Diss. ETH Nr. [16750]

# Infrastructure Support for RFID Systems

A dissertation submitted to ETH Zurich

> for the degree of Doctor of Sciences

presented by Christian Floerkemeier BA MEng, University of Cambridge born February 14, 1975 citizen of Germany

accepted on the recommendation of Prof. Dr. Friedemann Mattern, examiner Prof. Dr. Sanjay Sarma, co-examiner

2006

## Abstract

Radiofrequency identification systems (RFID) have appeared in a variety of applications over the past decades. Examples include keyless entry badges and animal tagging. More recently, low cost, standardized RFID systems have become more widespread in supply chain management and industrial automation. In these domains the technology has the potential to automate many labor-intensive processes because it does not require line-of-sight alignment and RFID tags can be identified over a distance of a few metres by the RFID reader. RFID systems can thus continuously and automatically monitor the flow of goods in many situations, including manufacturers' conveyor lines and and loading and unloading of trucks at dock doors. On a more abstract level, the technology holds the promise to bridge the gap between the real world of physical objects and the virtual world of computer systems. Due to its low-cost and invisible nature, RFID exemplifies the vision of ubiquitous computing, where information technology is seamlessly integrated into everyday objects.

The proliferation of RFID tags and readers also introduces a number of technical challenges, however. While traditional applications for radio-based identification, such as car immobilizers, usually feature no more than a single tag in the read range of a reader and the readers themselves are coarsely deployed and are not networked, this will be different once everyday items are equipped with RFID tags. In this thesis we contend that appropriate infrastructure support is required to maintain adequate performance levels, to protect the privacy of the individual, and to facilitate the management of large scale reader deployments.

In particular, this thesis provides

• transmission control strategies that optimize the throughput of the shared wireless channel and thus improve the identification speed of large tag populations. These transmission control strategies build on earlier work on conflict multiplicity estimation, but have been adapted to suit the characteristics of RFID. The transmission control strategies are evaluated experimentally, but also with the help of a scalable RFID simulation engine that we built and that supports different pathloss, fading, capture, and tag mobility models. Our evaluation shows that the Bayesian transmission strategies improve the throughput and thus reduce the time it takes to identify large tag populations when compared to existing approaches.

- an approach that integrates a subset of the widely accepted fair information practices into the communication protocols between RFID readers and tags. We argue that having RFID readers explicitly declare the scope and purpose of the tag data collection, as well as disclosing the identity of their operators, will allow both consumers and regulators to better assess and control the impact of everyday RFID encounters. Our analysis shows how the fair information principles of collection limitation, purpose specification, use limitation, openness, and accountability can be incorporated in today's RFID communication protocols without significant performance penalties. We also present the prototype of a watchdog tag that allows consumer interest groups and privacy-concerned individuals to judge whether a particular RFID reader deployment complies with the corresponding regulations by displaying and logging the information regarding the data collection broadcasted over the radio channel.
- a middleware design and implementation that addresses the applications' needs for aggregated and filtered data, delivered with appropriate notification latencies, but also taking into account the constraints imposed by passive RFID technology. The latter include the limited available bandwidth, the heterogeneous reader landscape, different types of tag memory organization, and the occurrence of false negative reads. Our implementation shows that an architecture built around a content-based router with a subscription feedback mechanism to the event producers is well suited to addressing the application requirements and constraints imposed by RFID technology.

Taken together, these support mechanisms can help to improve the quality of service in a world of ubiquitous RFID tags and readers.

# Kurzfassung

Während die Radiofrequency-Identification (RFID) in der Vergangenheit vor allem zur Zugangskontrolle und zur Tieridentifikation verwendet wurde, kommt die Technologie heute zunehmend in der Lieferkette und in der industriellen Fertigung zum Einsatz. Mussten betriebliche Informationssysteme in diesen Anwendungsgebieten bisher aufwendig von Menschenhand mit Hilfe von Tastatur und Barcode-Leser mit Daten versorgt werden, so ermöglicht die Verwendung von standardisierter, kostengünstiger RFID-Technologie nun die automatische und kontinuierliche Datensammlung, da kein Sichtkontakt zwischen RFID-Lesegerät und RFID-Transponder nötig ist und die RFID-Transponder über eine Distanz von mehreren Metern ausgelesen werden können. Aufgrund ihrer miniaturisierten Form und geringer Stückkosten verkörpern Funketiketten auf RFID-Basis die Vision des Ubiquitous Computing, die eine Durchdringung der Welt mit Informationstechnologie propagiert.

Die zunehmende Verbreitung von RFID-Transpondertechnik resultiert allerdings auch in einer Reihe von technischen Herausforderungen. Während sich bei klassischen RFID-Anwendungen, wie Wegfahrsperren in Fahrzeugen, häufig nur ein einzelner Transponder im Ansprechbereich des Lesegerätes befindet und die RFID-Lesegeräte nicht unter einander vernetzt sind, so wird dies beim grossflächigen Einsatz von RFID-Transpondertechnik in der Lieferkette anders sein. Dies gilt insbesondere dann, wenn in Zukunft Einzelartikel mit Funketiketten ausgestattet und RFID-Lesegeräte in Warenverteilungszentren und Supermärkten eingesetzt werden. Im Rahmen dieser Arbeit wurde untersucht, wie die Identifikation von RFID-Transpondern im Ansprechbereich des RFID-Lesegerätes beschleunigt werden kann, wie die Privatsphäre des Einzelnen gegen eine potentielle flächendeckende Uberwachung geschützt werden kann und wie das Management einer Vielzahl von RFID-Lesegeräten und der damit erfassten Daten vereinfacht werden kann.

Dazu liefert die vorliegende Arbeit folgende Beiträge:

- Um die Identifikation von RFID-Transpondern zu beschleunigen, sind Übertragungsstrategien entwickelt worden, die den Durchsatz über den Funkkanal optimieren. Diese Übertragungsstrategien bauen auf existierenden Arbeiten zur Vorhersage der Konfliktmultiplizität auf, sind aber vor allem an die besonderen Eigenschaften von RFID angepasst worden. Die Analyse der in dieser Arbeit neu entwickelten bayesschen Übertragungsstrategien zeigt, dass sie im Vergleich zu alternativen Ansätzen den Durchsatz erhöhen und somit die Zeit bis zur vollständigen Identifikation aller Transponder im Ansprechbereich des Lesegerätes verkürzen. Die Übertragungsstrategien werden dabei experimentell, aber auch mit Hilfe eines RFID-Simulators evaluiert, der sowohl Streckendämpfung, Schwund, Nah-Fern-Effekt, und Transpondermobilität modelliert.
- In Bezug auf die Datenschutzproblematik, die durch eine zunehmende Verbreitung der RFID-Technologie hervorgerufen wird, präsentiert die vorliegende Arbeit eine Methode, die die Fair Information Practises, die der aktuellen europäischen Datenschutzgesetzgebung zugrunde liegen, direkt in die Kommunikationsprotokolle zwischen RFID-Lesegeräten und Transpondern integriert. Dabei legen wir dar, wie Lesegeräte, die Informationen, wie z.B. den Zweck der Datensammlung oder die Identität des Datensammlers. über den Funkkanal verbreiten, dazu beitragen, dass sowohl Verbraucher als auch Regulierungsbehörden den Einsatz von RFID-Technik besser kontrollieren können. Die Arbeit zeigt im Detail auf, in welcher Weise die Fair Information Practises, wie Zweckbestimmung, limitierte Nutzung, Transparenz und Verantwortlichkeit in heutige RFID-Kommunikationsprotokolle eingebettet werden können, ohne dass damit signifikante Leistungseinbussen verbunden sind. Wir präsentieren ausserdem ein Watchdog-Tag, das es Konsumentenschützern und den um ihre Privatsphäre besorgten Verbrauchern erlaubt, aufgrund der über Funk verbreiteten Informationen zur RFID-Datensammlung, diese zu kontrollieren.
- Um das Management der RFID-Lesegeräte und der von diesen erfassten Daten zu vereinfachen, ist im Rahmen der Arbeit eine RFID-Middleware entwickelt worden. Diese berücksichtigt sowohl die Anforderungen der RFID-Anwendungen, nach bestimmten Gesichtspunkten aggregierte und gefilterte Daten zu erhal-

ten, als auch die Einschränkungen der RFID-Transpondertechnik. Letzteres schliesst die nur eingeschränkt zur Verfügung stehende Übertragungsbandbreite, die Heterogenität von Lesegeräten, unterschiedliche Datenspeicherorganisationen auf den Transpondern und das Auftreten von Lesefehlern ein. Die Arbeit zeigt auf, dass eine Software-Architektur, die auf einem Content-Event-Router aufbaut, gut geeignet ist, sowohl Anforderungen der Anwendungen zu erfüllen als auch die oben genannten Einschränkungen der RFID-Transpondertechnik zu berücksichtigen.

Zusammengenommen sollten die im Rahmen der Arbeit entwickelten und evaluierten Ansätze helfen, die Dienstqualität in einer Welt allgegenwärtiger RFID-Transponder und Lesegeräte zu verbessern und damit einen Beitrag zur Optimierung industrieller Automatisierungsvorgänge und logistischer Prozesse leisten.

# Contents

1.	Introduction						
	1.1.	1. Motivation					
	1.2.	.2. Contributions					
		1.2.1.	Speeding Up the Identification of Large Tag Pop- ulations	ર			
		122	Addressing BFID Privacy Concerns	5			
		1.2.2.	Managing BFID Systems	5			
	1.3.	Thesis	Outline	6			
2.	Radiofrequency Identification Technology						
	2.1.	Differe	entiation Criteria	9			
	2.2.	Short 2	History of RFID	12			
	2.3.	Operat	ting Principles of Passive RFID Systems	14			
		2.3.1.	Coupling in the Near Field	16			
		2.3.2.	Coupling in the Far Field	19			
		2.3.3.	Physical Layer	21			
		2.3.4.	Medium Access Schemes	22			
3.	Trai	Transmission Control Strategies 2					
	3.1.	Proble	m Statement $\ldots$ $\ldots$ $\ldots$ $\ldots$ $\ldots$ $\ldots$ $\ldots$ $\ldots$ $\ldots$	28			
		3.1.1.	Traffic Characteristics	30			
		3.1.2.	Early Cancellation of the Current Frame	33			
		3.1.3.	Limited Number of Frame Sizes Available	34			
	3.2. Bayesian Transmission Control Strategies			35			
		3.2.1.	Frame-by-Frame Bayesian Updating	35			
		3.2.2.	Slot-by-Slot Bayesian Updating	39			
	3.3.	Relate	d Work	43			
	3.4.	4. Evaluation Methods					
		3.4.1.	Experimental Set-Up	46			
		3.4.2.	Simulation Set-Up $\ldots$ $\ldots$ $\ldots$ $\ldots$ $\ldots$ $\ldots$	47			
	3.5.	Evalua	tion of the Bayesian Transmission Strategies	50			

		3.5.1. Evaluation of the Frame-by-Frame Bayesian Up-						
		dating	53 60					
		3.5.3. Comparison with Frame-By-Frame Transmission	•					
		Schemes	60					
		3.5.4. Comparison with the $Q$ Algorithm	61					
	3.6.	Limitations and Future Work	63					
	3.7.	Summary	64					
4.	Supporting the Fair Information Principles in RFID pro-							
		Dis Chatan and and Dalated Wards	<b>ונ</b>					
	4.1.	From England The Provide And Related Work	08 71					
	4.2.	Fair Information Practices	11					
	4.3.	4.2.1 On any set through Declar and Delies Identification	13					
		4.3.1. Openness through Reader and Policy Identification	14 76					
		4.3.2. Purpose Specification in the inventory Command	70					
		4.3.3. Use Limitation through Collection Types	19					
	4 4	4.5.4. Conection Limitation by Appropriate Tag Selection	82 04					
	4.4. 4 5	Discussion and Future Work	84 86					
	4.0. 4.C		80 07					
	4.0.	Summary	81					
5.	RFI	D Middleware Design 8	39					
	5.1.	Application Requirements	90					
	5.2.	Constraints Imposed by the Characteristics of RFID $$ .	94					
	5.3.	Design Implications	98					
	5.4.	$Implementation - the RFIDStack  . \ . \ . \ . \ . \ . \ . \ 1$	02					
		5.4.1. Data Dissemination, Filtering and Aggregation						
		with the Elvin Message Router $\ldots \ldots \ldots \ldots 1$	03					
		5.4.2. Writing to a Tag $\ldots$ $\ldots$ $\ldots$ $\ldots$ $\ldots$ $1$	06					
		5.4.3. Hardware Abstraction Layer	06					
	5.5.	Evaluation $\ldots$ $\ldots$ $\ldots$ $\ldots$ $\ldots$ $\ldots$ $\ldots$ $\ldots$ $1$	07					
	5.6.	Related Work	09					
	5.7.	Summary	15					
6.	Con	iclusion 11	17					
	6.1.	Speeding Up the Identification of Large Tag Populations $\ 1$	17					
		6.1.1.  Contribution  .  .  .  .  .  .  .  .  .	18					
		6.1.2. Limitations and Future Work	18					
	6.2.	Addressing RFID Privacy Concerns	19					

		6.2.1.	Contribution $\ldots \ldots 119$
		6.2.2.	Limitations and Future Work
	6.3.	Manag	ing RFID Systems
		6.3.1.	Contribution
		6.3.2.	Limitations and Future Work
Α.	RFII	D Simi	ılator 127
	A.1.	Compo	pnents
		A.1.1.	Physical Layer
		A.1.2.	Reader
		A.1.3.	Tag
		A.1.4.	Application Layer
	A.2.	Simula	tor Properties
		A.2.1.	General Simulator Properties
		A.2.2.	Protocol Properties
Bi	bliog	raphy	139

## 1. Introduction

Radiofrequency Identification (RFID) systems hold the promise of eliminating many existing business problems by bridging the economically costly gap between the virtual world of IT systems and the real world of products and logistical units. The proliferation of RFID tags and readers also introduces a number of technical challenges, however. In this first chapter we explain the need for infrastructure support to address these challenges. We outline the main contributions of this thesis and conclude the chapter with an overview of the remaining chapters.

### 1.1. Motivation

Radio-based identification has become common in a variety of applications where access control and robust data carriers without electrical contacts are required. Examples include keyless entry badges, pet tagging, and car immobilizers. More recently, RFID systems with an increased range have begun to find greater use in supply chain management and industrial automation. In these domains radio-based identification has shown itself to be a promising technology to track movements of goods because it does not require line-of-sight alignment and RFID tags can be identified over a distance of a few metres by the RFID reader. Since the RFID tags can thus be automatically identified – without a person scanning the object as in the case of traditional bar codes – industry observers expect significant labor savings from the use of the technology. In addition to these efficiency improvements, the tracking functionality provided by RFID, in conjunction with the appropriate information technology, gives unprecedented visibility to the supply chain. Such transparency can improve supply chain coordination, reduce inventory, and increase product availability. On a more abstract level, RFID holds the promise of bridging the gap between the real world of products and goods and the virtual world of information technology by providing a mechanism to continuously and automatically map the real world onto the virtual world. As such, the adoption

of RFID embodies the vision of ubiquitous computing, which is generally believed to be the next step in the evolution of computing. Its originators foresaw a world in which wirelessly networked computing devices would be seamlessly integrated into our everyday world and would provide useful services to humans in their lives.

The proliferation of RFID tags and readers also introduces a number of technical challenges, however. While classical RFID applications, such as car immobilizers or pet tagging, usually feature no more than a single tag in the read range of an RFID reader and the readers themselves are coarsely deployed and are not networked, this will be different once cases and individual items in supply chain and logistics applications are equipped with RFID tags. A gate reader at a dock door will typically have hundreds, if not thousands of tags in its read range. It will also operate in close proximity to other readers monitoring locations nearby. Since artificially slowing down the loading process or conveyor belt speed is not desirable from a business perspective, the fast identification of RFID tags is an important issue. However, high identification speeds are not only important in supply chain and logistics applications, but in any application featuring many RFID-tagged objects that are simultaneously within range of an RFID reader, e.g. a smart medicine cabinet [49] or even an RFID-enabled card game [50].

While the invisible nature of RFID technology has many benefits from an automation perspective, it is also the cause of some serious privacy concerns. The intended deployment of RFID tags on everyday items epitomizes for many the dangers of an Orwellian future: unnoticed by consumers, embedded microchips in our personal devices, clothes, and groceries can covertly be triggered to reply with their ID and other information, potentially allowing for a fine-grained yet invisible surveillance mechanism that pervades large parts of our lives. While some of the privacy threats associated with RFID are based on invalid assumptions regarding the capabilities of the technology, the feasibility of detecting an RFID tag carrying a unique identifier without line-of-sight over a distance of a few metres nevertheless introduces a set of serious privacy challenges, because data can be collected without any human involvement. Taking the human out of the loop, which is desirable from an automation perspective, means that explicit consent is no longer required from the individual carrying the objects, nor might he or she be aware of the purpose of the data collection or even the data collection itself. It is thus important to develop appropriate approaches that protect the individual carrying tags in his clothes or groceries against the potential misuse of the technology.

The proliferation of readers and tags not only mandates approaches that address these performance and privacy issues, however, but also requires an application-agnostic middleware that helps to manage large reader deployments and aggregates the captured RFID data. In novel application domains, such as supply chain management and logistics, there is no longer a 1-to-1 relationship between reader and application instance. In these domains many readers distributed across factories, warehouses, and distribution centers capture RFID data that need to be disseminated to a variety of applications. Each of these applications has different needs with respect to notification latencies, tag populations of interest, and aggregate types. An appropriate RFID middleware provides this functionality and successfully decouples readers and applications, but also considers the constraints imposed by passive RFID technology. The latter includes the limited available bandwidth, the heterogeneous reader landscape, different types of tag memory organization, and the occurrence of false negative reads.

### 1.2. Contributions

In the previous section we outlined the technical challenges that arise from the proliferation of low-cost, standardized RFID systems. In this section we present the major contributions of this thesis that address these issues. These contributions have also been published in [17, 45, 47, 48, 51, 52].

### 1.2.1. Speeding Up the Identification of Large Tag Populations

There are a variety of approaches to improving the speed at which RFID tags are identified. Most of them target the physical and medium access control (MAC) layer of RFID communication protocols. Examples include increased data transfer rates due to more efficient spectrum usage on the physical layer, and various so-called anti-collision algorithms that aim to minimize the time it takes to identify all tags in the range of the reader. These RFID anti-collision protocols are usually variants of contention-resolving tree algorithms [18] or ALOHA [2]. The performance of ALOHA protocols is known to depend heavily on

a transmission scheme that estimates the (unknown) number of stations transmitting and that controls access to the broadcast channel appropriately [110].

In this thesis, we present two novel Bayesian transmission control strategies

- that explicitly model the variant of framed ALOHA commonly used in RFID protocols, which only allows a limited number of frame sizes and permits the early cancellation of an ALOHA frame,
- and that make no assumption about the statistical distribution of the number of tags present and the tag arrival rate in the range of the reader. This is important because traffic in RFID applications tends to be variable and highly correlated rather than caused by many independent point-to-point transmissions – imagine a dock door with the occasional pallet of more than a thousand tagged items moving through it.

The two transmission control strategies we developed are evaluated experimentally, but also with the help of a scalable RFID simulation engine that we built and that supports different pathloss, fading, capture, and tag mobility models.

Our evaluation shows

- that the two transmission strategies have a higher throughput than existing approaches and when used to control RFID readers thus reduce the time it takes to identify large tag populations,
- that the increased throughput that can be achieved with our transmission schemes comes at the expense of a significant amount of computations, and
- that due to the unknown arrival and departure rates the second of the two Bayesian transmission strategies we developed, which incorporates the feedback from the reader on a slot-by-slot basis, performs significantly better than the strategy that waits for the end of a complete frame before the estimate is updated. This is due to the fact that the former scheme cancels a frame early if the frame size is estimated to be non-optimal.

### 1.2.2. Addressing RFID Privacy Concerns

While the integration of RFID tags in everyday items requires methods to speed up the identification of large tag populations, it also demands approaches that protect the privacy of the individual. For example, some consumer interest groups advocate a complete ban on RFID tags in the public parts of stores, in order to protect the privacy of the consumer. Although this approach will naturally protect the privacy of the individual, it falls short of an optimal solution even from a consumer standpoint, since it is not just the retail store that can benefit from the use of RFID tags, but also the consumer. In this thesis, a compromise is presented, inspired by our everyday lives, where we rarely encounter all-or-nothing tradeoffs but rather engage in meaningful exchanges that conditionally lead us to disclose parts of our personal data in return for more or less tangible benefits. We contend that having RFID readers explicitly declare the scope and purpose of the tag data collection, as well as disclosing the identity of their operators, will allow both consumers and regulators to better assess and control the impact of everyday RFID encounters.

In particular, we show

- that the widely accepted fair information principles of collection limitation, purpose specification, use limitation, openness, and accountability can be incorporated in today's RFID communication protocols without significant performance penalties and
- how a watchdog tag enables consumer interest groups and privacyconcerned individuals to judge whether a particular RFID reader deployment complies with the corresponding regulations by displaying and logging the information regarding the data collection broadcasted over the radio channel.

#### 1.2.3. Managing RFID Systems

The proliferation of RFID readers and tags also introduces the need for RFID middleware solutions designed to manage large deployments and the data captured.

