information will only be detected by the AIN. The AIN, in turn, could send a command to the device controller, and a red warning light could flash up at the receiving gate. In

this case, error detection creates additional network traffic as compared with local handling directly at the point of action. However, the huge reduction of traffic by filtering and aggregation clearly overcompensates for that effect, indicating the advantage of our distributed system architecture.

Device Administration and Management. The deployment of an Auto-ID solution usually includes a large number of RFID and sensor devices. Centralized administration tools to visualize, plan (capacity planning), deploy, configure, test, monitor, and upgrade remote devices is a prerequisite for the operation of large, highly distributed Auto-ID solutions. Our existing tools are a good starting point but more powerful tools are needed.

Deploying an Auto-ID Solution Is a Long-Term Task. The deployment of an Auto-ID solution will change the IT infrastructure, the business processes, and the operational processes of an organization. These fundamental changes cannot be done in a few weeks and may result in significant costs up front. It is essential for a company to have a long term migration plan addressing the required changes in the organization. Therefore, it is a good idea to start with a small pilot installation to learn about the required changes in an existing business environment before rolling out an Auto-ID solution on a large scale.

3.6 Open Issues

Based on our experiences with the existing prototype, we would like to point out the following open issues for future research in the area of smart items technology.

Different Qualities of Service. Different smart items applications require different qualities of service regarding event processing. For example, for high data quality, an Auto-ID infrastructure may have to provide end-to-end transaction support to guarantee that each observation message is only processed exactly once. That is, the system needs to guarantee that a predefined reaction to an event is executed exactly once - even in the case of a system or power failure. There is obviously a trade-off between higher degrees of reliability on the one hand and performance on the other. Accordingly, different qualities of service need to be defined and provided for different application classes.

Distributed Smart Items Infrastructure. The nature of smart items applications as well as scalability requirements may force a distributed system architecture. Although our existing Auto-ID Infrastructure allows the distribution of functionality between Device Controllers, Auto-ID Nodes, and backend systems, a full-fledged solution to the distribution problem needs to support the distribution and replication of functionality and data, requiring the sharing and synchronization of data across multiple nodes. The evaluation and adaptation of distribution and replication strategies developed in

distributed database systems, database caching, distributed event-based systems, and peer-to-peer systems could be a good starting point.

Seamless Integration of Environmental Sensors. Currently, most work in the area of smart items has focused on RFID and Supply Chain Management. To support application scenarios like product life-cycle management (PLM) or transportation, we need to seamlessly integrate other sensors like environmental sensors with RFID technology. From the application perspective, RFID readers and environmental sensors like temperature or light sensors simply provide event sources. From the perspective of the infrastructure, however, they are different. RFID readers are aperiodic event sources, whereas environmental sensors provide a stream of periodic events, that is, discrete readings of the corresponding environmental conditions. Conceptually such a sensor provides a current value for each point in time. The seamless integration of RFID and environmental sensors requires means to represent and resolve this mismatch.

Networked Embedded Systems. Smart items provide small embedded systems capable of independently collecting information from their environment, processing data, and communicating over wireless networks. With advances in memory capacity and processing power, these devices allow the execution of business logic at the periphery of a smart items infrastructure rather than in the middle layers or in a central backend system. Smart items can form entire networks of collaborating devices, thereby increasing reliability (through replication), efficiency, and flexibility. In addition to the question for new appropriate system architectures, efficient ways are required to model, generate, deploy, and manage business functions at the devices. Here approaches developed in the area of grid and peer-to-peer computing could be a good starting point for further research.

Privacy. The use of RFID technology, especially in retail, has raised a lot of discussion regarding privacy. The main concerns here are the possible profiling of customer behavior and the potential to track people. Although this discussion is not a purely technical one, on the technical side mechanisms are required that enable the efficient encoding of tag and sensor information, ensure data security, and allow the disabling of tags at predefined stages in a retail chain. The resulting technology needs to be an integral part of a sophisticated smart items infrastructure. To promote an open and transparent dialogue on RFID technology, SAP is hosting an open community, enabling customers, partners, vendors, and other stakeholders to share information on RFID-related topics. This community provides an open, Web-based platform divided into the areas technology, applications, data security, and privacy. Further discussion of privacy issues related to RFID is provided Chapter 11 by Garfinkel.

3.7 Summary

We have described our Auto-ID Infrastructure which was architected with scalability, flexibility, and usability in mind. Device Controllers allow the processing of event messages close to the periphery of the system; Auto-ID Nodes enable the execution of business logic in the infrastructure and integrate incoming observation messages with backend business processes. We have discussed our practical experiences with different pilot projects and the early ramp-up phase of the SAP Auto-ID component and summarized the main lessons learned. Smart item technology is very likely to change current business and operational processes, which will require changes in the IT infrastructure of many companies. Challenging issues remain that make this area an interesting topic for both hardware and software research.

References

1. I.F. Akyildiz, W. Su, Y. Sankarasubramaniam and E. Cayirci. "A Survey on Sensor Networks". *Computer Networks*, 38: 393–422. 2002.
2. K. Alexander, T. Gillian, K. Gramling, M. Kindy, D. Moogimane, M. Schultz and M. Woods. *IBM Business Consulting Services - Focus on the Supply Chain: Applying Auto-ID within the Distribution Center*, Auto-ID Center. White Paper IBM-AUTOID-BC-002. 2003.
3. C. Bornhövd, T. Lin, S. Haller and J. Schaper. Integrating Automatic Data Acquisition with Business Processes - Experiences with SAP's Auto-ID Infrastructure. *Proc. of the 30th Conference on Very Large Data Bases (VLDB)*. Toronto, Canada. 2004.
4. D.L. Brock. *Integrating the Electronic Product Code (EPC) and the Global Trade Number (GTIN)*, Auto-ID Center. White Paper MIT-AUTOID-WH-004. 2001.
5. S. Clark, K. Traub, D. Anarkat and T. Osinski. *Auto-ID Savant Specification 1.0*. Auto-ID Center. White Paper MIT-AUTOID-TM-003. 2003.
6. EPC Global. *EPC Tag Data Standards Version 1.1 Rev. 1.24*. EPCGlobal Standard Specification. 2004.
7. K. Finkenzeller. *RFID Handbook: Fundamentals and Applications in Contactless Smart Cards and Identification*. 2nd Edition. John Wiley & Sons, London. 2003.
8. C. Floerkemeier, D. Anarkat, T. Osinski and M. Harrison. *PML Core Specification 1.0*, Auto-ID Center Recommendation. 2003.
9. S. Haller and S. Hodges. *The Need for a Universal Smart Sensor Network*. Auto-ID Center. White Paper CAM-AUTOID-WH-007. 2002.
10. U. Hansmann, L. Merk, M.S. Nicklous and T. Stober. *Pervasive Computing: The Mobile World*. 2nd Edition. Springer SMB. Berlin, Heidelberg. 2003.
11. S. Miles, D.L. Brock and D. Engels. *Web Services WAN SIG: Proposals for Engineering the 'Silk Road of the Internet'*. Auto-ID Center. White Paper MIT-AUTOID-WH-04. 2003.

12. A. Thede, C. Schmidt, C. Merz. Integration of Goods Delivery Supervision into E-Commerce Supply Chains. *Second International Workshop on Electronic Commerce (WELCOM'01)*. Heidelberg, Germany. 2001.
13. H. Österle, E. Fleisch and R. Alt. *Business Networking*, 2nd Edition. Springer SMB. Berlin, Heidelberg. 2001.
14. D. Woodsand and J. Word. *SAP NetWeaver for Dummies*. John Wiley & Sons, London. 2004.