In this thesis, we show

• that the characteristics of passive RFID technology introduce constraints that are unique to the development of middleware for the RFID domain. These constraints include the limited communication bandwidth available to RFID readers, the occurrence of false negative reads, tag memory variations, and the heterogeneous reader landscape.

- that an RFID middleware design that uses a publish/subscribe system featuring full content-based routing and a subscription feedback mechanism to the event producers is well-suited to addressing these constraints and fulfilling the application requirements for filtered and aggregated RFID data.
- the strength and weaknesses of the RFIDStack, our middleware implementation that is based on the above design considerations. The RFIDStack uses the general purpose content-based router Elvin [112], which features a subscription feedback mechanism and a subscription language with predefined operators. We illustrate that this implementation is well-suited to addressing application needs and technology constraints. However, we also discuss characteristics of the content-based event router in its current version that limit the scalability of this approach.

### 1.3. Thesis Outline

In the second chapter, we provide an overview of RFID technology. We discuss the components of an RFID system and outline the fundamental principles of operation that influence characteristics such as range and identification speed – both of which represent important parameters in the remainder of this thesis. This chapter also compares different medium access schemes used in RFID protocols and motivates the work on transmission control strategies in the next chapter.

In the third chapter we propose two transmission control strategies that aim to increase the identification speed of large tag populations. After outlining the requirements that transmission control strategies have to meet, given the characteristics of RFID, we present two Bayesian transmission schemes. We continue with a description of related work, before we present the methods with which we compare the two Bayesian schemes against other existing approaches. We present in particular the RFID Simulation engine that we developed to evaluate the different transmission strategies. The chapter concludes with a discussion of the experimental and simulation results. The fourth chapter deals with the privacy concerns that arise from the adoption of RFID tags with a significant read range in today's retail environment. We discuss fair information principles and show how these can be integrated into RFID communication protocols. We also outline how this modification to existing protocols allows consumer interest groups and privacy-concerned individuals to judge whether a particular RFID reader deployment complies with the corresponding regulations, through the use of a watchdog tag. The chapter concludes with a discussion of the benefits and the limitations of the approach presented.

In the fifth chapter, we present the RFIDStack, a middleware design that aims to address both application needs and RFID technology constraints. We begin by analyzing application requirements and discussing RFID constraints – both of which make the development of middleware for RFID systems unique. We continue by discussing the resulting design considerations and a description of our middleware design and implementation – the RFIDStack. The related work section in this chapter describes in particular the work of the Auto-ID Center, an industry sponsored research program that developed a network infrastructure to manage the RFID data captured, and its predecessor EPCglobal.

# 2. Radiofrequency Identification Technology

Radiofrequency identification (RFID) is a generic term that is used to refer to a class of short range telemetry applications in which the purpose is to read the identity of a transponder attached to an object as it passes within the range of a reader device. Although all RFID systems consist of transponders, which store the identity of an object in the form of a serial number, and a reader, which reads from and possibly also writes to the tags, a wide variety of different systems have been developed. These can be classified in a number of ways, e.g. by operating principles, performance measures, and rationale for usage. The purpose of this chapter is to give an overview of the RFID design space and to provide the technical background for the remaining chapters of this thesis. We begin with a brief overview of key performance indicators, before presenting the history and applications of RFID. We continue with a discussion of the fundamental operating principles of the most common category of RFID systems, in which the transponders contain a microchip, but no battery – so-called passive tags. This analysis focuses in particular on two key performance indicators: range and identification speed – both of which are important in the context of this thesis.

## 2.1. Differentiation Criteria

The vast majority of RFID systems use a microchip in the tag: this microchip includes basic modulation circuitry and non-volatile memory. There are, however, also RFID systems available that operate without an integrated circuit – they are often referred to as chip-less – and are for example based on surface acoustic wave delay lines [26] or resonant circuits [43]. There are three broad categories of microchip-based RFID systems – passive, semi-passive and active systems. Passive tags do not contain an internal source of operating power, but are powered remotely by the reader (cf. Figure 2.1). They rely on communication



Figure 2.1.: Fundamental operating principle of passive RFID systems. Passive RFID systems are composed of two components – a reader and a passive tag. The tag is composed of an antenna and a microchip that includes basic modulation circuitry and non-volatile memory. The tag is energised by a time varying radiofrequency signal that is transmitted by the reader.

by means of reflected power for data transmission, since they have no active transmitter. Active tags have their own transmitter and an internal power source. There are also semi-passive tags, which use a battery to power the microchip, but rely on communication by means of reflected power.

The performance of an RFID system can be further characterized by the following properties:

- **Range** The maximum distance between reader and tag at which data can be communicated varies between less than 1cm and hundreds of meters. Read ranges larger than 10 m can today only be achieved with semi-passive or active tags.
- Variation in range under non-ideal conditions The range depends among other factors on the operating environment. In particular, the range of some of the passive RFID systems can be severely reduced under non-ideal conditions, e.g. in the vicinity of metal and other tags.
- **Data transfer rate** The rate at which data can be transferred from reader to tag and vice versa varies between a few and hundreds of

kilobytes per second.

- **Identification speed** The speed at which multiple tags in the range of the reader can be identified varies from a few to hundreds of tags per second. There are also some RFID systems that provide no means of detecting multiple tags in the range at the same time.
- **Tag memory** The type of memory used, the memory organization, and the total size vary between different tag types. All tags have in common that they carry a unique identifier. In some cases this identifier is a serial number that identifies the tag; in other cases the identifier describes the object to which the tag is attached. Tags usually feature read-only or read-write memory. Read-only tags carry no data except the identifier, which is programmed during microchip fabrication. There are also read-write tags carrying a unique serial number only. These are usually referred to as writeonce-read-many-times tags because the memory is locked after the initial write. There are also tags that feature additional random access memory for application data.
- Lifetime The operating life of semi-passive and active tags is limited by the battery. While the operating life of passive tags is in principle unlimited, it can be reduced in practice by mechanical wear.
- Tag and reader form factors Readers and reader antennas come in a variety of form factors including handheld, desktop, portal and tunnel readers as shown in Figure 2.2. Passive tags are physically small in size – typically less than 15 cm in diameter – and do not destroy the aesthetics of the tagged objects. Battery powered tags are usually significantly bulkier.

There are many interdependencies among the above parameters. The range of a passive RFID system is for example influenced by tag and reader form factor and desired data transfer rate. In Section 2.3 we will see that the range is also dependent on system parameters – such as frequency of operation, power consumption of the microchip, and reader sensitivity – and application parameters – such as orientation of the tag antenna with respect to the reader antenna, country of operation, and material in the vicinity. The distinction between passive and active systems is more straightforward. Active tags usually provide superior range at the expense of a limited operating life, higher costs, and larger



Figure 2.2.: Reader form factors. Handheld, gate and tunnel readers are shown.



Figure 2.3.: *Passive tag form factors.* LF, HF, and UHF tags are shown from left to right (not drawn according to size).

tags, while passive tags have in principle an unlimited lifetime, but a limited range.

### 2.2. Short History of RFID

Radio-based identification technologies have been around for decades. Friend or foe identification systems developed at the end of the Second World War have some similarity with today's RFID systems. The German air force used a roll-over manoeuvre to indicate their identity to their radar operators [3], which essentially constitutes an example of communication by means of reflected power – a communication principle published by Stockman in the academic literature in 1948 [118]. By rolling over at a predetermined signal the German pilots were effectively changing the polarization of the radar reflections, which was picked up by their own ground radars. The British air force soon afterwards developed an active identify friend or foe system that provided identification functionality, but did not use the backscatter principle [74]. The covert listening device developed by Theremin [130] also has some similarity to the underlying coupling and communication principle of today's RFID systems, although it was not conceived for identification purposes. Theremin's bug used inducted energy from radio waves to transmit an audio signal. The covert listening device changed the impedance of the antenna according to the sound in the vicinity and the reflections of the incident UHF signal were thus modulated accordingly. An interrogator could power the listening device remotely and receive the backscattered data. It thus constitutes an early example of a truly passive device that transmitted data via backscatter to endow the device with unlimited operational life.

The two decades after the Second World War were an era of exploration of radio-based identification techniques following technical developments in radar and radio. Commercial activities began in the 1960s and continued in the 1970s, when government laboratories, universities, and companies built the first RFID systems. Important developments included microwave-based systems such as [53, 70, 117] and those that exploited the principle of inductive coupling for object identification [60, 90]. Good overviews can be found in [24, 43, 74]. Application domains included automated toll payment systems and animal tagging. In the late 1980s RFID systems that offered faster data transfer rates were developed, which were increasingly used for payment systems and access control to buildings and in cars. The next generation of passive RFID technology was developed in the early 1990s by researchers at IBM. The micro-wave system offered not only high data rates, but also a significantly larger read range. It featured tags with a single integrated circuit – a capability previously limited to inductively coupled tags [74]. Because of the lack of standardized communication protocols, the high prices of the tags, and limited world-wide availability of the frequency band used, the technology prospered initially only in niche applications.

These long range, passive RFID systems received a boost when the Auto-ID Center, an industry-sponsored research programme headquartered at Massachusetts Institute of Technology, was initiated. Between 1999 and 2003, the Auto-ID Center gained the support of more than a hundred large companies – including large end-user companies such as Wal-Mart, Proctor and Gamble, Metro, and key RFID vendors – to develop open standards for low-cost tags and readers and the supporting network infrastructure [105, 106]. When the research program closed in October 2003, the technology was licensed to the Uniform Code Council, which created EPCglobal, as a joint venture with EAN International, to commercialize the technology. The efforts of the Auto-ID Center resulted in plans by some of the biggest retailers in the world - Wal-Mart, Tesco, Metro - and the US Department of Defence to use the passive, long-range RFID systems to track goods [99]. Supply chain management applications in particular benefit from the standardized, low cost RFID technology, since it allows the automatic tracking of objects as goods move across the supply chain [27]. RFID systems can continuously and automatically monitor the flow of goods in many scenarios, including manufacturers' conveyor lines, loading and unloading of trucks at dock doors, and handling of palletized goods within warehouses and distribution centres (cf. Figure 2.4). While labour savings from avoidance of manual bar code scanning make the technology attractive for this use alone [22, 79], there are a number of other economic benefits commonly associated with the use of this automatic identification technology. These include reduced inventories [42], theft prevention [5], out-of-stock monitoring [57], and counterfeit detection [115]. On a more abstract level, RFID technology holds the promise of bridging the gap between the real world of products and goods and the virtual world of enterprise resource planning and supply chain management systems. It enables a continuous and automated mapping of the real world, providing increased visibility and transparency [41].

## 2.3. Operating Principles of Passive RFID Systems

In the previous section, we saw that passive RFID systems are well established in a variety of applications such as preventing car theft, collecting road tolls without the need for vehicles to stop, gaining keyless entrance to buildings, and animal tagging. The recent development of standardized, low cost tags with improved range is today triggering the use of the technology in new applications. In these, e.g. supply chain management, everyday items are equipped with disposable tags that allow for automatic and invisible tracking. The proliferation of such



Figure 2.4.: Supply chain applications at distribution centers and in retail stores [9].

low-cost tags on everyday items means, however, that increasingly often large tag populations need to be identified and the privacy of the individual carrying the tags – for example embedded in his clothes and groceries – is at risk. In order to assess these problems and to provide an adequate background for the approaches proposed in the remaining chapters of this thesis, we will discuss the operating principles of passive RFID with a strong focus on the parameters that influence range and identification speed.

Passive RFID systems are commonly categorised according to the frequency of operation. RFID systems operate in unlicensed radiofrequency bands – most of which are ISM bands. Common frequency bands for RFID systems are <135 kHz (LF), 13.56 MHz (HF), 915 MHz (868 MHz in Europe) (UHF) and 2.45 GHz (MW). Electromagnetic field theory predicts that the distance  $r = \lambda/(2\pi) = c/(2f\pi)$  is of significance in determining the nature of the fields surrounding a magnetic or electrical dipole. Within this distance – in the so-called near field – the dominant fields are reactive and quasi-static. Outside of this distance, the dominant fields are those associated with energy propagation by electromagnetic waves. This region is known as the far-



Figure 2.5.: RFID frequency bands [73]. RFID systems operate in unlicensed frequency bands. The most common ones are <135 kHz (LF), 13.56 MHz (HF), 915 MHz (868 MHz in Europe) (UHF) and 2.45 GHz (MW).

field region. In the following two subsections, we will discuss how the frequency of operation dictates the type of coupling and consequently maximal range and susceptibility to variation in range.

#### 2.3.1. Coupling in the Near Field

LF and HF systems, where the near field encompasses the operating range, realize coupling through quasi-static fields. In principle, both inductive and capacitive coupling between reader and tag are possible. In practice, nearly all LF and HF RFID systems rely on inductive coupling to power the tags. For this purpose, the reader's antenna coil generates a strong alternating electromagnetic field, which penetrates the cross-section of the tag coil area, which is some distance away from the coil of the reader (cf. Figure 2.6). The induced voltage in the transponder's antenna coil is rectified and serves as the power supply for the microchip. On the tag there is usually a resonance circuit, whose frequency is tuned to the operating frequency of the RFID system. At resonance, the induced voltage produced across the tuned tag will thus be significantly enhanced compared to frequencies outside the resonant bandwidth. The arrangement of the tag coils represents a transformer where there is only a very weak coupling between the two coils.

Successful communication between the reader and the tag requires that the supply voltage across the microchip is high enough for the operation of the data carrier and that the signal to noise ratio of the tag reply is high enough to be decoded by the reader. The minimum interrogation field strength required to power the tag depends on the frequency of operation, tag coil area, the number of windings of the



Figure 2.6.: Inductively coupled RFID system. The reader generates a changing magnetic field that penetrates the tag coil. The integrated circuit in the tag is supplied by the voltage induced in the tag coil.

tag coil, and the minimum supply voltage before rectification. The required integration field strength is at a minimum when the transmission frequency of the reader corresponds to the resonance frequency of the tag. Manufacturing tolerances and interaction with nearby transponders and other metallic objects often lead to a deviation in the tag resonance frequency [40]. Such a deviation from the transmission frequency of the reader will result in a higher minimum interrogation field strength. The result is a lower maximum activation range.

The magnetic field strength at a given position depends on the reader antenna geometry and the current in the reader coil. The choice of current and antenna geometry is constrained by radio regulations that specify maximum field strength. At HF, these electromagnetic compatibility regulations are enforced in the far field and allow only minimal radiations (cf. Figure 2.7). However, near and far fields scale differently with distance. In particular, the near field energy density per unit volume decreases as the inverse sixth power from the antenna. This means that very close to the reader antenna substantial energy densities can be obtained, but they diminish rapidly as distance increases. The result is that under current regulations the operation of HF systems is confined to the near field and short distances (< 1 m), unless large label antenna sizes are chosen. In practice, the maximum current through a larger reader coil is also limited by the large voltage across the reactive components, which needs to be sustained in the series resonance circuit at the reader.

So far, we have considered the properties of inductively coupled systems that determine the maximum energy range of the system. How-



Figure 2.7.: Magnetic field strength limits at 10 metre measurement distance for the 6.78 MHz and 13.56 MHz bands [38]. The power levels of the f1 band are reserved for RFID and Electronic Article Surveillance Applications only.

ever, successful data transmission from the tag to the reader requires not only a powered tag microchip, but also the successful reception of the tag replies. Tag-to-reader communication takes place via a technique called load modulation. If a tag is placed within the alternating magnetic field of the reader coil, the tag will draw energy from the magnetic field. This results in a voltage change across the reader coil. By switching a load resistor on and off at the tag coil, the amplitude of the voltage change can be modulated. Due to the weak coupling between the two coils, the voltage changes across the reader coil, which represent the tag reply, are nevertheless smaller by several orders of magnitude than the amplitude of the carrier wave in the reader coil. To facilitate decoding, load modulation by means of a subcarrier is frequently used at HF. Whether a tag reply can be successfully received by the reader depends on a number of factors, such as the sensitivity of the receiver, the coding and modulation of the tag reply, and the conditions of the operating environment (electrical noise, etc.). The interested reader is referred to [40] for a quantitative analysis.

#### 2.3.2. Coupling in the Far Field

The analysis in the previous section showed that passive tags operating in the near field have a limited read range, typically less than 1 m, unless large label sizes are employed. In this section, we will see that RFID systems operating in the far field can achieve superior reading distances.

Given reader and tag antennas, the ratio of power received by the tag antenna,  $P_{tag}$ , to power input to the reader antenna,  $P_{reader}$ , is given by the Friis Transmission Equation

$$\frac{P_{tag}}{P_{reader}} = G_{tag}G_{reader} \left(\frac{\lambda}{4\pi d}\right)^2 \tag{2.1}$$

where  $G_{reader}$  and  $G_{tag}$  are the antenna gain of the reader and tag antennas, respectively,  $\lambda$  is the wavelength, and d is the distance between reader and tag. The antenna gains are with respect to isotropic transmitters. However, the above equation applies only under the following ideal conditions:

- The antennas are in unobstructed free space.
- $P_{tag}$  is understood to be the available power at the tag antenna terminals. It will not be fully delivered to the tag microchip unless the microchip is impedance matched to the antenna.
- $P_{reader}$  is understood to be the power delivered to the reader antenna. It will not equal the reader output power unless the reader impedance is impedance matched to the reader antenna.
- The antennas of reader and tag are correctly aligned and polarized.

The above equation for free space propagation shows that the attenuation is frequency dependent. Thus assuming that all other properties are equal, a UHF RFID system will have a larger range than a system operating at MW. The maximum theoretical activation range is in practice decreased by three types of losses:

**Absorption** Since most RFID systems are deployed indoors and there is not always a line-of-sight path between tag and reader, the free-space assumption is usually not valid. The electromagnetic wave supplying tags with power is for example completely reflected by perfect conductors and partially reflected by perfect dielectrics [98]. Lossy dielectrics between reader antenna and tag antenna will also absorb some of the incident radiation. The result is that in practice, depending on the material between tag and reader, the read range is significantly less than predicted by Equation 2.1.

- Multipath fading Even if there is line-of-sight between reader antenna and tag, small-scale fading effects can increase and decrease the read range. Multipath fading is caused by interference between two or more versions of the transmitted reader signal, which arrive at the receiver at slightly different times [98]. These multipath waves combine at the receiver to result in a signal which can vary widely in amplitude and phase. Due to the constructive and destructive effects of multipath waves, a tag moving past a reader antenna can pass through several fades in a small period of time. If the tag passes through such a field null, it will lose power and possibly also its state. The typical rapid variations in the received signal level due to small-scale fading as a tag is moved through a small corridor past a reader antenna are shown in [88].
- **Polarization losses** The activation range is further significantly reduced by polarization losses, since the precise orientation of tags relative to the reader antenna is usually not known. Even when the reader is transmitting with a circularly polarized antenna, the transponder fails to be adequately powered when the axis of the tag dipole antenna is aligned with the propagation direction of the emitted electromagnetic wave. Circular polarized antennas also introduce an additional loss of 3dB. A promising approach to alleviate this orientation dependence is the use of two tag antennas that are orthogonally polarized and attached to the same microchip [120].
- **Impedance mismatch** The activation range predicted by the Friis transmission equation is in practice further reduced by impedance mismatch between tag antenna and microchip.

In the section on near-field coupling we saw that the range of an RFID system depends not only on the maximum distance at which the tags remain powered, but also on an adequate signal-to-noise ratio of the tag signal at the reader. RFID systems operating in the far field rely on communication by means of reflected power [69]. The microchip

modulates the reflection characteristics of the antenna by altering the load connected to the antenna. This changes the proportion of incoming power that is reflected by the antenna. The ratio of the power backscattered by the tag and the power received by the reader can again be estimated using the Friis transmission equation. Given that today's readers have a sensitivity of -90 dBm and tags require around  $30-50\mu$ W power, far field RFID systems are limited by the forward channel.

#### 2.3.3. Physical Layer

In the previous section, we discussed the basic operating principles of RFID systems operating in the near and far field. In particular, we focused on the properties that influence read range. In this section, we concentrate on the characteristics that determine the data transfer speed of an RFID system. We also provide an overview of existing radio regulations that limit bandwidth, transmit power, and transmission duration because of their significant effect on communication in RFID systems. The spectrum mask specified in the radio regulations has a strong influence on the data transfer speeds that can be achieved with RFID systems. The sidebands that result from the modulation of the carrier wave by the reader are limited to the maximum field strength and power levels shown in Figure 2.7. The actual data rate in a protocol is however a complex tradeoff, also influenced by other factors than the radio regulations in the particular frequency band: coding and modulation techniques; the required level of robustness against noise; and then the desired operating range. Some RFID protocols even support multiple data rates to achieve optimum performance, given the partly conflicting constraints [32, 62]: in a low-noise environment with wide bandwidth availability, such as under US regulations in the UHF frequency band, the reader would choose a high data rate, while in another country or in an environment with considerable interference, a lower data rate would be more appropriate to communicate with the tags. Since RFID tags are only passive transmitters that backscatter some of the power they receive from the reader, the constraints for the return link are slightly different than for the forward link: return signals are typically a few orders of magnitude weaker than the signal on the forward link; bandwidth for tags replies is not always regulated; and tag-to-reader data rates can be higher. The high performance mode of the HF RFID protocol specified in ISO 18000-3 [64] takes advantage of the latter by not only having tags reply at a higher data rate, but also on 8 different frequency channels. In other frequency bands, such as UHF, such out-of-band tag replies are prohibited by radio regulations, e.g. ETSI EN 302 208 in Europe [35]. Novel coding and modulation techniques can also lead to significantly improved data rates. The high-performance mode of the ISO Standard 18000-3 is a further good example. It uses phase-jitter-modulation on the forward channel and is thus able to use the little available bandwidth in the 13.56 MHz (HF) band (cf. Figure 2.7) more effectively.

### 2.3.4. Medium Access Schemes

At the physical level, the data rates on both the forward link (reader-totag) and the return link (tag-to-reader) heavily influence the speed at which tags can be identified. They determine how long the transmission of a command to the tag and its corresponding reply (e.g. with its ID) take. In RFID applications that simultaneously feature multiple tags in the read range of the reader, the identification speed is also influenced by the choice of medium access (MAC) algorithm, which controls access to the shared radio channel. This section will give an overview of the most common MAC approaches in RFID protocols and explain our base protocol selection – framed Aloha – for which we developed our efficiency improvements in the subsequent sections.

The general problem faced by an RFID reader is that a number of unknown tags are present in the range. Each of these is marked with a unique ID from a large address space. In order to inventory the tag ID and to retrieve additional data stored in the tag memory, the reader needs to singulate each individual tag. The specific medium access problem encountered in RFID applications is characterized by:

- the limited tag resources, which imply that the tags cannot sense whether another tag is transmitting and whether a collision has occurred.
- the restrictive radio regulations that limit the amount of reader signalling and also forbid out-of-channel tag replies in some frequency bands.

Due to the limited resources on the tags and the limited bandwidth available, most RFID systems use time division multiplexing techniques [58]. An exception to this rule is the communication protocol of



Figure 2.8.: Framed Slotted ALOHA.

ISO 18000-2 Mode B, which uses a combination of TDMA and FDMA, as mentioned above. Since the tags present in the read range represent only a tiny subset of the complete number of the global tag population, polling techniques that check for the presence of each possible ID are prohibitively time consuming. More commonly found MAC schemes include variants of ALOHA and tree search; these are discussed below:

- ALOHA: The most basic RFID MAC scheme is a variant of pure ALOHA. Tags announce themselves to the reader by sending their ID once they detect the presence of a reader. The reader, which transmits a continuous unmodulated carrier signal, is detected by the fact that the tag powers up. In this tag-talks-first (TTF) protocol the reader only listens to the tags that continue to transmit their IDs at random intervals. The advantage of this scheme is that these read-only TTF tags are extremely simple, resulting in a small chip and low cost tags. The reader design is also straightforward, since there is no need for any synchronization or command transmission. The main drawbacks of this protocol are the fixed broadcast probability of the tags, the low throughput associated with this scheme and the restriction to read-only tags. The problem of a lack of control over the broadcast probability has been addressed in subsequent versions of this protocol, which include some reader signalling to mute tags that successfully replied and measures to implement carrier sensing [28]. The low channel utilization associated with ALOHA remains an issue for applications with large tag populations.
- Framed ALOHA: Framed ALOHA, a variant of slotted ALOHA, is also a commonly used MAC scheme in RFID protocols. Figure 2.8

shows the slotted frame structure. The reader usually announces the start of the next frame to the tags and informs them about the number of slots available in that frame. The tags then choose one of the slots at random for their replies. The major feature distinguishing framed ALOHA from standard slotted ALOHA is that a tag is permitted to transmit its ID at most once per frame. The framed ALOHA implemented in RFID protocols is usually a slightly modified version of the one described in the networking literature. The reader usually transmits a short synchronization signal after each slot, indicating the start of the next one. Often, a successful receipt of a tag is also acknowledged immediately after the slot. The tag will then enter an inventoried state and will not reply in subsequent frames unless it receives a reset signal.

• Tree-Search Protocols: Contention resolving tree search protocols are another common category of multiple access schemes used in RFID communication protocols. One can broadly distinguish a number of different kinds. There are those tree-search protocols that require no memory on the tag. The reader transmits a prefix and all tags where the ID matches the prefix respond with their ID [78]. If the replies are synchronized – an assumption that is not valid for all RFID systems – the reader can even determine the position of bit collisions. In other systems, the tags keep their state in memory (active, temporarily inhibited, inventoried) and the reader only needs to send the next bit as it traverses the tree of possible IDs [8].

The use of a reservation system [14] is a common way to substantially improve the throughput of multi-access channels in RFID protocols [58]. In [32], framed ALOHA is only used as the contention mode for short packets that reserve longer noncontending slots for the transmission of the longer unique tag ID. Let us assume here that tag ID packets require one time unit each for transmission and that reservation packets require  $\nu = 1/8$  units (neglecting overheads of reader commands and turn-around-times). According to [14], the maximal throughput U in tag IDs per time unit achievable in such a scheme is then given by the following equation, where  $U_r = e^{-1}$  represents the maximum throughput of framed ALOHA:

$$U = \frac{1}{1 + \frac{v}{U_r}} \approx 0.75 \tag{2.2}$$
All of the above MAC schemes represent a reasonable compromise, given the partly conflicting requirements. MAC schemes that achieve a higher throughput in less variable environments, such as TDMA, are not appropriate given the traffic characteristics of RFID applications. The protocols based on ALOHA suffer from their non-deterministic nature, i.e. the reader can only detect all tags with a probability that approaches unity. More serious in practice is the fact that they require some kind of control to achieve a decent throughput. Since the true number of tags powered and ready to transmit is usually unknown, transmission control strategies are required to estimate the true number of tags and then transmit the corresponding broadcast probability – encoded in the frame size – to the tags. In the next chapter, we will discuss two novel approaches that address this issue.

# 3. Transmission Control Strategies for Framed-ALOHA Based RFID Protocols

While traditional RFID applications usually feature no more than a single tag in the read range of an RFID reader, this will be different once cases and individual items in supply chain and logistics applications are equipped with RFID tags. A gate reader at a dock door will typically have hundreds, if not thousands of tags in its read range. Since artificially slowing down the loading process or conveyor belt speeds is not desirable from a business perspective, the fast identification of RFID tags is an important issue.

In the previous section, we illustrated that the speed with which large tag populations can be identified depends on a scheme that controls the transmission probability of the tags. The main contributions of this chapter are two transmission control strategies for RFID communication protocols, based on framed ALOHA, that ensure that the limited communication bandwidth available can be used efficiently to identify large tag populations. The proposed transmission control schemes build on earlier work by [100] on Bayesian broadcast strategies, but have been adapted to suit the characteristics of RFID. They have in particular been designed for framed ALOHA, a MAC scheme frequently used in RFID communication protocols [32, 63, 95]. The transmission strategies also make no assumption about the statistical distribution of the number of tags in the read range. This is important because the number of tags in RFID applications tends to be variable and the traffic is highly correlated rather than caused by many independent point-topoint transmissions – imagine a dock door with the occasional pallet of more than a thousand tagged items moving through. The chapter also presents experimental evidence and simulation results showing that the proposed transmission control schemes provide a superior throughput when compared to existing approaches.

The chapter is organized as follows: we first identify the main factors

influencing tag estimation and throughput in RFID systems, which include the unusual traffic characteristics of RFID applications and the particularities of the variant of framed ALOHA commonly used in RFID. Based on this analysis, we then present the two Bayesian transmission control schemes which specifically take these unique characteristics of RFID systems into account. Before we conclude the chapter with an evaluation of the transmission schemes, we present related work.



Figure 3.1.: Framed Slotted ALOHA. The reader initiates a frame with a start-offrame (SOF) signal that broadcasts the frame size. Because of the large proportion of collisions in the first frame, the frame size is increased to 8 slots in the second frame. The figure also shows the reader feedback that distinguishes slots with no, a single, or more than a single tag reply (collision).

# 3.1. Problem Statement

The previous chapter showed that there are a variety of approaches to improving the speed at which RFID tags are identified. Most of them target the physical and medium access control layer of RFID communication protocols. Examples include increased data transfer rates due to more efficient spectrum usage on the physical layer and various socalled anti-collision algorithms that aim to minimize the time it takes to identify all tags in the range. These RFID anti-collision protocols are variants of contention-resolving tree algorithms [18] or ALOHA [2]. In framed ALOHA, which is used in a number of RFID communication protocols [32, 63, 95], the reader begins its interrogation round by announcing the frame size to the tags (cf. Figure 3.1). Each tag se-



Figure 3.2.: Number of Timeslots vs. Expected Throughput in Framed ALOHA.

lects one of the available slots at random and transmits a (temporary) identifier. According to [111], the expected throughput U of framed ALOHA with N tags and L slots in a frame is given by:

$$U(N,L) = \frac{N}{L} \left(1 - \frac{1}{L}\right)^{N-1}$$
(3.1)

It is evident from the above equation that the throughput depends on the appropriate choice of frame length L, given the number of tags N in the read range. Figure 3.2 shows the well-known upper bound of the throughput of  $e^{-1}$  (as N becomes large) that is characteristic for slotted ALOHA and also applies to framed ALOHA. The maximum throughput occurs at L = N.

Since the number of tags present is usually not known, the performance of framed ALOHA depends on a transmission scheme that estimates the (unknown) number of tags based on feedback from the reader and chooses a corresponding frame size. This feedback from the reader comprises the number of slots in which no, a single, and more than one tag replied as shown in Table 3.1. The latter is referred to as a collision because the data received by the reader are garbled.

The total number of tags  $N_t$  that reply in a frame at time t is given by

Feedback type	Description		
Empty Slot	no tag is transmitting during the slot		
Single Reply Slot	exactly one tag uses the channel and it		
	reply is successfully received. It will not		
	reply again in the subsequent frames.		
Collision slot	two or more tags transmit in the same slot		
	but none of the individual tag replies can		
	be reconstructed at the reader. All tags		
	need to retransmit in the next frame.		

Table 3.1.: Ternary feedback model.

$$N_t = \gamma c + s \tag{3.2}$$

with

$$\gamma \ge 2 \tag{3.3}$$

where s and c denote the number of single tag replies and collisions, respectively. The exact value of  $\gamma$  is usually unknown because the reader cannot detect how many tags replied if there are two or more tag replies. The tags involved in a collision are backlogged and retransmit their identifier in the subsequent frame. The number of tags  $N_{t+1}$  that transmit in the subsequent frame is the sum of the backlogged tags that remain powered and newly arriving tags.

$$N_{t+1} = \gamma c - n_D + n_A \tag{3.4}$$

where  $n_D$  and  $n_A$  denote the departing tags and newly arriving tags, respectively.

Before presenting two novel transmission strategies, we describe the characteristics of the RFID domain that introduce novel constraints. We focus in particular on the traffic characteristics, the possibility to interrupt a frame before the last slot of the frame is reached, and the limited available frame sizes.

### 3.1.1. Traffic Characteristics

In RFID applications, the tag arrival and departure rates are influenced by application parameters and RFID system design choices. Figure 3.3(a) shows a medicine cabinet that is equipped with an HF RFID system. At time  $t_0$  the reader is switched on and all tags in which a sufficient operating voltage is induced are instantly powered (cf. Figure 3.3(b)). Since the tags are not moved until the end of the identi-



(b) Number of Tags Powered vs. Time

Figure 3.3.: *RFID enabled Medicine Cabinet [49].* The RFID reader is switched on at time  $t_0$  and the tags attached to the products are powered. The tags remain powered as the identification progresses because they are not moved.

fication round, the number of tags powered remains constant over the identification period.

In supply chain operation, UHF readers are often used in a portal configuration as shown in Figure 3.4(a). Tagged objects are placed on a pallet and moved past the reader antenna. In the previous chapter, we saw that UHF RFID systems are affected by the strong fading component characteristic for indoor wireless channels in this frequency band [98], which leads to frequent field nulls. Since tags do not carry a battery, they will frequently lose power and possibly also their state, as they move past the reader. Figure 3.4(c) shows the received signal power vs. position for a single tag that is placed inside the pallet [89].



(a) Portal Reader [89] (b) Pallet with Tagged Cases [89]



(c) Received Power as Single Tag Passes Reader Antenna [89]



(d) Number of Tags Powered vs. Time

Figure 3.4.: Bulk identification in a warehouse application. The tagged cases are moved through a portal to which an RFID reader antenna is attached. As the tags move past the reader antenna, the received signal strength varies (cf. Figure 3.4(a), causing the tags to lose power. The exact number of tags powered at any position of the pallet depends on the detailed set-up.

The number of tags that are powered at a certain moment of time



Figure 3.5.: Framed ALOHA in the EPCglobal Class 1 Generation 2 RFID Protocol (Source: EPCglobal). The reader initiates a frame with a Query command that includes the number of slots in the frame. Each subsequent slot in the frame is initiated with a QueryRep command. The reader can thus start a new frame after any slot by issuing a new Query command.

thus varies significantly due to the movement of the tags past the reader and the frequent field nulls within the interrogation volume (cf. Figure 3.4(d)). The actual arrival and departure rates depend on a number of factors, including antenna properties, multipath effects, material properties of the tagged objects, tag orientation, density and speed.

Common to both scenarios is that the number of tags present is unknown initially. In the first scenario, where medication in a cupboard is inventoried by an HF reader, there are, however, no changes to the number of tags present as the identification progresses after the initial step input. In the second scenario, on the other hand, the number of tags powered is initially unknown, but this number also varies considerably as tags arrive and depart during the identification process.

#### 3.1.2. Early Cancellation of the Current Frame

Framed ALOHA usually means that acknowledgements are only sent after the end of each frame. This is however not true for the RFID domain. There is usually a reader command after each slot (cf. Figure 3.5). Transmission control schemes consequently do not have to wait until the end of a frame to change broadcast probabilities by set-



Figure 3.6.: Maximum throughput vs. number of tags replying in a frame. The figure shows that the optimum throughput which can be achieved if the number of tags transmitting is known will be reduced to 35% for some tag estimates, if only frame sizes that are powers of two are available.

ting the appropriate frame size. They can simply cancel a running frame and initiate a new one.

Figure 3.5 also illustrates the reservation system of the EPCglobal UHF Class 1 Generation 2 RFID Protocol [32]. There is only contention for the short slot following a Query or QueryRep command, in which the tag replies with a 16 bit long random number (RN16). The transmission of the unique tag identifier (EPC) and additional data, such as CRC and PC, happens after the transmission time has been reserved.

#### 3.1.3. Limited Number of Frame Sizes Available

The RFID domain not only introduces unusual traffic characteristics; the variant of framed ALOHA used in RFID protocols such as [32, 63, 95] also differs from the framed ALOHA commonly described in the networking literature [111] because not all frame sizes are available. To reduce the complexity of the tags, the available frame sizes are limited to powers of two. This results in a reduction in the maximum throughput to 35% from the maximum of  $e^{-1} \approx 37\%$  for some tag estimates (cf. Figure 3.6).



Figure 3.7.: Frame-by-Frame and Slot-By-Slot Frame Size Updating. In the latter approach, the estimate of the number of tags transmitting in the frame is updated after each slot.

# 3.2. Bayesian Transmission Control Strategies

In this section, we present two transmission strategies that address the characteristics of the RFID domain mentioned in the previous section. Both schemes explicitly model medium access in framed ALOHA and compute the probability that a certain number of tags are present based on the feedback from the reader. They make no restrictive assumption about the probability distribution of the random variable that represents the number of tags powered and ready to transmit. The two strategies differ in that the second approach updates the estimated probability distribution of tags present after each slot, while the first approach updates the probability distribution after each frame only (cf. Figure 3.7).

### 3.2.1. Frame-by-Frame Bayesian Updating

In this subsection we present the first of two transmission control strategies. It assumes that the feedback from the reader is only available at the end of a frame. The individual steps of the broadcast scheme are adapted from [100] to suit the nature of framed Aloha and RFID:

1. Compute the frame length L based on the current probability

distribution of the random variable N that represents the number of tags transmitting.

- 2. Start frame with L slots and wait for tag replies.
- 3. Update probability distribution of N based on evidence from the reader at the end of the frame. The evidence comprises the number of empty, singly-occupied, and collision slots in the last frame.
- 4. Adjust probability distribution N by considering newly arriving tags and departing tags including the ones which successfully replied and do not transmit in subsequent slots.

#### Computing the Optimum Frame Size

In step 1 of our procedure the optimum frame length is computed, given the probability distribution of N. We choose the frame length L which maximizes the expected throughput U (cf. Eqn. 3.1).

$$E(U(L)) = \sum_{i=0}^{n_{max}} U(N=i,L) Pr(N=i)$$
(3.5)

This approach is computationally feasible because the available frame sizes are limited to powers of 2 in RFID protocols using framed ALOHA, such as [32, 63, 95].

#### Bayesian Updating of the Probability Distribution

Let H, S, and C denote random variables indicating the number of empty, success (singly-occupied), and collision slots in a single frame with L slots and N tags. After the frame is completed and the feedback in terms of H, S, and C is available, the number of tags that replied is estimated. According to Bayes' rule, the probability that N tags have been transmitting in the frame at time t, given all evidence  $z_{1:t}$ including that from the past frame, is then given by

$$Pr(N|z_{1:t}) = \alpha Pr(N|z_{1:t-1}) \cdot Pr(z_t|N)$$

$$= \alpha Pr(N|z_{1:t-1}) \cdot Pr(C, H, S|N)$$

$$(3.6)$$

where  $\alpha$  is a normalizing constant.

#### Computing the Conditional Probability Distributions Pr(C, H, S|N)

Let us first consider the problem of determining the number of ways T(n, c, h, s, L) to distribute *n* distinguishable tags into *L* distinguishable slots 1, 2, 3, ..., L with the first *c* slots containing at least 2 tags, the next *s* slots containing exactly a single tag, and the remaining *h* slots with no tag reply. The exponential generating function<sup>1</sup> for T(n, c, h, s, L) is given by

$$G(x) = \left(\frac{x^2}{2!} + \frac{x^3}{3!} + \frac{x^4}{4!} + \frac{x^5}{5!} + \dots\right)^c x^s$$
(3.7)

By the expansion of  $e^x$ ,

$$G(x) = (e^{x} - (1+x))^{c} x^{s}$$
(3.8)

T(n, c, h, s, L) is given by the coefficient of  $\frac{x^n}{n!}$  in the expansion of G(x). The number of ways V(n, c, h, s, L) to distribute *n* distinguishable tags into *L* distinguishable slots 1, 2, 3, ..., L with *c* slots containing at least 2 tags, the *s* slots containing exactly a single tag, and the remaining *h* slots with no tag reply, is then simply given by

$$V(n, c, h, s, L) = T(n, c, h, s, L) {\binom{L}{c}} {\binom{L-c}{s}} {\binom{L-c-s}{h}}$$
  
=  $T(n, c, h, s, L) \frac{L!}{c!s!h!}$  (3.9)

since there are  $\frac{L!}{c!s!h!}$  different permutations of the collision, single reply, and empty slots.

The conditional probability distribution Pr(C, H, S|N) is given by the ratio of the number of outcomes in the event space – V(n, c, h, s, L) – and the number of outcomes in the sample space – the number of ways to distribute n tags in L slots

$$Pr(C, H, S|N) = \frac{V(n, c, h, s, L)}{L^n}$$
(3.10)

To illustrate the above, let us consider an example outcome of a frame. For two collisions in the first two slots, followed by a single occupied and an empty slot, the exponential generating function becomes

<sup>&</sup>lt;sup>1</sup>Generating functions are a powerful tool to count the number of arrangements or patterns. A good overview of generating functions and their applications is available in [101].

$$G(x) = \frac{x^5}{4} + \frac{x^6}{6} + \frac{5x^7}{72} + \frac{x^8}{45} + \frac{17x^9}{2880} + \frac{41x^{10}}{30240} + \dots$$
(3.11)

The number of ways T(5, 2, 1, 1, 4) that 5 tags can be arranged in the four slots with two collisions in the first two slots, a single slot in the third slot, and no reply in the fourth slot is then given by the coefficient of  $\frac{x^5}{5!}$ 

$$T(5,2,1,1,4) = \frac{5!}{4} = 30 \tag{3.12}$$

The conditional probability distribution Pr(C = 2, H = 1, S = 1 | N = 5) is then given by

$$Pr(C=2, H=1, S=1|N=5) = \frac{T(5, 2, 1, 1, 4)\frac{4!}{2!1!1!}}{4^5} = \frac{45}{128} \quad (3.13)$$

#### Modelling Newly Arriving and Departing Tags

Once the posterior tag number distribution  $Pr(N_t|z_{1:t})$  is calculated, we still need to incorporate the successful transmissions of the last frame. This only applies to RFID protocols where tags transition to a quiet state after successful identification. Under these circumstances, successful transmissions result in a reduction in the number of tags which reply in the next frame. This means that we simply need to drop the first s entries of the posterior tag distribution in order to compute  $Pr(N_{t+1} = n|z_{1:t})$ :

$$Pr(N_{t+1} = n|z_{1:t}) = Pr(N_t = (n+s)|z_{1:t})$$
(3.14)

The number of tags that transmit their ID in the next frame can also change because new tags arrived and others disappeared during the last frame. The exact probability distribution of newly arriving and departing tags,  $P_A(n)$  and  $P_D(n)$ , depends on the application characteristics and technology parameters as mentioned earlier. We will consider two extreme cases. If a number of tags are placed within the range of a reader and no tags are removed or added until all tags are identified, there is no need to update the probability distribution at all. This corresponds to the example of the medical cabinet mentioned earlier (cf. Figure 3.3). On the other hand, there are scenarios where tags continuously move through the range of the reader, e.g. a dock door scenario with a UHF portal antenna (cf. Figure 3.3). Here, new tags arrive because they leave a deep fade or are powered for the first time; some tags depart because they lose power as they enter a deep fade or disappear from the vicinity of the reader all together. We can compute the probabilities for  $N'_{t+1}$  then as:

$$Pr(N'_{t+1} = n) = \sum_{j=0}^{n} Pr(N_{t+1} = j)P_A(n-j) \qquad (3.15)$$
$$+ \sum_{j=n+1}^{n_{max}} Pr(N_{t+1} = j)P_D(j-n)$$

where the conditioning evidence  $z_{1:t}$  is omitted.

#### 3.2.2. Slot-by-Slot Bayesian Updating

The Bayesian scheme presented in the previous section updates the tag estimate at the end of the frame. This approach is especially useful when the transmission strategy is not implemented on the reader device and the reader hardware reports the number of empty, single, and collision slots after the completion of a frame. Due to the possibility of cancelling a frame early and the unknown number of tags arriving and departing, we will present another transmission control strategy that evaluates the current frame size as the frame progresses. This allows us to cancel a current frame and initiate a new frame with a more appropriate frame size. The individual steps of the broadcast scheme are adapted to include steps that provide this functionality:

- 1. Compute the frame length L based on probability distribution Pr(N).
- 2. Start frame with L slots and wait for tag replies.
- 3. Update Pr(N) based on evidence from the reader at the end of each slot.
- 4. Adjust Pr(N) for tags that are departing during the current frame because they lost power.
- 5. If frame length L is optimal, given Pr(N), continue with the next slot and go back to step 3. Otherwise, cancel current frame.

6. Adjust Pr(N) by considering the arrival of "new" tags and the departure of tags that were successfully identified.

Steps 1 and 2 of the above procedure follow the same principles outlined in the previous section, where we discussed a Bayesian approach that updates Pr(N) at the end of each frame only.

#### Bayesian Updating of the Probability Distribution

Bayes' rule is again used to update the probability that n tags are replying in the current frame, given all evidence  $z_{1:t}$  from previous frames and the evidence  $y_{1:j}$  from the first j slots in the current frame:

$$Pr(N|y_{1:j}, z_{1:t}) = \alpha Pr(N|y_{1:j-1}, z_{1:t})$$

$$\cdot Pr(y_j|N, y_{1:j-1}, z_{1:t})$$
(3.16)

where  $\alpha$  denotes a normalizing constant. Since consecutive frames are considered to be independent given the number of tags transmitting, the following holds

$$Pr(y_j|N, y_{1:j-1}, z_{1:t}) = Pr(y_j|N, y_{1:j-1})$$
(3.17)

#### Computing the Conditional Probability Distribution $Pr(y_j|N, y_{1:j-1})$

Let us first consider the problem of determining the number of ways  $T_C(n, c, h, s, L)$  to distribute *n* distinguishable tags into *L* distinguishable slots 1, 2, 3, ..., L with the first *c* slots containing at least 2 tags, the next *s* slots containing exactly a single tag, the next *h* slots with no tag reply, the j = (c+h+s+1) th slot containing at least 2 tags, and the remaining L-c-h-s-1 slots containing an unconstrained number of tags. The exponential generating function for  $T_C(n, c, h, s, L)$  is given by

$$F_{C}(x) = \left(\frac{x^{2}}{2!} + \frac{x^{3}}{3!} + \frac{x^{4}}{4!} + \frac{x^{5}}{5!} + \dots\right)^{c+1} x^{s} \qquad (3.18)$$
$$\cdot \left(1 + x + \frac{x^{2}}{2!} + \frac{x^{3}}{3!} + \frac{x^{4}}{4!} + \frac{x^{5}}{5!} + \dots\right)^{L-c-s-h-1}$$
$$= (e^{x} - (1+x))^{c+1} x^{s} e^{(L-c-s-h-1)x}$$

Similarly, for the jth slot featuring a single reply, the exponential generating function for  $T_S(n, c, h, s, L)$  is given by

$$F_{S}(x) = \left(\frac{x^{2}}{2!} + \frac{x^{3}}{3!} + \frac{x^{4}}{4!} + \frac{x^{5}}{5!} + \dots\right)^{c} x^{s+1}$$

$$\cdot \left(1 + x + \frac{x^{2}}{2!} + \frac{x^{3}}{3!} + \frac{x^{4}}{4!} + \frac{x^{5}}{5!} + \dots\right)^{L-c-s-h-1}$$

$$= (e^{x} - (1+x))^{c} x^{s+1} e^{(L-c-s-h-1)x}$$
(3.19)

In the event that the jth slot is empty, the exponential generating function for  $T_H(n, c, h, s, L)$  is given by

$$F_{H}(x) = \left(\frac{x^{2}}{2!} + \frac{x^{3}}{3!} + \frac{x^{4}}{4!} + \frac{x^{5}}{5!} + \dots\right)^{c} x^{s}$$

$$\cdot \left(1 + x + \frac{x^{2}}{2!} + \frac{x^{3}}{3!} + \frac{x^{4}}{4!} + \frac{x^{5}}{5!} + \dots\right)^{L-c-s-h-1}$$

$$= (e^{x} - (1+x))^{c} x^{s} e^{(L-c-s-h-1)x}$$
(3.20)

The number of ways  $T_T(n, c, h, s, L)$  to distribute *n* distinguishable tags into *L* distinguishable slots 1, 2, 3, ..., L with the first *c* slots containing at least 2 tags, the next *s* slots containing exactly a single tag, the next *h* slots with no tag reply, and the remaining L - c - h - s slots containing an unconstrained number of tags, can be computed with the following exponential generating function

$$F_T(x) = \left(\frac{x^2}{2!} + \frac{x^3}{3!} + \frac{x^4}{4!} + \frac{x^5}{5!} + \dots\right)^c x^s \qquad (3.21)$$
$$\cdot \left(1 + x + \frac{x^2}{2!} + \frac{x^3}{3!} + \frac{x^4}{4!} + \frac{x^5}{5!} + \dots\right)^{L-c-s-h}$$
$$= (e^x - (1+x))^c x^s e^{(L-c-s-h)x}$$

The number of ways of distributing the *n* tags into the *L* slots with the above constraints is given by the coefficient of  $\frac{x^n}{n!}$  in the expansion of the corresponding generating function F(x).

Note that

$$T_T(n, c, h, s, L) = T_C(n, c, h, s, L) + T_S(n, c, h, s, L)$$
(3.22)  
+  $T_H(n, c, h, s, L)$ 

The conditional probability distribution  $Pr(y_j = collision | y_{1..j-1}, N)$ that the jth slot is a collision slot given that there were c collisions, h empty slots, and s single occupied slots in the j - 1 previous slots in the frame is then given by

$$Pr(y_j = collision|y_{1..j-1}, N) = \frac{T_C(n, c, h, s, L)}{T_T(n, c, h, s, L)}$$
(3.23)

Similarly for the conditional probability  $Pr(z_j = empty|C, H, S, N)$ and  $Pr(z_j = single|C, H, S, N)$ ,

$$Pr(y_j = empty|y_{1..j-1}, N) = \frac{T_H(n, c, h, s, L)}{T_T(n, c, h, s, L)}$$
(3.24)

$$Pr(y_j = single|y_{1..j-1}, N) = \frac{T_S(n, c, h, s, L)}{T_T(n, c, h, s, L)}$$
(3.25)

Note that we only consider a single arrangement of the c collision, h empty slots, and s single slots when we compute T(n, c, h, s, L). This is feasible because the conditional probability  $Pr(y_j|y_{1..j-1}, N)$  is the same for all different arrangements of collision, empty, and single occupied slots.

#### Adjusting Pr(N) for tags that are departing during the current frame

Because of the frequent field nulls, some tags which received the initial start of frame signal will lose power before they can reply in the slot randomly selected. The result is that the average number of tags that reply decreases as the frame progresses. To take into account this early departure of transmitting nodes, we use the mechanism proposed in Eqn. 3.15, but apply it after each slot. No new tags can arrive during a frame because tags which power up during the frame will have missed the initial start of frame signal. The probability of n newly arriving tags  $P_A(n)$  is thus 0 for all n > 0.

#### Evaluation of the current frame size

After each slot, we compute the expected throughput in the next slot based on the updated probability distribution of N for the different frame size according to Eqn. 3.5. If the expected throughput with a different frame size is larger by a certain margin, we cancel the current frame and initiate a new one with the modified size. We use hysteresis to prevent oscillations between two frame sizes that exhibit marginally different performance.

# 3.3. Related Work

In the previous two sections we presented two Bayesian transmission strategies that control the frame size to optimize the throughput. The two schemes address in particular the challenges of the RFID domain mentioned earlier. The idea of controlling the ALOHA channel with a transmission strategy is not new, however [72]. Work focussing in particular on controlling an ALOHA channel with an additional frame structure has been carried out by Schoute [111] and Wieselthier [129]. Schoute developed a backlog estimation technique for framed ALOHA which is exact under the assumption that the frame size is chosen in such a way that the number of stations which transmit in each time slot is Poisson distributed with mean 1. The backlog after the current frame  $B_t$  is then simply given by:

$$B_t = 2.39c \tag{3.26}$$

where c is the number of collisions in the current frame. Due to the unknown distribution of the number of tags and limited number of available frame sizes, the assumption made by Schoute leads to deviations between the estimate and the true number of tags present, whenever the above assumption is not valid. The comparison with the transmission scheme proposed in this thesis, which does not make this simplifying assumption at the expense of additional computations, shows how this restriction affects performance.

Schoute also proposes a method to estimate the conditional probability distribution of collisions and single occupied slots in a frame, given the number of stations and the frame size. While we use a combinatorial method to estimate the above probability distribution, the scheme in [111] uses a method recursive in the number of tags to compute Pr(C, S|N) for a given frame size L:

$$Pr(C, S|N) = Pr(C = c, S = s|N = n - 1) \cdot \frac{c}{L}$$
(3.27)  
+  $Pr(C = c - 1, S = s + 1|N = n - 1) \cdot \frac{s + 1}{L}$   
+  $Pr(C = c, S = s - 1|N = n - 1) \cdot \frac{L - c - s + 1}{L}$ 

The probability of c collision slots and s single reply slots, when n tags are transmitting in a frame with L slots, is thus the sum of the following three events, which are mutually exclusive:

- 1. n-1 RFID tags transmit yielding c collisions and s single reply slots, and the nth RFID tag picks one of the collision slots.
- 2. n-1 RFID tags transmit yielding c-1 collisions and s+1 single reply slots, and the nth RFID tag picks one of the single reply slots.
- 3. n-1 RFID tags transmit yielding c collisions and s-1 single reply slots, and the nth RFID tag picks one of the empty slots.

The recursive relation of Eqn. 3.27 needs to be initialized with

$$Pr(C = c, S = s | N = 0) = \begin{cases} 1 & \text{if } (c, s) = (0, 0) \\ 0 & \text{for any other pair } (c, s). \end{cases}$$
(3.28)

In [129], Wieselthier et al. present a performance evaluation of framed ALOHA with capture. It is based on a combinatorial technique that computes the probability that there are i single occupied slots, j slots with two replies, k slots with three replies, etc., in a frame. Their work is based on the assumption that the number of slots per frame L is fixed and access to the channel can be controlled by adapting the probability with which backlogged stations respond in subsequent frames. In our analysis, the frame size is variable – though limited to powers of two – but all backlogged tags respond in subsequent frames.

More recently, Vogt [123] and Zhen et al. [132] also studied framed ALOHA in the context of RFID. Zhen et al. use the approach proposed by Schoute to estimate the number of tags. Vogt presents a backlog estimation procedure that selects the tag number estimate that minimizes the error between the observed number of empty h, singly-occupied s, and collision slots c and the expected values E(H), E(S), E(C):

$$\min_{N} \left| \begin{pmatrix} h \\ s \\ c \end{pmatrix} - \begin{pmatrix} E_{N}(H) \\ E_{N}(S) \\ E_{N}(C) \end{pmatrix} \right|$$
(3.29)

While this approach does not assume a fixed multiplicity of conflict as the scheme proposed by Schoute does, it only considers the observations from the current frame and neglects past evidence. Vogt compares the



Figure 3.8.: Q algorithm presented in [32] (Source: EPCglobal).

above estimation algorithm to an estimation strategy that represents a lower bound:

$$B_t = 2c \tag{3.30}$$

The Q Algorithm defined in [32] represents another transmission control strategy. It keeps a representation of the current frame size which is multiplied by a constant  $\beta$  whenever a collision occurs and which is divided by  $\beta$  whenever an empty slot is detected. A successful slot leaves the estimate unchanged. The estimated backlog  $B_t$  after the current frame is given by

$$B_t = 2^{Q_{fp}} = 2^{Q_t + a(c-h)} = 2^{Q_t} \beta^{c-h}$$
(3.31)

The new frame size  $L_{t+1}$  is then chosen as

$$L_{t+1} = 2^{round(\log_2 B_t)} \tag{3.32}$$

While the Q algorithm requires only modest computational resources, it does not specify a method to compute the crucial control parameter  $\beta = 2^a$ . It only provides a range of suitable values ( $1.07 \le \beta \le 1.41$ ). The above backlog estimation procedure can also be carried out after each slot. This allows the early cancellation of frames with suboptimal frame sizes, whenever the optimum frame length computation of Eqn. 3.32 indicates this. After  $Q_{fp}$  has been updated accordingly, the algorithm compares the updated  $Q_{fp}$  against the Q of the current frame size  $(2^Q)$ . If the current frame size is deemed suboptimal, it is cancelled and a new frame with an optimized length is initiated.

The work by Rivest [100] is also closely related to our work on transmission control strategies for slotted ALOHA. Rivest introduces an elegant pseudo-Bayesian transmission strategy by approximating the probability distribution of the number of stations N with a Poisson distribution with mean v. Each station keeps a copy of v and during each slot, transmits a packet with a probability of 1/v. It decrements v by 1 if the current slot is empty or a success, increments v by 1.39 if the current slot is a collision, and sets v to  $\max(v + \lambda)$ , where  $\lambda$  is an estimate of the arrival rate. While our transmission scheme builds on the Bayesian approach outlined in [100], our mathematical model has been adapted to suit framed ALOHA and does not assume that the random variable N denoting the number of tags is Poisson distributed. Due to the large number of arriving and departing tags in RFID applications, the Poisson assumption leads to a slow response to changes in the number of tags present. In [54], Frigon et al. present a pseudo-Bayesian algorithm which extends to framed ALOHA and mixed priorities. It also assumes that the number of tags present follow a Poisson distribution.

Krohn et al. [71] recently presented an approach which presents a fast technique to estimate approximately the number of tags present in the read range. Their work is based on the assumption that there are empty slots and occupied slots only.

# 3.4. Evaluation Methods

In this section, we present the methods we used to compare the two Bayesian transmission strategies presented earlier against the schemes outlined in the related work section. We begin by outlining our experimental evaluation method, which comprises our experimental set-up. In the second part of this section, a simulation engine is presented that implements the EPCglobal UHF Class 1 Generation 2 protocol and that supports different path loss, fading, capture, and tag mobility models.

### 3.4.1. Experimental Set-Up

The experimental evaluation relies on an HF RFID system that uses the Philips I-Code1 air interface protocol [95]. 64 HF Philips I-Code1

tags [95] are placed on an antenna attached to a Ridel5000 reader made by Softronica [113] (cf. Table 3.2 and Figure 3.9). To test the transmission control strategies under different traffic conditions, we used five different frame sizes (16, 32, 64, 128, 256). This corresponds to the following five different average traffic rates 4, 2, 1, 0.5, 0.25 respectively. Each frame is initiated with an 'Unselected Read' command that communicates the frame size to the tags and triggers their replies. At the end of the frame, we collect the number of empty, singly-occupied, and collision slots that the reader reports (cf. Table 3.3). The experiment is repeated 1000 times for each frame size. These measurement data are then fed to the different transmission control strategies for evaluation. Since we limit our experiment to a single frame, we compare the transmission schemes based on the quality of the tag number estimate. The better the estimate, i.e., the closer to the known number of tags transmitting, the more appropriate the choice of the next frame size in a real world application and the higher the throughput.

# 3.4.2. Simulation Set-Up

To further evaluate the performance of the two transmission strategies, in particular for UHF RFID systems, we developed an RFID simulation engine. The RFID simulator implements the EPCglobal UHF Generation 2 Class 1 Protocol and features a pathloss, fading, capture, and tag mobility model.

### **RFID Simulator Design**

The RFID simulator runs as a collection of simulation entities in the simulation runtime JiST, which was developed by Rimon Barr [11]. JiST is a discrete event simulation engine that runs over a standard Java virtual machine. JiST represents an approach to building discrete

Property	Value
Frequency	$13.56 \mathrm{~MHz} \mathrm{~(HF)}$
Reader	Softronica Ridel5000 [113]
Reader Antenna Size	$300 \text{ mm} \ge 400 \text{ mm}$
Tag IC	Philips I-Code1 SL1 [95]
Tag Antenna Size	$39 \mathrm{mm} \ge 39 \mathrm{mm}$
Communication Protocol	Philips I-Code1 [95]
Reader Transmit Power	4 W

Table 3.2.: Equipment Overview.



(b) 64 HF I-Code1 Tags placed on top of the reader antenna

Figure	3.9.:	Exper	iment	al Setup.	RFID	reader,	reader	antenna,	and	64	tags.
										_	

Data	Frame	No.	$\mathbf{Empty}$	Single	Collision
Set	Size	of	$\mathbf{Slots}$	$\mathbf{Reply}$	Slots
		Tags		$\mathbf{Slots}$	
1	32	64	4	14	14
2	32	64	3	13	16
3	32	64	5	11	16
998	32	64	3	14	15
999	32	64	2	11	19
1000	32	64	4	12	16

Table 3.3.: Example Data Sets.

event simulators, called virtual machine-based simulation, that unifies traditional systems and language-based simulator designs. Simulation code that runs on JiST need not be written in a specific simulation language, nor need it be littered with system calls to support runtime simulation functionality (cf. Figure 3.10 and 3.11). Instead, JiST converts an existing virtual machine into a simulation platform by em-



Figure 3.10.: JiST System Design [13]. Simulations are compiled (1), dynamically modified by a byte-code rewriter (2), and executed in the Java virtual machine.

bedding simulation time semantics at the byte-code level. Our RFID simulator is thus written in Java, compiled using a regular Java compiler, and run over a standard, unmodified virtual machine. The JiST simulation engine is also efficient and performs well when compared to existing highly optimized simulation runtimes both in time and memory consumption [13].

The pathloss, fading, and mobility models we implemented build on the Scalable Wireless Ad-hoc Network Simulator [12] that is also implemented on JiST. There is also support for different capture models, such as a stochastic model [129] and different power models [98]. The simulation engine supports the entire command set of the EPCglobal UHF Generation 2 Class 1 Protocol except for a few optional commands that allow access to user memory. The basic architecture of the RFID simulator is shown in Figure 3.12.

The RFID simulation engine currently only supports a single reader. While the directivity of the reader antenna can be specified, it is assumed that all tag antennas are isotropic. The simulation is also limited to two dimensions in space. Multipath effects cannot be modelled explicitly, but need to modelled by statistical fading models, such as Rician or Rayleigh fading.

#### Simulation Scenario

To compare the different transmission strategies, the above RFID simulator is used in the dock door scenario shown in Figure 3.14. The  ${\bf public \ final \ class \ RFIDTagMac \ extends \ RFIDMac \ } \{$ 

public void send(Message msg, boolean replyExpected) {

```
TagReply reply = (TagReply) msg;
// add preamble to tag reply
reply.addPreamble(new TagPreamble());
// compute transmit time for tag reply
long duration = transmitTime(reply);
// compute turnaround time for tag reply
long timeToWait = waitTime();
// schedule tag reply transmission
JistAPI.sleep(timeToWait);
radioEntity.transmit(reply, duration);
```

Figure 3.11.: Code example of the RFID Simulator. JiST transparently introduces simulation time execution semantics to programs written in plain Java.

movement and identification of two pallets carrying 200 UHF tags is simulated. The multipath fading is modelled statistically by a Rician distribution, which is commonly used to describe the small-scale fading envelope, when there is a strong line-of-sight component [98]. The other simulation parameters are listed in Table 3.4. Figure 3.14 shows the number of tags that are powered at any moment in time as the pallet carrying the tags is moved past a reader in the dock door application.

# 3.5. Evaluation of the Bayesian Transmission Strategies

The evaluation of the transmission strategies is split into two parts. We begin by showing how the Bayesian scheme that operates on reader feedback at the end of the frame compares to other schemes that also operate on a frame-by-frame basis. We compare in particular the

}



Figure 3.12.: Overview of RFID Simulator Architecture [131]. The solid boxes represent simulation entities which communicate with each other by exchanging simulation events.

Parameter Name	Value
Path Loss Exponent	2
Rician factor	6 dB
Reader power (EIRP)	$3300 \mathrm{mW}$
Reader antenna 3dB beamwidth	$60^{o}$
Reader capture ratio	32  dB
Reader sensitivity	-80 dBm
Minimum tag input power	-15.22 dBm
Tag backscatter factor	0.25
BER	$10^{-3}$

Table 3.4.: Simulation parameters.

scheme proposed by Schoute, Vogt, and the lower bound estimate, which we discussed in the related work section, with our Bayesian update scheme. In the second part of the evaluation, we illustrate the difference between the Bayesian slot-by-slot technique and the frameby-frame technique. We also compare our Bayesian approach to the Q algorithm, which is part of the EPCglobal UHF Class 1 Generation 2 [32] specification and which can also operate on a slot-by-slot basis.



Figure 3.13.: Loading dock application. Pallets containing 200 randomly distributed RFID tags are moved past a single reader at a constant speed of 3 m/s.



Figure 3.14.: Simulated arrival pattern for the loading dock scenario of Figure. In the UHF frequency band, field nulls caused by small-scale fading effects lead to frequent power losses of the tags. The x-axis in the figure above denotes the position of the centre of the pallet relative to the reader.



Figure 3.15.: Experimental evaluation using 64 HF Philips I-Code1 Tags. The Bayesian approach is compared to the scheme of Schoute, Vogt and the lower bound estimate. In this experiment, the Bayesian approach uses evidence from a single frame only.

#### 3.5.1. Evaluation of the Frame-by-Frame Bayesian Updating

In this section, we evaluate the Bayesian transmission control scheme that operates on a frame-by-frame basis. Figure 3.15 shows the results of an experimental evaluation. 64 HF Philips I-Code1 [95] tags are placed in the read range of an RFID reader as described in Section 3.4. The figure shows the variation of the estimate as a function of traffic. It compares the algorithm proposed by Schoute (cf. Eqn. 3.26), the algorithm by Vogt (cf. Eqn. 3.29), the lower bound estimate (cf. Eqn. 3.30), and our Bayesian frame-by-frame scheme.

To evaluate the transmission strategies in different traffic scenarios, we used five different frame sizes (16, 32, 64, 128, 256). The different traffic rates shown in Figure 3.15 correspond to the ratio of the fixed number of tags to the frame size. The results illustrate how the Bayesian transmission scheme and the approach proposed by Vogt exhibit superior performance under the various traffic rates evaluated. The two schemes dynamically adjust the multiplicity of conflict, whereas the multiplicity of conflict in the schemes proposed by Schoute and the lower bound estimate is fixed at 2.39 and 2, respectively. In the Bayesian and Vogt scheme, the multiplicity of conflict is a function of the evidence and the frame length.

Lower bound:	$N_t =$	2c+s
Schoute:	$N_t =$	2.39c + s
Vogt:	$N_t =$	$f(z_t, L)c + s$
Bayesian Frame-by-Frame:	$N_t =$	$f(z_{1t}, L)c + s$

Note that the above equations do not actually represent the Vogt and Bayesian Frame-by-Frame algorithms, but represent a simplification that is used to summarize their operation. Figure 3.15 shows that the algorithm proposed by Vogt also provides a reasonable estimate, but shows greater variance. The variance of the Bayesian scheme is further decreased by including all evidence  $z_{1..t}$  from past frames, which is not done in the scheme proposed by Vogt. It relies on the evidence from the last frame only  $z_t$ . This feature of the Bayesian scheme is illustrated in Figure 3.16, which shows how the variance of the distribution is reduced as reader feedback of multiple frames is incorporated in the estimate. In this experiment, the tags are not silenced after the successful identification and it is assumed that no new tags arrive during the identification process. The figure also shows the number of collision slots (C), the number of slots with a single tag reply (S), and the number of empty slots (H) in each frame.

The more accurate estimates of the Bayesian frame-by-frame and Vogt approaches translate into a higher throughput in the subsequent frame. Based on the estimated number of tags, the Bayesian algorithm chooses a frame with 64 slots, while the other two schemes choose a frame with 32 slots, assuming a traffic rate of 4. This corresponds to an expected throughput increase of 37% for the Bayesian scheme over the other two in this particular case (cf. Figure 3.17).

The experimental results (cf. Figure 3.15) also show that the Bayesian frame-by-frame scheme predicts that there are fewer RFID tags transmitting than the 64 tags placed on the reader antenna as shown in Figure 3.9. We associate this discrepancy with a combination of the following causes:

**Capture Effect.** The capture effect refers to the ability to receive a tag reply correctly, despite the presence of other tag replies transmitted simultaneously [102]. The experimental results shown in



Figure 3.16.: Evolution of Bayesian estimate as the evidence from the current frame is combined with evidence from the previous frames. Successfully identified tags are not silenced and no new tags arrive. The frame size is fixed at 16 slots and the reader feedback regarding the last frame is shown in each figure.

Figure 3.18 indicate that there will also be captured tag replies in inductively coupled RFID systems, if the difference in received signal strength is large enough. As the distance between two HF tags that are forced to reply in the same slot passes a certain threshold, the tag reply from the tag closer to the reader antenna is correctly received with a certain probability, i.e. captured. Note that the tag further away is still powered up and replying. Although some RFID systems can distinguish a captured reply from a single reply



Figure 3.17.: Expected Throughput vs. Number of Timeslots in a frame for 64 HF Philips I-Code1 Tags. The figure shows how the superior estimate translates into an increased throughput.

under some conditions as shown in [94], we have to assume that the reader provides feedback without capture in most cases. Our Bayesian algorithm will underestimate the true number of tags if there is capture, but the reader cannot detect the presence of weak tag replies.

In the experiment shown in Figure 3.9, we tried to eliminate the capture effect by keeping the distance of separation between the RFID tags and the reader antenna constant. The variations in magnetic field strength across the reader antenna and the manufacturing tolerances of the RFID tag components might still result in variations in the received signal strength, however.

RFID communication protocols based on framed ALOHA can benefit from this effect because they can achieve throughput values that exceed the theoretical maximum throughput associated with slotted and framed ALOHA  $e^{-1}$ . Figure 3.19 shows the effect of capture on the throughput. The model assumes a capture probability of 20%, for simplicity's sake. This is an unrealistic capture model, but illustrates the result of the capture effect: the maximum throughput is increased and occurs at a frame length L < N.

While the capture effect results in an underestimate of the number of tags, the influence on the overall throughput is small because



Figure 3.18.: *Experimental evidence of the capture effect.* Two HF tags are forced to reply in the same slot.



Figure 3.19.: Expected Throughput vs. Number of Timeslots for 64 tags without capture and with a capture probability of 20%. The figure illustrates that the optimum throughput no longer occurs at L = N, but at L < N. With the above capture probability and the 64 tags, 51 timeslots would result in the optimum throughput.

the underestimate is compensated by not including the capture effect in the computation of the expected throughput, which would mandate choosing a frame size that is not equal, but smaller than the number of tags present.

- **Errors in reader feedback**. The underestimate of the true number of tags can also be caused by wrong feedback from the reader. Our Bayesian approach assumes perfect feedback from the reader. In practice, the reader does not always successfully distinguish slots with no reply, a single reply, and a reply with multiple tag replies. The reader occasionally interprets a collision slot as a single tag reply with a bit error. Whenever a collision is misclassified as a single tag reply with a communication error or no tag reply, the schemes presented above will underestimate the true number of tags transmitting.
- Not all tags reply in all frames. While all RFID tags were successfully tested individually before the experiment, occasionally some tags would not transmit in a frame. Unfortunately, it is extremely difficult to determine which tags actually transmitted during the experiment.

It is straightforward to include feedback with capture as well as communication errors in our transmission scheme. With K and E denoting random variables that represent the number of captured slots and error slots respectively, the conditional probability Pr(C, S|N) becomes

$$Pr(C, S|N) = Pr(C = c + k, S = s + e|N)$$
(3.33)

The number of captured replies also needs to be considered when the number of tags that will not reply in subsequent frames because they have been successfully identified is taken into account. The posterior distribution  $Pr(N_{t+1} = n|z_{1:t})$  is then updated by dropping the first s + k entries:

$$Pr(N_{t+1} = n|z_{1:t}) = Pr(N_t = (n+s+k)|z_{1:t})$$
(3.34)

Eqn. 3.1 also needs to be updated to include the capture effect. This is done by assuming that the capture probability  $\frac{k}{c+k}$  in the subsequent frame is identical to the one in the past frame. The expected throughput becomes the sum of the expected number of single reply and captured slots divided by the total number of slots.



(b) Too many empty slots

Figure 3.20.: Frames with a low throughput due to an over- and underestimate of the number of RFID tags transmitting. The slots following the startof-frame (SOF) signal are either mostly empty or collision slots.

$$U(N,L) = \left(E(S_{N,L}) + \frac{k}{k+c}E(C_{N,L})\right)L^{-1}$$
(3.35)

While both the Bayesian frame-by-frame scheme and the scheme proposed by Vogt provide good estimates of the number of tags, the estimate is only available at the end of the frame. This means that a wrong estimate that leads to a frame full of collisions or empty slots cannot be interrupted (cf. Figure 3.20). In applications such as the medicine cabinet shown in Figure 3.3, this will only happen in the first frame of the identification process, since there are no new tags arriving and departing due to power losses during the identification progress. In applications such as the dock door scenario shown in Figure 3.4, the number of tags transmitting varies considerably during the identification process. In the following section, we show how the overall throughput can be improved by the slot-by-slot Bayesian technique of Section 3.2.2, which updates the estimate of the number of tags transmitting after each slot.

### 3.5.2. Evaluation of the Slot-By-Slot Bayesian Updating

To analyze the impact of reducing the response time of the transmission scheme from a single frame to a few slots only, we will focus on the second Bayesian algorithm presented in this thesis. The characteristics of the RFID domain outlined in Section 3.1 suggest that a transmission scheme is desirable that can interrupt a running frame, if the current frame size is considered non-optimal. Our analysis is split into two parts. We begin by showing the performance improvement that results from our Bayesian slot-by-slot updating scheme, and continue by comparing our Bayesian algorithm against the Q Algorithm of the EPCglobal UHF Class 1 Generation 2 protocol.

## 3.5.3. Comparison with Frame-By-Frame Transmission Schemes

Frame-based methods such as the approaches evaluated in the previous section provide good backlog estimates based on the number of tags that replied in the previous frame. Nevertheless, the arrival and departure of tags caused by the frequent field nulls and tag movements themselves will significantly increase the uncertainty about the true number of tags ready to reply in the next frame in applications such as the dock door scenario presented earlier. In this subsection, we demonstrate that the Bayesian transmission control strategy that is capable of evaluating the evidence on a slot-by-slot basis provides a better throughput because it cancels a frame once the scheme determines that the original estimate is not correct.

We use the simulation scenario outlined in Section 3.4 to compare the performance of the Bayesian frame-by-frame and slot-by-slot transmission strategies quantitatively in the presence of unknown arrival and departure rates. Two pallets carrying 200 UHF tags are moved past a single antenna and the transmission control strategy changes the frame size accordingly, until all 200 tags are identified.

The different transmission schemes are compared to what we termed the "perfect estimate at the end of frame" (cf. Figure 3.21). This transmission scheme knows at the end of each frame exactly how many tags responded in the last frame and chooses the next frame size accordingly. The divergence from the theoretical maximum throughput evident in Figure 3.21 results from the unknown number of newly arriving and departing tags in the consecutive frame. The simulation results also show


Figure 3.21.: Divergence of the simulated throughput from the theoretical maximum of 37% for three different transmission schemes. The figure illustrates the improvement that can be expected from transmission schemes which operate on a slot-by-slot basis.

that the performance of the Bayesian frame update algorithm is essentially the same as the perfect estimator. Its throughput is 11% below the maximum theoretical throughput of framed ALOHA. The Bayesian scheme that updates the probability distribution of tags transmitting in the frame after every slot achieves the highest throughput (34% on average), which is close to the theoretical maximum of 37%. Due to the limited number of different frame sizes, the actual theoretical maximum will even be slightly less (cf. Figure 3.6).

#### 3.5.4. Comparison with the Q Algorithm

In the previous subsection, we demonstrated the throughput improvement resulting from the use of our Bayesian slot-by-slot technique. In this subsection, we compare the Bayesian slot-by-slot scheme with the Q Algorithm, which is part of the EPCglobal UHF Class 1 Generation 2 [32] specification and which can also operate on a slot-by-slot basis (cf. Figure 3.8).

While the Q algorithm requires only modest computational resources, it does not specify the value of the crucial control parameter  $\beta$ , which is used to update the estimated number of RFID tags:

Since this parameter depends on a past estimate that incorporates past evidence, the evidence from the current slot, and the current frame size, its choice is not trivial. A value of  $\beta$  that is chosen too large will lead to significant overshoots, while a  $\beta$  that is too small will reduce the swiftness of a response to a change. The Q Algorithm also assumes that the frame size is optimal, once the ratio of collision to empty slots is equal to one. However, computing the expected number of collisions and empty slots at L = N, which is the criterion for optimum throughput, implies a ratio of collision to empty slots that is smaller than 1.

Our Bayesian slot-by-slot updating algorithm, on the other hand, explicitly models framed ALOHA. It indirectly uses all available information, i.e. all past evidence, including the evidence from the last slot and the current frame size, to compute the multiplicity of conflict. However, the Bayesian algorithm does require significant computing resources, although some of the computations can be precomputed and stored in the memory of the reader device.

In Figure 3.22, both slot-by-slot transmission schemes are compared. Our implementation of the Q algorithm uses a value for  $\beta$  which is set to  $2^{\frac{0.8}{\log_2 L}}$ . The throughput achieved with the Q algorithm in our simulation scenario is on average 33%, which represents a 9.5% divergence from the maximum throughput of 37%. The Bayesian slot-by-slot algorithm provides only a slightly better performance (Figure 3.22).

Figure 3.22 also shows the performance of the Q Algorithm, if applied after each frame. The figure demonstrates that the performance of the Q algorithm is poor in our scenario, given our choice of the constant  $\beta$ . The performance can, however, be significantly improved when changes to Q are restricted to incremental changes (denoted (incremental) in Figure 3.22). Under these conditions the oscillations of the Q algorithm are damped and the simulated throughput is similar to the other framebased transmission schemes.



Figure 3.22.: Performance of the Q Algorithm, which is part of the EPCglobal UHF Class 1 Generation 2 [32] specification.

## 3.6. Limitations and Future Work

The increased throughput that can be achieved with our transmission schemes comes at the expense of a significant amount of computations. Alternative transmission schemes make certain assumptions about the distribution of the number of tags present or simply assume a fixed multiplicity of conflict. This reduces the resources required to estimate the number of tags transmitting and to choose the frame size accordingly. While our Bayesian approaches permit some computations to be made a-priori, there is a significant amount of computation remaining that needs to be carried out online. In the frame-by-frame Bayesian scheme, the conditional probability distribution Pr(C, S|N) can be computed and stored in memory. This means that we only need to update the posterior probability by computing  $Pr(N|z_{1:t})$  after each frame:

$$Pr(N|z_{1:t}) = \alpha Pr(N|z_{1:t-1}) \cdot Pr(C, S|N))$$
(3.36)

Once the posterior distribution is known, the most appropriate frame choice needs to be computed according to Eqn. 3.5.

To simplify the computation even further, the choice of the most appropriate frame size given a posterior could be carried out by computing:

$$L_{t+1} = 2^{round(\log_2 E(N|z_{1:t}))}$$
(3.37)

While we did use an HF RFID system to evaluate our transmission schemes, we relied on simulations at UHF. Our approaches have thus not been validated experimentally with a UHF RFID system. Future work should thus aim to implement the transmission schemes presented in this chapter in a UHF RFID reader. The simulation engine could also be upgraded to include tag antenna directivity and multiple reader antennas.

In our evaluations, we assumed that a reader can operate independently in a given channel. In practice, there will be other readers operating in close vicinity, which will possibly interfere. Future work might thus also consider the effect of such reader collisions on the performance of the transmission schemes. Furthermore, a transmission scheme that chooses an appropriate frame size for a number of readers which are synchronized in order to deal with the limitations of the listen-before-talk schemes introduced in some countries could be part of future investigations.

Our Bayesian models assume a feedback model where the reader can successfully distinguish between no, a single, a single, but corrupted tag reply, and more than a single tag reply. In practice, it might be difficult to always successfully distinguish a corrupted single tag reply from a collision where more than a single tag replied. Future versions of our transmission schemes should thus include the possibility and ideally the likelihood of such a wrong classification.

In our model, we also assume that all slots have the same fixed length. In practice, a reader might close an empty or even collision slot early. The cost of an empty or collision slot might thus vary. Future versions of the above transmission strategies might anticipate this possibility and explicitly model the early closure of empty and collision slots. Our analysis also neglects the overhead associated with reader commands that initiate a frame.

# 3.7. Summary

As the number of objects which are equipped with RFID tags increases, it is becoming increasingly important to identify large tag populations quickly. This mandates among other things a high throughput over the shared radio channel. The throughput performance of RFID medium access protocols, such as ALOHA, depends, however, on a transmission scheme that estimates the (unknown) number of stations transmitting. The number of RFID tags transmitting remains uncertain, since RFID readers cannot detect the multiplicity of conflict if more than two RFID tags transmit simultaneously.

In this chapter, we provide two transmission strategies that speed up the identifications of RFID tag populations by accurately estimating the unknown number of RFID tags transmitting and thus improving the throughput over the shared radio channel. The proposed transmission control schemes build on earlier work on Bayesian broadcast strategies, but have been adapted to suit the characteristics of RFID. The latter include the frequent use of a variant of slotted ALOHA, known as framed ALOHA, in RFID protocols and the unknown tag arrival and departure rates. Both of the two approaches presented in this thesis differ from related work because the transmission schemes make no assumption about the statistical distribution of the number of tags present and the tag arrival rate in the range of the reader. The schemes also include all past evidence from the RFID reader. The two schemes differ in that the first Bayesian transmission strategy updates the estimate of the number of tags transmitting at the end of a complete frame only, while the other Bayesian scheme computes the estimate after each individual slot of a frame.

The transmission control strategies are evaluated experimentally and via simulations. The simulations rely on an RFID simulation engine we developed that supports different pathloss, fading, capture, and tag mobility models. The experimental evaluation shows that the Bayesian scheme can accurately estimate the number of RFID tags transmitting over a broad range of traffic scenarios. The frame-by-frame Bayesian scheme outperforms existing schemes that assume a fixed multiplicity of conflict. The Bayesian scheme also provides a tighter estimate than a scheme proposed by Vogt [123], which also dynamically adjusts the multiplicity of conflict based on feedback from the reader. The frame-by-frame Bayesian scheme is in particular useful for applications where the number of RFID tags is initially unknown, but the number of RFID tags powered does not vary significantly until all RFID tags are successfully identified.

In applications where the number of RFID tags transmitting varies significantly from one frame to the next, the second Bayesian algorithm presented in this thesis is more suitable. The simulation results show that due to the unknown arrival and departure rates in these applications, a transmission strategy that incorporates the feedback from the reader on a slot-by-slot basis performs significantly better than a strategy that waits until the end of a complete frame before the estimate is updated. This is due to the fact that the former scheme will cancel a frame early if the frame size is estimated to be non-optimal. In our simulation, the slot-by-slot Bayesian scheme realizes a throughput of 34%, which is close to the theoretical maximum throughput of framed ALOHA of 37%.

The increased throughput that can be achieved with our transmission schemes comes at the expense of a significant amount of computations. Alternative transmission schemes make certain assumptions about the distribution of the number of tags present or simply assume a fixed multiplicity of conflict. This significantly reduces the resources required to estimate the number of tags transmitting and to choose the frame size accordingly. While our Bayesian approaches permit some computations to be made a-priori, there is a significant amount of computation remaining that needs to be carried out online.

# Supporting the Fair Information Principles in RFID protocols<sup>1</sup>

The impending ubiquity of RFID tags requires not only support mechanisms to provide adequate performance as discussed in the previous chapter, but also measures to address privacy concerns associated with unobtrusive tags on everyday items. When Mark Weiser envisioned computing capabilities everywhere, embedded in the environment in such a way that they can be used without being noticed, he also acknowledged that the invisible nature of the computing devices will make it difficult to know what is controlling what, what is connected to what, and where information is flowing [128]. This tension between the conflicting requirements of control and privacy on the one hand and usability and performance on the other, predicted by Weiser, are well illustrated by the privacy concerns associated with the planned deployment of RFID technology in supermarkets and retail outlets.

In this chapter, we present an approach that addresses this privacy threat by integrating a subset of the widely accepted Fair Information Practices into the communication protocols between RFID readers and tags. We argue that having RFID readers explicitly declare the scope and purpose of their tag data collection, as well as disclosing the identity of their operators, will allow both consumers and regulators to better assess and control the impact of everyday RFID encounters. Our analysis shows how the Fair Information Practices of collection limitation, purpose specification, use limitation, openness, and accountability can be incorporated in today's RFID communication protocols without significant performance penalties. We also present the prototype of a watchdog tag that allows consumer interest groups and privacyconcerned individuals to judge whether a particular RFID reader deployment complies with the corresponding regulations by displaying and logging information regarding the data collection broadcasted over

 $<sup>^{1}</sup>$ This chapter is based on joint work with Roland Schneider and Marc Langheinrich [51]



Figure 4.1.: Consumer fears associated with RFID [66](Source: Ari Juels, RSA).

the radio channel.

The rest of the chapter is organized as follows. After discussing the problem statement and briefly restating the Fair Information Practices and their role in today's privacy legislation, we show how the requirements defined by the Fair Information Practises could be embedded into the reader-to-tag communication of existing RFID standards. We then present an early prototype of a "watchdog" tag, a small personal device that can be used in conjunction with our protocol extensions to further increase the transparency of the identification process. We conclude with a discussion of our approach, giving special attention to its efficiency, as well as outlining future work.

#### 4.1. Problem Statement and Related Work

While the invisible nature of RFID technology has many benefits from an automation perspective, it is also the cause of some serious privacy concerns. The intended deployment of RFID tags on everyday items epitomizes for many the dangers of an Orwellian future: unnoticed by consumers, embedded microchips in our personal devices, clothes, and groceries can covertly be triggered to reply with their ID and other information, potentially allowing for a fine-grained yet invisible surveillance mechanism that pervades large parts of our lives (cf. Figure 4.1).

RFID	Automation Benefit	Privacy Drawback
Property		
No line- of-sight required	No human involvement is re- quired to aim the reader at the RFID tag or vice versa. The result is that RFID system can automatically and continuously identify RFID tags passing by, even if the RFID tag is not fac- ing the reader. RFID tags can also be embedded in the pack- aging, which protects them from	The owner of a tagged object or product might be unaware of the presence of the tag and the process of its memory being accessed by a concealed reader. Since there is no human super- vision required, there is also no implicit consent and the purpose of the data collection and the data collector himself might re- main unknown
Read range of a few me- ters	wear and dirt. No human is required to place the tags in the direct proximity of the reader or vice versa. RFID tags can simply be identified as they pass the reader at some distance, e.g. as they are moved by a fork lift.	The owner of a tagged object might be unaware that an RFID reader is reading the data from the RFID tag, since there is no need to bring the reader close to the tag as in the case of a bar code read.
Data stor- age on the tag	Each tag can carry a unique iden- tifier from a global address space, such as an electronic product code (EPC) [30]. The result is that any tagged object can be uniquely identified. In addition to the unique identifier, the tag can also carry additional appli- cation data, which might com- prise the product manufacturer and the product type.	Clandestine scanning may thus allow tracking of a person's lo- cation across multiple sites and identification of the kind of ob- jects he or she is carrying.

Table 4.1.: Automation benefits and privacy drawbacks of the three major RFID characteristics.

While some of the privacy threats associated with RFID are based on invalid assumptions regarding the capabilities of the technology or simply a not very plausible scenario [77], the feasibility of detecting an RFID tag carrying a unique identifier without line-of-sight over a distance of a few metres nevertheless introduces a set of serious privacy challenges [107]. The capabilities of RFID listed in Table 4.1 allow for automatic and continuous data collection without any human involvement. Taking the human out of the loop, which is desirable from a automation perspective, means that explicit consent is no longer required from the individual carrying the objects, nor might he or she be aware of the purpose of the data collection or even the data collection itself (cf. Table 4.1). Stajano [116] compared this threat to X-ray vision where one can set up a (concealed) reader and use it to identify the type and size of clothes colleagues are wearing. It is thus important to develop appropriate approaches that protect the individual carrying tags in his clothes or groceries against potential misuse of the technology.

Initial proposals to address these privacy concerns included technical solutions such as the kill command that deactivates individual tags. Due to the invisible nature of RFID communication, it is questionable whether the consumer will trust such a kill operation that is again difficult to verify, due to its invisible nature. In our opinion, a mechanical kill operation, as proposed by Karjoth et al. [68] seems more promising. Some consumer interest groups even advocate a complete ban on RFID tags in the public part of stores [97]. Although the latter approach will naturally protect the privacy of the individual, it falls short of an optimal solution even from a consumer standpoint, since it is not just retail stores that can benefit from the use of RFID tags, but also consumers. The magic medicine cabinet [49], the magic wardrobe [56], and the often-cited smart fridge are just some of the consumer applications that would benefit from post-point-of-sales item-level RFID tagging.

Other technical solutions have focussed on proper access control, which only allows authorized parties to decode the global identifier stored on the tag. Examples include key-based access control [114] and cryptographic hash functions that hide the unique identifier on the tag by transmitting a meta-ID [61, 127], which can only be translated into the unique identifier by authorized parties. As Langheinrich points out in [77], these approaches rely on appropriate key management and thus require significant effort from an end user worried about her privacy.

In this thesis, we argue for a middle ground inspired by our everyday lives, where we rarely encounter all-or-nothing tradeoffs, but rather engage in meaningful exchanges that conditionally lead us to disclose parts of our personal data to service providers in return for more or less tangible benefits. By incorporating the basic principles of the widely accepted fair information practices at the reader-to-tag protocol level, RFID-system operators will be able to deploy readers that only collect tag data relevant to the current application, while small personal devices could additionally provide consumers with a detailed look at a reader's operator and its purpose for collecting data, potentially allowing for explicit consent before any tag information is read out. Future tags might even be able to decide independently whether or not to reply to a reader's query, based on its stated ID, purpose, and target range. Having RFID readers explicitly declare the scope and purpose of the tag data collection, as well as disclosing the identity of their operators, will allow both consumers and regulators to better assess and control the impact of everyday RFID encounters.

# 4.2. Fair Information Practices

The Fair Information Practices (FIP), published by the Organization of Economic Cooperation and Development (OECD) in 1980 [93], are a well established set of guidelines for consumer privacy. They have their roots in a 1973 report of the "United States Department for Health, Education, and Welfare (HEW)" and were drawn up by the OECD to better facilitate the cross-border transfer of customer information as part of trade between its member states. The eight principles can be summarized as follows:

- 1. Collection limitation: Data collectors should only collect information that is necessary, and should do so by lawful and fair means, i.e., with the knowledge or consent of the data subject.
- 2. Data quality: The collected data should be kept up-to-date and stored only as long as it is relevant.
- 3. Purpose specification: The purpose for which data is collected should be specified (and announced) ahead of the data collection.
- 4. Use limitation: Personal data should only be used for the stated purpose, except with the data subject's consent or as required by law.
- 5. Security safeguards: Reasonable security safeguards should protect collected data from unauthorized access, use, modification, or disclosure.
- 6. Openness: It should be possible for data subjects to learn about the data collector's identity, and how to get in touch with him or her.
- 7. Individual participation: Data subjects should be able to query data collectors as to whether or not their personal information

has been stored, and, if possible, challenge (i.e., erase, rectify, or amend) this data.

8. Accountability: Data collectors should be accountable for complying with these principles.

The FIP form the basis for many of today's privacy laws, such as the EU Directive 95/46/EC [37], which provides the framework for the national privacy laws of all EU-member states. For example, Article 6 of the Directive requires data collectors to collect only as much information as necessary (also called the *proportionality principle* or the principle of *data minimization*) while article 7 requires them to obtain the unambiguous consent of the data subject before collection.

It is undisputed that the act of reading out one or more RFID tags can constitute a data collection, meaning that existing privacy laws also apply to the communication between tags and their readers. This has also been recently pointed out by the International Community of Data Protection and Privacy Commissioners [1]. At the outset, this would mean that RFID readers would need to be openly announced with the help of public signs and placards explaining the purpose and extent of the data collection, as well as the identity of the data collector [55]. While adequate from a legal point of view, presenting the necessary information in such a way risks the consumer ignoring it, as the ubiquitous privacy policy links on today's Web sites have demonstrated. This is because of two important drawbacks of such an out-of-channel solution: firstly, data subjects need to actively seek out such information, which might otherwise be easily overlooked. Secondly, even when accessible, reading and understanding this information puts an added burden on the consumer, as it is often written in dense legal prose.

On the Web, the Platform for Privacy Preferences Project (P3P) aims at alleviating these two drawbacks [25]. Developed under the auspices of the World Wide Web Consortium (W3C), P3P integrates machine readable privacy policies into the browser-to-server protocol, thus allowing the user's Web browser to automatically read the privacy policy of a Web site, compare it with the user's preferences, and subsequently take action on behalf of the consumer (e.g. facilitating or preventing transfer of personal data, or advising the user in an easily understandable manner). Our goal is to implement a similar mechanism into the protocol between RFID tags and their readers, in order to lessen the burden on the consumer by having her tags (and optionally a personal

Principle	Support
(1a) collection limitation	through selection mask
(1b) consent	with watchdog tag (optional)
(2) data quality	out of scope (use privacy-aware DB)
(3) purpose specification	through purpose declaration
(4) use limitation	with collection types
(5) security safeguards	encryption (future work)
(6) openness	through reader and policy ID
(7) participation	out of scope (use privacy-aware DB)
(8) accountability	through reader and policy ID

Table 4.2.: Support for the FIP in our reader-to-tag air interface. About half of the principles can be embedded directly at the protocol level.

mobile device carried with her) read and process privacy related information autonomously.

Some of these principles, such as individual participation or data quality, will need support primarily in the storage back-end, for example with the help of privacy-aware databases [4, 75, 76]. However, the majority of the principles could be supported directly at the point of data collection, i.e. when the reader interrogates the tags. Table 4.2 lists the level of technical support for the FIP that our extended reader-to-tag air interface offers. Obviously, most of this support can also be achieved through non-technical means, e.g. a notice about tagreading taking place could also be simply announced through an easily noticeable sign. However, by incorporating such principles directly into the underlying protocol, both consumers and data collectors can more easily follow them, thus strengthening existing legal protection by providing the means to verify and thus enforce corresponding regulations.

# 4.3. Supporting the FIP in Existing RFID Standards

In this section we outline how existing RFID standards can be modified to satisfy the principles of *collection limitation*, *purpose specification*, *openness*, *use limitation*, and *accountability*. The extensions are illustrated using the ISO-Standard 18000 Part 6 Type A as an example, though they can equally well be applied to other RFID standards. Figures 4.2 and 4.3 show examples of the inventory command (Init\_round\_all) and process, respectively, as defined in the ISO-Standard 18000 Part 6 Type A [63]. The protocol uses framed ALOHA

Protocol extension	Init round all	SUID flag	Round size	CRC-5
1 bit	6 bits	1 bit	3 bits	5 bits

Figure 4.2.: The inventory command, Init\_round\_all, as specified in ISO 18000-6 Type A. The command frame consists of a field that indicates the number of time slots that are available for a reply (round size), various flags, and a cyclic redundancy check (CRC) to detect transmission errors [51].



Figure 4.3.: The inventory process, as specified in ISO 18000-6 Type A. The reader initiates a round of tag replies by issuing an Init\_round\_all command. Energized tags respond by selecting one of the available time slots at random to transmit their ID [51].

to singulate individual tags. After a tag has participated in the interrogation round and has been successfully identified, it transitions to an inventoried state (cf. Figure 4.4).

#### 4.3.1. Openness through Reader and Policy Identification

None of today's RFID standards allow tags to identify the reader they are communicating with. Anonymous broadcast by the reader is certainly desirable from a performance point of view, since the reader's goal is to identify as many tags by their UID as possible in a certain period of time. The transmission of any additional data such as the identification number of the reader will thus reduce the speed at which



Figure 4.4.: Simplified tag state transition diagram. As soon as tags enter the reader's RF field, they move into the "ready" state and reply to the reader's "inventory" command. Once the reader has inventoried tags in its read range, it can access them individually [51].

tags can be detected. Without knowledge about the device that is collecting data, it is, however, impossible to satisfy the principles of *openness* and *accountability*. In order to address these FIP requirements also at the air interface, we include a unique reader policy ID (RPID) in the reader's inventory command, which uniquely identifies both the reader and its operator, as well as the policy in place. Having an explicit reference to the policy allows us to provide additional information about a policy over a separate channel and also facilitates dispute resolution by allowing customers to directly identify the policy used.

The RPID itself is encoded in a three-tier format, specifying the following three fields: the data collector ID, the policy ID, and the reader ID (cf. Figure 4.5). With this structure, our solution follows closely the well-established EPC format and its general identifier encoding (GID-96) [31]. Even though we are not identifying products, but data collectors and their policies, this symmetry could potentially benefit the administration of the data collector IDs, as their identical format would allow data collectors to reuse the existing "General Manager Number" [31] of their EPCs (data collectors that do not already have such a number could acquire it similarly to the way they obtain an EPC identifier). Moreover, the existing ONS architecture [33], which provides a look-up functionality for captured EPCs, could also be used transparently to resolve our reader policy references.

The policy ID follows directly after the data collector ID, giving data collectors a 24 bit value for identifying policies. Data collectors are free to substructure this value in any way they like, as they can do for the last value, the actual reader device ID, which comprises 36 bits. Useful substructures would be a division by country, region, city, or store, thus simplifying both policy publishing and reader localization from this ID. In our prototype, we use the policy ID to acquire more detailed policy information over wireless LAN, while the reader ID is resolved to its designated approximate location, in order to allow the (manual) detection of reader ID spoofs (e.g. a reader of a retail outlet on 5th Ave. suddenly appearing ten blocks south of this address).



Figure 4.5.: The modified inventory command, Init\_round\_all, of ISO 18000-6 Type A featuring an additional field for the reader policy identifier, the purpose declaration, collection type, and an additional checksum (CRC) [51].

Figure 4.5 shows a summary of our reader and policy identification code, and illustrates its usage again using the inventory command of the ISO 18000 Part 6 Type A protocol as an example.

#### 4.3.2. Purpose Specification in the Inventory Command

The FIP require that the purpose for which personal data is collected should be specified no later than at the time of data collection. P3P addresses this issue by providing a list of 12 abstract purpose types that describe why data is being collected, relevant to the specific web site that the policy describes [25]. Although RFID needs to be treated slightly differently, in the sense that in most cases the user will be unaware of the data collection taking place, as well as of the actual data being collected, many of the P3P purpose definitions can be applied equally well to the RFID domain.

Type (Pos)	Description
access control $(0)$	Tag IDs are scanned for the purpose of access con-
	trol, e.g. by identifying a pass holder or by autho-
	rizing the validity of an access key.
anti-counterfeiting $(1)$	Readers read out data stored on the tags to assert
	the genuineness of a piece of merchandise.
anti-theft $(2)$	Readers scan for tags that are attached to items
	that have not been paid for.
asset management $(3)$	In contrast to inventory purposes, tags are read to
_ 、 ,	provide a picture of the whereabouts of assets, in-
	stead of monitoring changing stock quantities.
contact $(4)$	Tag contents are read out in order to determine
	a contact channel to the customer, e.g. a mobile
	phone number or email address.
current (5)	Tags are read to provide a service that was explic-
	itly desired by the individual, e.g. when placing
	shopping items on a kiosk in order to calculate to-
	tals, or for disabling (killing) tags.
development (6)	This purpose should be used during system testing
1 ()	and development only.
emergency services (7)	The system is monitoring tags in order to provide
5, ()	rescue workers with occupancy information.
inventory (8)	A shelf monitoring its contents, e.g. in order to
5 ( )	provide out-of-stock notices to a central system.
legal (9)	Law enforcement or other legal obligations require
	the system owner to read out tag IDs. Additional
	information required on legal grounds should be
	made available to the customer.
payment (10)	The current action involves payment, e.g. at check-
	out when tag IDs are read for billing purposes.
profiling (11-13)	Data is collected for profiling or ad-hoc personal-
	ization. See Table 4.4 for individual values.
repairs and returns $(14)$	Warranty and manufacturing details are read out
1	in order to facilitate or speed up a repair or return
	process.
other $(15)$	None of the above purposes fits. Further informa-
	tion should be accessible, e.g. in the form of a sign
	or explicit contractual agreement.
()	tion should be accessible, e.g. in the form of a sign or explicit contractual agreement.

Table 4.3.: RFID purposes declarations.Data collectors can combine 15 different<br/>purpose declarations for RFID reader queries.

Contrary to Web services, however, some purposes such as *admin* or *current* are much more difficult to assess in an RFID environment. For example, the current purpose is usually implicitly defined by the Web interaction the user is currently experiencing, e.g. the shopping cart checkout in a Web shop, while administration is usually defined by keeping Web server log files. In an RFID context, however, many different "current" or "admin" purposes can be envisioned: a smart shelf might issue read commands for inventory purposes (in a supermarket) or for asset tracking (e.g. for multimedia equipment that employees can check out from a central repository), both of which could be called administrative purposes. "Current" purposes can equally vary, from a payment purpose at a self check-out station to a repair and return purpose at a customer information station.

Consequently, we have expanded some of the existing P3P purposes while dropping others, in order to better reflect the more implicit interactions present in RFID systems. Table 4.3 lists the 14 purposes we identified as useful declarations in this context, even though additional purposes might become necessary in the future. This list is therefore only an initial suggestion that should be repeatedly validated by realworld prototypes, and subsequently standardized by an appropriate standardization body.

Apart from the "profiling" purpose, all purposes are encoded as single bit values that can be arbitrarily combined in our 16 bit number, indicating that data are collected for multiple purposes. The profiling purpose uses three bits to encode one of five possible profiling purpose types that are mutually exclusive (see table 4.4).

For example, a smart shelf application that monitors its contents for out-of-stock warnings, as well as providing data for anonymous instore movement information (e.g. to see where consumers spend most of their time), would need to declare both the "inventory" and the "pseudo-analysis"-profiling purposes. A corresponding smart shopping cart that would provide customers with shopping suggestions, based on its contents, would declare "pseudo-decision"-profiling. And a selfcheckout station that allows customers to wirelessly pay for their goods, while also associating the purchased items with the customer's loyalty card, would consequently declare the "payment," "anti-theft," and "individual-decision" profiling purposes.

Type (Bits)	Description
ad-hoc-tailoring (011)	This applies to immediate and anonymous tailoring, e.g. providing shopping recommendations based on the current content of a shopping basket, or sug-
pseudo-analysis (100)	tomer has taken into the dressing rooms. The collected data are used to learn about the in- terests or other characteristics of individuals. This
	may help to reveal the interests of visitors in differ- ent areas of a store. For example a store's shelves could be newly arranged based on the collected ag- gregated data.
pseudo-decision (101)	This information will be used to make customiza- tion decisions based on the interests of individuals, without actually identifying them. For example, a shop could suggest items to a customer based on his or her previous visits (without actually identifying
individual-analysis (110)	that person). The data collected is used in combination with the identified data of an individual, allowing a profile of a certain customer to be generated. This could help to reveal the interests of visitors based on their age, social situation, or other relevant demographic data. Identification could occur in combination
individual-decision (111)	with a consumer or credit card. The information is used to determine individual preferences and to link them with identified data. This profile allows personalized suggestions, based on the individual's interests collected from previ- ous visits, combined with personal information, e.g. from a consumer loyalty card.

Table 4.4.: *Profiling purposes.* Profiling purposes are mutually exclusive, as profiling types lower in the table (i.e., with higher bit-codes) can potentially include all of the above types.

#### 4.3.3. Use Limitation through Collection Types

The principle of RFID reader-to-tag interactions (i.e. readers issuing an inventory command and tags replying with their IDs) makes it difficult to create privacy-friendly monitoring applications even if no identifying tag information needs to be collected as part of the envisioned application. Imagine an RFID system that tries to keep track of the number of people on a certain station platform, in order to avoid overcrowding. Even though RFID tags entering and exiting the area might reply to reader commands with their IDs, the application only needs to keep track of individual tags (e.g. an RFID-based train pass) without having to actually know their specific ID. Additionally, even when identifying information is collected, consumers will typically become much more concerned if this information is not only used locally, but also correlated across multiple readers in order to track an item's (or a person's) movements over time.

To allow data collectors to differentiate between various collection needs, i.e. whether or not they actually require the serial number of individual tags, or whether they intend to track multiple occurrences of the same tag across different locations, we additionally define four distinct collection practices that must be declared as part of a reader's inventory command:

- 1. Anonymous Monitoring: Collecting state information about the items in the vicinity of a particular location, without the need to actual identify tags by their unique serial number. Examples would be simple sensor applications (e.g. an automatic door opener) or counting tasks (e.g. monitoring the number of items in a certain area).
- 2. Local Identification: Tag IDs are collected in order to provide a localized service, e.g. a smart medicine cabinet or smart fridge that monitors its contents. Although unique IDs are collected (e.g. for resolving them to human readable descriptions), the application does not require (nor attempt) the correlation of events across different locations.
- 3. *Item Tracking:* Collecting information about the location of an item for the purpose of monitoring its movements. Note that this potentially enables people to be tracked through constellations. However, in order to differentiate between these different intentions, the separate "tracking person" declaration should be used, if people are tracked by the items they carry.
- 4. *Person Tracking:* Collecting information about the location of a person. Note that although item-level tracking can potentially reveal the location of a person, data collectors will only need to declare this if they actually collect RFID tag information for this purpose. It is up to legal frameworks to force data collectors to anonymize item-tracking data so that it cannot be used for person tracking.

Together with a corresponding purpose, collection declarations further facilitate the accurate assessment of any RFID scan event. This not only helps data subjects to better understand the *intentions* behind a data collection process, but can also be used to selectively allow tags to remain *anonymous* whenever possible. Anonymous replies are already part of some RFID protocols, e.g. ISO 18000 Part 6 Type A, although the reason for using them is usually, again, efficiency, not data privacy. To detect collisions, a 64 bit or longer unique ID is usually not needed and just decreases the number of individual tags that can be successfully detected per unit of time. The anti-collision routine can thus first use the tag's random short identifier to single it out from the set of present tags, before requesting additional data, which might include the unique but static serial number. We propose that this kind of an anti-collision protocol could become the default, whenever "anonymous monitoring" intentions are declared, thus explicitly providing tag anonymity and unlinkability.

Even without any specific support in the tags themselves: declaring, say, "local identification" would still provide the data subject with the additional level of assurance that her movements would not be tracked across different locations (though this might not preclude the keeping of log files that could be later combined, e.g. as part of a criminal investigation). Obviously, none of these declarations constitute proof that the data collector stating them is actually following them. However, as with the purpose declarations, any explicit privacy policy declaration provides a lever to threaten wrongdoers with legal action – just as is the case with today's printed policies.

With reference to the examples from the previous section, the smart shelf tracking inventory and the performance of anonymous movement analysis of customers within the store would thus need to declare a collection practice of "person tracking", even though these traces are anonymous (pseudo-analysis). The smart shopping cart would use "local identification", as it would use the identity of the items in the cart to locally decide what other products to suggest to the user. Note that it does not matter whether this decision process is actually done on the shopping cart itself or wirelessly via a remote system, as long as the tracked tags are not correlated to other carts or shelves. A smart check-out station would need to declare "person tracking" again, if a consumer loyalty card is scanned at the point of sale.



Figure 4.6.: The modified inventory process. The reader first selects a tag population, before initiating a round of tag replies by issuing the modified Init\_round\_all command. Previously selected tags (tag 1, 2 and 4) respond in a randomly chosen slot [51].

#### 4.3.4. Collection Limitation by Appropriate Tag Selection

The first of the fair information principles requires data collectors to limit the amount of data they collect to what is absolutely necessary (today, the EU directive makes this a legal requirement in most European countries). Consequently, rather than asking *any* tag present to respond to a reader query and then filtering out the tags of interest on the application level, we want readers to limit their initial query to target only relevant tags in the first place, thus realizing the collection limitation principle already at the protocol level.

As an example of how this would work in practice, let us look at the frequently considered usage scenario of a supermarket smart shelf, whose purpose is to detect whether it is stocked with sufficient supplies of a particular item. Instead of issuing indiscriminate read commands, which might also pick up tags in the clothing of nearby shoppers, the shelf reader will target only tags of products stacked on the shelf, such as a particular brand of razor blades. Optionally, the shelf reader could occasionally run a separate request that targets *all* of the supermarket's products in order to detect misplaced items.

To implement this functionality in our reader-to-tag-protocol, we make use of a similar mechanism that is typically used to singularize a particular tag from a set of tags in range (e.g. the Group-Select and Group-Unselect commands in ISO 18000 Part 6 Type B). However,



Figure 4.7.: Modified tag state transition diagram. After getting energized the tag enters the ready unselected state. The tag moves into the selected state once it receives a matching "select" command. Only selected tags will respond to an "inventory" command by the reader [51].

instead of using a selection mask to facilitate and potentially speed up the inventory process, we use selection masks to restrict tag ID collection by the reader to relevant tags for privacy reasons.

Once tags appear in the range of a reader and get energized, they initially begin in an "unselected" state. Unselected tags will need to be explicitly selected before replying to any inventory, read or write command from the reader. Tags become selected only after receiving a select mask that matches their data in memory. Readers thus begin any command cycle with one or more select commands that first determine the tag population that is the target of the query (see figure 4.6). Once selected tags have been "inventoried", readers can issue actual access commands (see Figure 4.7).

The Select command contains the following parameters (see Figure 4.8):

- Pointer, length, and mask (PLM). Pointer and length address a certain tag memory range. The mask, which must be "length" bits long, contains a bit string that the tag must compare against the contents of the specified memory location.
- Selection type. The selection type indicates whether tags that

match the PLM should enter the selected state or return to the ready, but unselected state.

Note that an appropriate selection of tags that fulfils the requirement of the collection limitation principle will only be feasible if the tag IDs follow a known structure that allows for a certain grouping, e.g. a common prefix for a certain product from a particular manufacturer. This is the case in the currently favoured EPC system, where ID ranges are grouped by manufacturer ID and product type. If there is no such information encoded in the identifier, it needs to be available in the remaining portion of the tag memory and accessible during the selection process, as random tag IDs would be difficult to select efficiently.

Protocol extension	Select	State flag	Pointer	Mask length	Mask value	CRC16
1 bit	6 bit	1 bit	8 bits	8 bits	variable	16 bits

Figure 4.8.: The new Select command enables readers to select a subset of tags within the read range. The state flag indicates whether a tag with a matching mask should enter or leave the selected state [51].

In the following section, we show how the feature set outlined in this section – i.e. the reader policy ID, the purpose and collection type declaration, and the selection mask – can significantly increase transparency in today's RFID scenarios.

# 4.4. Watchdog Tag

In order to make full use of the additional information now present in the reader protocol, we use a so-called "watchdog tag" to provide transparency to the otherwise invisible tag detection process. Simply speaking, the watchdog tag is a sophisticated version of an ordinary tag, as it features an additional battery, a small screen, and potentially even a long-range communication channel. The watchdog tag's main task is to decode the commands transmitted by a reader, and make them available on the screen of the device for inspection by the user, as shown in Figure 4.9, or to log all data transfers and provide consumers with detailed summaries whenever needed. While the watchdog tag could be carried by the user as a separate device, its functionality

🌮 Internet Explorer		🎊 Internet Explorer		
ndog Tag	Watchde	og Tag		
13:18:24, 03/06/2004	Time/Date	13:18:24, 03/06/2004		
5F.4A886EC.8EC947.24A68E4F6	Data Collector	Example Store Inc.		
Inventory, Pseudo-decision	Policy ID	8EC947		
Person Tracking	Reader ID	24A68E4F6		
, and the second	Reader Location	Aisle 6		
**,7B3E747,3DBA49,********* **,7B3E747,3D91E1,*********		98, Main Street Example City, EC 21508		
**.7B3E747.3D86B4.********	Purpose	Inventory, Pseudo-decisior		
	Collection Type	Person Tracking		
	Target Selection	Close Shave Men Close Shave Lady Close Shave Super		
	Net Explorer     13:18:24, 03/06/2004     5F.4A886EC.8EC947.24A68E4F6     Inventory, Pseudo-decision     Person Tracking     **.7B3E747.3DBA49.*********     **.7B3E747.3D91E1.*********     **.7B3E747.3D86B4.*********	Internet   Inventory, Pseudo-decision   Person Tracking   **.7B3E747.3DBA49.********   **.7B3E747.3D91E1.********   Purpose   Collection Type   Target   Selection		

Figure 4.9.: Watchdog Tag screen shots. The screen shot on the left shows data collected by the watchdog tag over the RFID channel. If a separate communication channel is available, these raw data can be resolved to a more expressive, human readable format as shown in the screen shot on the right [51].

could also be integrated into a mobile phone, allowing it to leverage the existing display, battery, memory capacity and long-range communication features of the phone.

Without the privacy features in the protocol, the watchdog tag would only be able to inform the user that some anonymous reader is scanning for tags in a certain vicinity. Due to the privacy features introduced in the RFID protocol, this notice can now include the operator's ID, the purpose and type of data collection, and the target range of tags. If a separate long range communication channel is available (e.g. wireless LAN or GSM), the watchdog tag can additionally translate the data transmitted over the RFID channel into a more expressive format, as shown in Figure 4.9. Of course, any such lookup would require an appropriate backend infrastructure, e.g. the ONS architecture developed by the Auto-ID Center [33]. In addition, providing the reader location in a human readable format allows for a simple, manual detection of reader ID spoofs. More sophisticated watchdog tags featuring an integrated location system could potentially detect reader ID spoofing automatically.

The above screen shots were taken from our initial watchdog prototype, which serves as our design test bed for our protocol extension. Built on top of a standard Windows CE PDA, it uses the built-in wireless LAN to retrieve human readable descriptions. While we are currently working on a separate antenna design that allows us to interface our PDA directly with the RFID reader's communication channel, we have so far been simulating the complete RFID protocol over the wireless LAN as well (with a PC posing as a virtual RFID reader).

### 4.5. Discussion and Future Work

Even with our proposed protocol extensions, unauthorized read attempts by readers not conforming to our specification will still be possible. While consumers carrying a watchdog tag might be able to actively jam or block the tag-to-reader communication [67], for example based on user preferences regarding the reader's ID (e.g. following an online lookup), the average consumer would still need to resort to explicitly disabling her tags in order to completely prevent misuse. However, even without any additional devices, the required selection mechanism at the protocol level supports the core principle of *collection limitation*, while the compulsory identification string facilitates the principles of *openness* and *accountability*, thus providing the same level of protection as today's compulsory forms, signs, and placards announcing the privacy policy of the data collector. While they might be ignored in our daily routine, their presence forms an important legal lever once a dispute over the proper use of personal data arises.

Our proposed protocol extensions are easily realized even with today's readers, as they only require updates to the reader's firmware, since the physical layer remains unaltered. While tags would require changes to their logic, these should be straightforward to implement, as the physical layer is not affected and only slight alterations to the medium access layer and the command set would be necessary. Our extensions do, however, affect the performance of an RFID system. The addition of the RPID, purpose code, and collection type require the additional transmission of 130 bits. At a data transfer rate of 30 kBit/s, typical for reader-to-tag signalling of systems operating in the UHF band, the execution time of any command is prolonged by 4.3 ms. This delay is comparable to the time it takes for a single tag to reply with its ID, assuming symmetrical data transfer rates. In modern RFID systems that typically read several dozens, if not hundreds of tags at a time, losing a single tag slot thus seems negligible. For an RFID system that features a slow data transfer rate, e.g. 1.6 kBit/s as specified in ISO 15693 (HF), the delay is more significant, approximately 80 ms.

However, in many situations such a delay would be outweighed by the shortened reply times, as the Select command allows the reader to ignore tag IDs that are of no interest to the application in the first place. Newly arriving tags in the read range will have to wait for the next select command before they can be inventoried by a reader.

Future tags could implement the resurrecting duckling model proposed by Stajano [116], where tags would only respond to a "mother" reader, but ignore requests from all others. Instead of killing tags at checkout [8], stores would transfer their "mother" rights to the customer's reader, thus allowing for safe post-sales RFID usage. Additionally, such "mother" readers could inhibit replies by "their" tags for non-desired purposes and intentions to unknown readers by programming the tags accordingly.

While technical approaches to addressing the threats to privacy resulting from the proliferation of RFID systems, such as the one presented in this thesis, can help to protect the rights of the individual, their contribution to influencing the negative consumer attitude towards RFID should not be overestimated. In order to positively influence risk perception, education about the true capabilities of the technology and the development of a trust relationship between data collectors and individuals are required as well. In the long term, the availability of suitable applications that let the consumer benefit from the proliferation of RFID tags on everyday items might play the most important role in addressing consumer fears.

## 4.6. Summary

The work presented in this chapter is intended to help build future privacy-aware RFID standards that are not only optimized for performance and low cost, but also satisfy the fair information principles (FIP). The key idea of our approach is to augment the communication protocol between RFID readers and tags with a feature set that identifies the reader to provide *openness* and *accountability*, enables RFID operators to disclose a *purpose specification* and collection type, and supports a selection mechanism to facilitate the principle of *collection limitation*. In concert with a watchdog tag or a similar device, selective jamming can support the principle of explicit *consent*, while the integration of readers into an overarching privacy-infrastructure such as "pawS" [75] would allow the enforcement of *use limitation*, *data quality*, and *participation* principles. Such infrastructures provide a readily available, practical and simple solution to many of today's RFID privacy concerns, while the possible integration of the watchdog tag functionality into future mobile phones might in future even make the detection of an RFID reader, its policy, and location as easy as detecting signal strength and operator IDs on a mobile phone today.

# RFID Middleware Design – Addressing Both Applications Needs and RFID Constraints

The proliferation of readers and tags not only mandates approaches that address performance and privacy issues, but also requires an application agnostic middleware that helps to manage large reader deployments and aggregates the captured RFID data. In novel application domains, such as supply chain management and logistics, there is no longer a 1-to-1 relationship between reader and application instance. In these domains many readers distributed across factories, warehouses, and distribution centers capture RFID data that need to be disseminated to a variety of applications. Each of these applications has different needs with respect to notification latencies, tag populations of interest, and aggregate types.

In this chapter we analyze the requirements an RFID middleware component should meet in order to manage large deployments of readers and the amount of data these readers capture. The main contribution in this chapter is a middleware implementation, the RFIDStack, that addresses these requirements, but also considers the constraints imposed by RFID technologies. We illustrate in particular how an architecture built around a content-based routing system can deal with the application needs and the constraints of passive RFID technology.

The chapter is organized as follows: Section 5.1 discusses the application requirements RFID middleware should address. Section 5.2 provides a brief overview of RFID technology and outlines the constraints imposed by the characteristics of RFID. In Section 5.3, we discuss the corresponding design implications. In Section 5.4, we present the RFIDStack, an RFID middleware platform that addresses the requirements and constraints outlined previously. Before we present related work in Section 5.6, we discuss the strengths and limitations of our implementations in Section 5.5. The chapter concludes with a summary of our contribution in Section 5.7.



Figure 5.1.: *RFID readers in a distribution center feeding captured data to different applications.* 

## 5.1. Application Requirements

RFID technology promises to help automate many processes in supply chains [27], but also in other domains such as maintenance [80], and production [86]. The common feature of all these applications is that they benefit from dedicated RFID middleware that provides data processing, routing, and reader management functionality. Such a middleware can for example preprocess the raw data captured by the readers before these applications interpret the data and turn them into more meaningful information. In traditional RFID applications, there was little need for an RFID middleware because the RFID readers were not networked, e.g. readers in animal tagging applications or car immobilizers, and the RFID data were only consumed by a single application, e.g. to grant access to a building.

The following paragraph outlines a supply chain scenario – a distribution center of a pharmaceutical company (cf. Figure 5.1) – that represents common applications of RFID technology. The corresponding use case diagram is shown in Figure 5.2. We use this scenario and the use cases to derive the application requirements that an RFID middleware should address.



Figure 5.2.: Use case diagram of RFID usage scenario in a distribution center

Goods arrive at the distribution centre and are identified by the reader at the dock door. The captured information is transmitted at once to the supply chain management (SCM) system, which provides track and trace functionality. The goods are placed in the warehouse, where RFID readers scan the inventory. At regular intervals the readers trigger the legacy warehouse management systems to update the inventory counts of the corresponding product categories. Goods are picked from the warehouse and packed at the corresponding pick and pack station. An RFID reader monitors the tagged items currently packed so that a local application can support staff with a near real-time comparison of items actually packed and the items on the pack list. Before the shipments are loaded into the trucks at the loading dock, they pass a reader that scans the tag on the pallet and passes this information to the supply chain management system, which sends an advance shipping notice (ASN) to the recipient of the shipment. On a nightly basis all tag IDs of the items packed and shipped are transmitted to the healthcare authorities to comply with pedigree legislation. To maintain an adequate service level, the readers report exceptions to a remote system monitor.

Filter by	Description
Reader Identifier	This filter type allows the application to
	specify that it is only interested in data
	from a particular set of readers.
Tag Identifier and Data	The application can define the tag popula-
	tion that it is interested in, e.g. a restric-
	tion to tags attached to pallets.

Table 5.1.: Filter types required [48].

*RFID* system integrators can inspect a configuration of a reader and reconfigure reader devices remotely.

Based on an analysis of many different RFID applications including the above, we have identified the following requirements an RFID middleware should meet:

Data dissemination. The information captured by a reader is usually of interest not only to a single application, but to a diverse set of applications across an organization and its business partners. The captured RFID data must thus be broadcast to the entities that indicated an interest in the data. Different latencies need to be supported, since the desired notification latency depends upon the application type. Applications that need to respond immediately to local interaction with the physical objects require a short notification latency that is comparable to the observation latency. Legacy applications that are not designed to handle streaming data might need to receive batched updates on a daily schedule.

Data filtering and aggregation. A common feature of all applications that make use of the captured data is the desire to receive filtered and aggregated RFID events rather than raw streams of RFID data. Different applications are interested in a different subset of the total data captured, based on the reader and the tag involved. Since RFID permits identification at the instance-level rather than at the class-level, the fine-grained RFID data need to be aggregated for applications that cannot deal with the increased granularity. The different filter and aggregation types that should be supported by an RFID middleware are listed in Table 5.1 and 5.2.

Writing to a tag. Some tags feature not only memory space for an identifier, but for additional data. Middleware solutions should thus provide means to write to and read from this additional memory. This additional memory can then be used to store application data such as expiry dates in order to facilitate data exchange, where no network

Aggregate type	Description
Entry & Exit	This aggregate type reduces a number of
	successful reads of a tag to the best esti-
	mate of when the tag appeared and disap-
	peared from the read range.
Count	Applications can prefer to receive informa-
	tion about the total number of items of a
	specific detected category rather than the
	individual ID of each object.
Passage	This event indicates the direction in which
	a tagged object is moved as a tag moves
	from one reader to another. Applications
	prefer to receive a passage event rather
	than being forced to interpret a sequence
	of entry and exit events from two individ-
	ual readers.
Virtual readers	When an application does not distinguish
	between two readers, this aggregate type
	allows it to virtually join their read range.

Table 5.2.: Aggregate types required [48].

access is available. In a broader sense, writing to a tag also includes the initial write to the tag to program its ID, and killing a tag to permanently disable it.

Reader integration in IT-service management. The proliferation of readers mandates their integration into an existing IT-service management concept that performs incident, change, and configuration management. This includes in particular reader management functionality to monitor the health of the reader and to retrieve and to modify the reader configuration remotely.

*Privacy.* In the previous chapter we saw that the intended deployment of RFID-based tracking solutions in today's retail environments epitomizes for many the dangers of an Orwellian future: unnoticed by consumers, embedded RFID tags in our personal devices, clothes, and groceries can unknowingly be triggered to reply with their ID and other information, potentially allowing for a fine-grained yet unobtrusive surveillance mechanism that would pervade large parts of our lives. An RFID middleware should consider these consumer fears and the legal guidelines that apply for data collection. If future RFID communication protocols support the feature set proposed in the previous chapter, the RFID middleware will also need to inform the reader about the purpose of data collection and about the reader and policy identifier.

Other application requirements that relate to security and perfor-



Figure 5.3.: Overview RFID Middleware (adopted from [104]). Our analysis focuses on the device and data management aspects.

mance are not discussed here in detail, since they are not unique to RFID and have already been mentioned in [16, 121]. The data and device management requirements listed above result in RFID data which are coherent and less noisy. From an application perspective, it is also desirable to provide a mechanism that interprets the captured RFID in a given business context and that turns the low level RFID event into the corresponding application events (cf. Figure 5.3). The relevant requirements have been studied in detail by a number of researchers [16, 23, 103, 126]. Our present analysis is thus restricted to the application requirements towards the device and data management part of an RFID middleware, which reduces the data volumes and filters the noise, but does not make inferences about the captured RFID data.

# 5.2. Constraints Imposed by the Characteristics of RFID

Before describing how the application requirements listed in the previous section can be met, we will outline the constraints imposed by the characteristics of RFID. We believe that these constraints have a significant impact on the design of an RFID middleware and introduce aspects that are unique to the RFID domain. Any RFID middleware design that fails to include these will result in inefficient data capture and consequently low quality data.

In Chapter 2, we saw that there are a wide variety of different RFID systems that address the requirements of individual applications, e.g. with respect to range, transmission speed, susceptibility to environmental interference and cost. Different passive RFID systems can be distinguished by the frequency band they operate in, the coding, modulation, and medium access techniques used and the supported command set. Since RFID design is generally driven by tradeoffs between different properties – e.g. read range, data transfer rates, identification speed, tag and reader form factor – there is no single RFID technology that proves superior in all possible application domains.

For the purpose of this chapter, the characteristics common to all passive RFID systems are most important, since they have a strong impact on RFID middleware design. These include:

Limited communication bandwidth. RFID systems rely on the availability of unlicensed frequency bands. Radio transmissions in these unlicensed bands are governed by local radio regulations. In the UHF frequency band, which is particularly suitable for supply chain applications due to its superior read range, the European radio regulations ETSI EN 302 208 permit the use of fifteen 200 kHz-wide channels between 865.0 MHz and 868.0 MHz by RFID readers [35]. However, there are only 10 sub-bands in which the maximum radiated power of 2 W ERP is available, and which thus allow for the largest read range. To minimise interference in these bands, ETSI EN 302 208 introduces the concept of "Listen Before Talk (LBT)". Prior to transmission, the RFID reader has to listen for the presence of another signal within the intended sub-band of transmission.

The threshold level set for this Listen-Before-Talk scheme, -96dBm in the worst case, implies that it is unlikely that two readers will be able to operate in the same channel at the same time within one facility [39, 81]. Since large distribution centres might need to operate as many as 100 readers, it is evident that readers need to co-ordinate their activities somehow to avoid missing tags that pass by while the reader is not operating.

Another constraint is the bandwidth available per channel, which limits the data transmission rate between readers and tags. It restricts



Figure 5.4.: Select command in the EPCglobal UHF Class 1 Generation 2 Protocol [32] (Source: EPCglobal). A particular tag population is selected before the inventory process is initiated with the Query command

the number of tags that can typically be identified per second to the order of tens or hundreds – even with efficient transmission strategies as outlined in the previous chapter. A tag labelling a shipment that arrives on a pallet carrying more than a thousand tagged items can thus easily be missed, unless the identification of the "shipment" tag is prioritized. To facilitate the latter, some RFID protocols permit the selection of a certain group of tags based on data stored on the tag [32, 65]. In [32], there is a *Select* command that selects a tag population based on a number of different criteria, before the inventory process is initiated by the reader (cf. Figure 5.4). The result is that the limited communication bandwidth available is utilized efficiently, since only the tag population of interest is inventoried.

These bandwidth restrictions apply not only to systems operating at UHF, but also to other frequency bands. In Chapter 2, we saw that the 13.56 MHz ISM band provides a single frequency channel only, but at the same time features propagation characteristics that result in a shorter reuse distance.

Reliability issues. Due to field nulls, e.g. caused by multipath fading (at UHF) or absorption by objects in the range of the reader, there is no guarantee that a tag will stay powered while in the assumed range of the reader. The result is a false negative read. Such false negative reads can also be caused by collisions on the air interface and by transmission errors [17]. The false negative reads result in the fact that a tag will not be continuously detected on consecutive scans by a reader (cf. Figure 5.5), although the tag remains in the assumed range


Figure 5.5.: Example of a series of tag reads showing regular reads by an RFID reader (HF) of six tags which are present in its read range [17]. The "missing" reads indicate that not all tags are detected on each scan (false negative reads). The figure also illustrates two different causes of false negative reads: interference problems (1) and collisions on the air interface (2).

of the reader.

Tag memory. The design of RFID middleware is also impacted by the memory structure on the tags. The memory on the microchip embedded in the tag usually contains a unique identifier. This can either be a random serial number or an identifier code that incorporates information about the tagged object, e.g. its manufacturer. Most microchips also feature small amounts of additional random access memory. Due to the increased power required to write to the EEPROM on the microchip, the maximum distance between reader and tag for a "write" operation is a fraction of that for a "read" operation. In [69], Karthaus et al. present a UHF RFID tag microchip with EEPROM that consumes 12.5  $\mu$ W in read mode and 35  $\mu$ W in write mode.

Heterogeneous reader landscape. The diverse computing and net-

working capabilities of readers is also an important RFID consideration. Low cost readers usually support only a single antenna and a serial RS232 interface. These reader types are connected to a computer which hosts the application directly or forwards the captured data over a network connection. More sophisticated reader devices support several antennas, a TCP host interface, and ample computing resources for on-device data processing. All RFID readers can usually also be parameterized to some extent. This covers network interface parameters such as the reader's ID and port, RF parameters such as transmit power, frequency hop sequences, noise levels, and air protocol specific parameters like data rate, coding type, and MAC properties.

## 5.3. Design Implications

Due to the diverse set of applications that consume the captured RFID data and the resource limitations of RFID readers, an event-based middleware that decouples readers and applications is appropriate for RFID. Readers produce RFID events and deliver them to the messaging system; it is the responsibility of the messaging system to get the messages to their intended destinations (cf. Figure 5.6). In such a publish/subscribe concept the producer, the reader, does not need to track which applications are supposed to receive a certain message. Likewise, applications consuming RFID data do not need to maintain communication channels with individual readers, but can simply specify which events they are interested in by submitting subscriptions to the messaging system. The result is that there is no need for the readers to hold a reference to subscribing applications, nor do they need to know how many subscribers there are. Similarly, applications need only a single reference to the event service.

However, the application requirements and constraints characteristic for the RFID domain mandate a set of special features. In this section, we outline these design features. They represent the foundations of our RFIDStack implementation, which is presented in the following section.

Full content-based routing. Applications are only interested in a subset of the total data captured. This subset can be specified using reader ID, tag ID, and possibly tag data (cf. with Figure 5.7 and Table 5.1). In order to carry out the filtering within the messaging system itself, the nature of RFID events demands the use of a messaging system that provides full content-based routing rather than channel or topic-based



Figure 5.6.: Decoupling RFID readers and applications by a publish/subscribe messaging system.

routing. Otherwise, the entire message content would need to be replicated in the subject, or alternatively applications would be forced to carry out some of the filtering locally. They would for example need to subscribe to a 'reader' topic feed and discard the messages featuring tags of no interest. Hierarchical addressing schemes that allow for the organization of topics according to a containment relationship provide more flexibility than simple channel based systems, which offer flat addressing, where different channels represent disconnect event spaces [36]. In such topic-based event routing services, the notifications can be filtered if the reader ID and tag ID are hierarchically organized as topics. However, it remains impossible to filter notifications within the event router according to the data on the tag or the aggregate type, e.g. certain aggregate events. Unless a content-based publish/subscribe mechanism is used, some filtering will thus have to be carried out within the application.

Scope and expressiveness of the subscription language. Subscription languages of notification services can be classified based on their scope and expressive power [21]. On the one hand, there are those classes of subscription languages that consider only a single field in a single

Notification:	
readerID:	urn:epc:id:gid:12.244.345
tagID:	urn:epc:id:gid-96:20.300.4000
timestamp:	2002-11-06T13:04:34
Subscription:	
<pre>(readerID == (begins-with</pre>	<pre>''urn:epc:id:gid:12.244.345'' &amp;&amp; (tagID, ''urn:epc:id:gid-96:20.'')))</pre>

Figure 5.7.: *RFID read event notification and matching subscription*. In its simplest form, the notification comprises information about the reader which detected a certain tag at a certain point of time.

notification. These are commonly found in channel- and topic-based event routing systems, such as TIBCO [122]. As shown in the previous paragraph, due to the structure of RFID read events they are not suitable for an RFID middleware. By the definition of content-based routing systems, the corresponding subscription languages usually consider multiple fields in a single notification. However, there are also some subscription languages that support compound filters or patterns which are matched against multiple notifications based on both their attribute data and on the combination they form [83, 10].

The expressive power of a subscription language is concerned with the kind of operators which can be used to form subscription predicates. The least expressive subscription languages support a simple equality predicate only, while most subscription languages of content-based routing systems feature logical and arithmetic operators and string comparisons. Some subscription languages even allow for user-defined operators [20], where subscribers can provide a predicate object able to filter events at runtime. The disadvantage of such executable code is that the resulting filters are difficult to optimize [36].

To filter the raw RFID data captured, the subscription language of an RFID publish/subscribe system should consider all fields in a notification and support logical operators and string comparisons (cf. Figure 5.7). The aggregation functionality described in Table 5.2 could possibly be realized with a subscription language that supports patterns which are matched against one or more notifications, based on both their attribute data and on the combination they form. The appropriate pattern might thus correlate a number of notifications to



(b) with subscription feedback

Figure 5.8.: Decoupling RFID readers and applications by means of a publish/subscribe messaging system. In order to filter over the air interfaces and use the scarce bandwidth efficiently, it is essential that the RFID readers receive feedback about the subscriptions currently registered with the messaging service.

detect, e.g. a passage event (cf. Table 5.2). While user-defined operators have been used in some RFID middleware implementations, e.g. in [91], we believe that they are not necessary to meet the application requirements for filtered and aggregated RFID data.

Subscription feedback mechanism. While the decoupling of RFID event consumers and producers is desirable, the limited bandwidth available to RFID requires a feedback mechanism for readers to determine whether applications are interested in the RFID data they produce (cf. Figure 5.8). Such feedback can then lead to an appropriate adaptation of the queries exercised by the reader over the air interface, e.g. targeting a particular tag population at a higher sampling rate or switching off completely to make the bandwidth available to another reader. The filtering of the RFID data is then no longer carried out in the event router, but over the air interface (cf. Figure 5.4). If such a feedback mechanism is missing and readers simply co-ordinate access to the radio channel independently of the application's needs, the quality of the captured data will suffer. A reader configured to read any tag might miss a fast-moving pallet tag – potentially the only tag an application is interested in. Likewise, a reader listening for tag replies and occupying a radio channel – though no application desires its data – will potentially cause a dock door reader unable to find a free channel to miss an outgoing shipment. Such a subscription feedback mechanism is also beneficial from a privacy perspective. It prevents RFID readers from collecting data and possibly logging those data without any specific application indicating an interest in the data.

*Reliable Messaging* On the one hand, there are applications, e.g. a point of sale system or a magic medicine cabinet, that request the raw RFID data stream with a minimal notification latency. Since these applications are only interested in the current state and need to respond to the detection of individual tags in "real-time", there is little need to retain messages matching their subscriptions, which accumulate while the connection is being re-established. In this type of application, historical information is irrelevant and it is not necessary to time-decouple event subscriber and producer. The situation is different for applications that receive a batched update comprising sometimes "a day's worth of RFID events". These require the messaging service to guarantee that a message will be stored until the consumer is able to receive it. Here, it is essential that the event service decouples the readers and applications in space and time. Other qualities of service commonly considered in publish/subscribe systems, such as priorities and transactions [36], are not required for RFID middleware.

## 5.4. Implementation – the RFIDStack

In this section, we present the RFIDStack, a middleware platform that we implemented based on the design considerations discussed in the previous section. We also illustrate how the RFIDStack addresses the application requirements and RFID constraints mentioned in Section 5.1 and 5.2. We show how the RFIDStack provides means to utilize the restricted bandwidth available to RFID systems efficiently, given the application needs for filtered and aggregated data. There is also dedicated support for the heterogeneous reader landscape and the different memory structures on RFID tags.



Figure 5.9.: Deployment diagram of the our RFID middleware implementation [48]. The diagram shows how an event-based messaging system decouples applications and readers. It also features the virtual tag memory system and the surrogate concept that address the limitations of tag memory and the heterogeneous reader landscape respectively. To utilize the scarce bandwidth effectively, the filtering is done on the air interface where possible or otherwise directly in the messaging system.

# 5.4.1. Data Dissemination, Filtering and Aggregation with the Elvin Message Router

The messaging component of the RFIDStack, our RFID middleware implementation, relies on the content-based router Elvin [112]. Elvin can be described as a pure notification service [7], in which notification producers push messages to the service, which in turn delivers them asynchronously to consumers. When a notification arrives at the routing service from the producer, it is compared against the subscriptions currently registered by the different consumers and the notification is forwarded to those consumers whose subscriptions it satisfies. Elvin routes undirected, dynamically typed messages, which consist of a set of named attributes of simple data types. The decoupling of event producers and subscribers improves scalability since it removes any explicit dependencies. Removing these dependencies reduces the coordination between readers and applications and facilitates the communication in this distributed environment, which is asynchronous by nature.

The Elvin subscription language supports [85] a regular expression matching operator and a number of simpler string matching routines, e.g. 'begins-with' and 'contains'. Arithmetic and logical operations are also available within the router. There is no support for user-defined operators and the scope of the subscription language is limited to a single notification.

Through its quenching functionality Elvin provides the subscription feedback mechanism we require to inform the readers about the subscriptions currently registered [112]. The Elvin router permits readers to define filters over the subscription database, constraining quench updates to those subscriptions of interest to the reader. The filter is expressed as a list of notification attribute names that must be present in the subscription. The readers can register, modify, or remove quench filters in a similar way to subscriptions.

In our implementation, the subscription feedback – provided by the quenching functionality of the Elvin router – allows us to carry out the filtering on the air interface, whenever possible, due to bandwidth considerations (cf. Figure 5.9). Filtering over the air interface is enabled by a number of RFID communication protocols, e.g. [32, 65], as demonstrated earlier. The quenching functionality also means that a reader whose data are of no interest will not operate and will not occupy any of the scarce bandwidth.

We previously outlined the need for the aggregation of the captured RFID data to reduce the flood of raw tag reads to more meaningful events such as the first appearance of a tag in the read range and its subsequent disappearance (cf. Figure 5.10). Aggregation is also needed to address the problem of temporary false negative reads and to smooth the data accordingly [17]. The RFIDStack supports the four different aggregation types listed in Table 5.2: Entry/Exit, Count, Passage, and Virtual Readers. The aggregation functionality is currently realized via surrogates to which the readers are connected (cf. Figure 5.9). In the future, more powerful readers can carry out this functionality themselves, while less powerful readers will continue to rely on a surrogate to carry out the aggregation. Aggregates that are based on data captured by more than a single reader, e.g. passage events, are computed by



Figure 5.10.: Entry and exit event illustration [48]. The row labeled A shows the frames in which the tag under test was present in the read range and should ideally have been detected (dark boxes). The row below shows the frames in which the tag was actually detected by the HF reader. Row C shows the assumed presence of the tag and the point of time where the entry and exit events are generated.

a separate aggregation entity within our architecture (cf. Figure 5.9). This is a result of the limited scope of the Elvin subscription language, which only operates on a single notification and thus cannot be utilized to specify aggregate events. The aggregation component in the RFID-Stack is triggered by the appropriate quench filter when an application registers for a particular type of aggregate. The aggregation component then subscribes to the messaging service to receive the appropriate notifications from the corresponding readers. These notifications are correlated and the new 'super' event notifications are again sent to the messaging router.

Elvin was not only chosen as an event router for our RFID middleware implementation because it has a built-in quenching mechanism and allows for content-based routing, but also because it is a pure notification service that only decouples subscribers and producers in space, but not in time. This avoids overheads commonly associated with messaging systems that provide strong reliability guarantees [36] and thus reduces the notification latency for those applications that require "realtime" updates. As notifications are not stored in the event router for applications which failed or are disconnected, guaranteed delivery has to be realized via dedicated modules that store events and replay them to the applications. In the RFIDStack, there is a buffering component that supports applications that require batch updates rather than notifications with a minimum latency. It collects raw tag reads or aggregates over the time period specified and delivers these in one batch via the messaging service at the desired point of time. The support for disconnected applications within the Elvin event router has also been extensively discussed in [119].

### 5.4.2. Writing to a Tag

The RFID middleware should ideally make writing to an RFID tag as easy as writing data to a hard disk of a computer. The virtual tag memory service (VTMS) in our system facilitates this by shielding the application from the particularities of RFID tag memory (cf. Figure 5.9): limited memory size, different memory organizations, reduced write range. Applications simply provide key-value pairs that should be written to a set of tags. The RFID middleware then checks with the VTMS for the appropriate tag memory block and page to write to the given key. If the write succeeds, the RFID middleware will acknowledge this to the application and will store a backup copy of the data in the virtual representation of the tag in the VTMS. If the tag memory gets corrupted at a later stage or the application wants to access the memory of the tag while the tag is outside the range of any reader, the RFID middleware can make the data available via this virtual memory. If the write to the tag fails due to insufficient power, the key-value pair will be stored in the VTMS and flagged as 'open'. The RFID middleware will retry the write command at a later point of time. If there is insufficient memory space, the application will receive the appropriate error message and the key-value will be stored in the virtual tag memory only. The application can also indicate that the virtual memory of a tag can only be accessed once the tag is in the read range of the particular reader.

To realize the 'write'-functionality the applications produce the corresponding event notifications and subscribe to receive an acknowledgement upon completion of the operation. The readers themselves subscribe to receive 'write command' notifications from the applications. The result is that applications and RFID readers are both event producers as well as event consumers.

### 5.4.3. Hardware Abstraction Layer

To address the idiosyncrasies of the different RFID readers available, a hardware abstraction layer was developed in joint work with M. Lampe [108]. This hardware abstraction layer runs on a surrogate that connects over a proprietary communication protocol to the reader hardware. Typical transport protocols include serial and TCP. The hardware abstraction layer currently supports a limited number of different reader platforms.

## 5.5. Evaluation

Tables 5.3 and 5.4 show that our implementation addresses the majority of the application requirements and constraints of passive RFID technology presented in the earlier sections of this chapter. The performance of the Elvin content message router is also adequate. A single instance of the above event router can process more than 10,000 messages per second [84]. In the worst case this corresponds to 100 readers continuously detecting 100 tags per second and feeding the raw data non-aggregated to the messaging system. In a more realistic scenario, a single non-federated event router should thus be able to deal with at least 1000 readers.

To support even larger numbers of event producers and notification volumes, Elvin provides a mechanism to farm out the subscription evaluation to multiple servers in a tightly-coupled local area federation. This includes methods for handover of connections to facilitate load sharing. However, the current version of the Elvin router is not yet suitable for wide-area deployment [112], where very large numbers of RFID readers, which are no longer deployed within the same organizational boundaries, are connected to a single content based routing service.

One reason why we chose the Elvin message router for our implementation was its quenching functionality. It provides us with the subscription feedback mechanism, which is essential for the efficient use of the limited bandwidth. While the readers do not need to maintain connections to the different applications, the quenching functionality provides a mechanism for the readers to determine whether the subscribers are interested in their notifications. To receive a quench update, the readers specify a quench filter over the subscription database, constraining quench updates to those subscriptions capable of matching notifications generated by the client. This quench filter is expressed as a list of attribute names that have to be present in the subscription. However, Elvin does not allow quench filters that specify both attribute name and the corresponding attribute value [112]. The result is that each

Application Re-	Addressed by
quirement	
Data dissemination	the use of a publish-subscribe messaging sys-
	tem that decouples readers and applications.
Data filtering	the use of a messaging system that allows for
	full content based routing. The subscription
	language of the messaging system allows for
	the filtering of all fields in a single notification
	and supports string comparison and logical
	operators.
Data aggregation	the use of a dedicated aggregation compo-
	nent that computes the aggregates that inter-
	pret read events from different readers. En-
	try/Exit and count aggregates are computed
	on a reader surrogate.
Writing to a tag	the provisioning of a write notification, which
	are produced by the applications and con-
	sumed by the readers.
Integration into IT	the use of simple heartbeat notifications
service management	which are produced by the readers to indicate
	their health to a system monitor subscribed
	to these heartbeat messages.
Privacy	the use of a subscription feedback mechanism
	(quenching) that prevents RFID readers col-
	lecting data without any application indicat-
	ing an interest in the data.

Table 5.3.: Addressing the application requirements. The table shows to what extent the RFIDStack fulfills the application requirements.

quench update contains a large proportion of the entire subscription database due to the notification structure of RFID events. The readers then have to filter the quench update locally to check whether there is a subscription that is relevant.

The Elvin subscription language also does not support compound filters that analyze more than a single notification. The result is that aggregates cannot be computed within the message router. In our implementation, aggregates are either computed by a separate aggregation component or directly by the readers or their surrogates.

The desirable integration of RFID readers into an existing IT-service management concept that performs incident, change, release, and configuration management is only possible to a very limited extent with our implementation. There is currently only a simple heartbeat event that a system monitor can subscribe to in order to monitor the health

RFID Constraint	Addressed By	
Limited communica-	filtering over the air interface enabled by the	
tion bandwidth	subscription feedback mechanism in the pub-	
	lish/subscribe messaging system	
Reliability issues	the computation of entry/exit aggregates	
	that reduce the volume of data generated to	
	events that reflect the status change of a tag	
	$({\rm appearance/disappearance}).$	
Tag memory varia-	the implementation of a virtual tag memory	
tions	system that abstracts from the different tag	
	memory organizations and provides redun-	
	dancy.	
Heterogeneous reader	the surrogate concept where (embedded)	
landscape	computing devices support those RFID read-	
	ers with limited computing resources	

Table 5.4.: Addressing the constraints of passive RFID technology. The table shows to what extent the RFIDStack addresses the constraints of passive RFID technology.

of a reader. The proper integration of the RFID readers into an existing IT-service management concept would require methods to query and modify the existing configuration of a reader, mechanisms to remotely update the software on a reader, and exception reporting functionality that is compatible with non-RFID devices. There is currently ongoing work within the RFID community to develop such a standardized reader management concept that builds on the established Simple Network Management Protocol [6].

### 5.6. Related Work

The need for an RFID middleware and the specific application requirements towards such an RFID middleware have been discussed in a number of publications [16, 48, 104]. Initially, the concept of a distributed networking infrastructure for RFID (cf. Figure 5.11) was proposed by the Auto-ID Center, an industry-sponsored research program to foster RFID adoption [105], which coined the term EPC Network.

Our work is closely related to the middleware component in the original EPC Network, called Savant [91] (cf. Figure 5.11). The Savant concept relied on a hierarchical system architecture with Edge Savants and Internal Savants. The Edge Savant is a Savant type that is connected to the readers and owes its name to its logical placement in



Figure 5.11.: Overview of EPC Network as originally envisioned by the Auto-ID Center [44].

the EPC Network. While the Savant software addresses many of the application requirements presented in Section 5.1, e.g. it features functionality for coping with the idiosyncrasies of different kinds of readers and for cleaning the data, there is only limited built-in functionality that specifically addresses the constraints of passive RFID technology. The event management system of the Edge Savant does not provide a feedback mechanism that allows the readers to adjust their queries to the subscriptions registered. This means that the scarce bandwidth might not be used efficiently: readers might operate although no application requests their data or the reader inventories the entire tag population when there is only a need to search for a particular kind of tag, e.g. a tag that identifies a shipment.

There are also some substantial design differences to the implementation we presented in this chapter. The Edge Savant provides both messaging and persistence in one server, while we separate message routing and message persistence. Central to our design is a pure notification service that does not provide any datastore functionality. Buffering and data storage are provided by dedicated components that subscribe to the notification service. This reduces the overhead associated with reliability guarantees and thus minimizes notification latencies. Short notification latencies are in particular important for applications that require "real-time updates", e.g. the pick and pack station mentioned earlier. In the Edge Savant implementation, there is also no subscription language with predefined operators. Instead, the Edge Savant concept allows applications to register event filters programmed in the Java programming language, which can operate on a combination of notifications. This approach increases the expressiveness of the subscription language at the expense of performance and scalability [21]. We believe that a subscription language with predefined operators is expressive enough to support RFID filtering and possibly aggregation and that there is no need for the additional flexibility provided by user-defined operators programmed in the Java programming language.

The RFID infrastructure envisioned by the Auto-ID Center also features a common vocabulary called Physical Markup Language [105]. The purpose of the core part of the physical markup-language (PML) Core) [46] is to provide a standardized format for the exchange of data captured by the sensors in an RFID infrastructure, for example, RFID readers (cf. Figure 5.11). PML Core comprises a set of XML schemas [124] that define the interchange format for the transmission of the data values captured. PML Core focuses on observables – physical properties and entities that are capable of being observed or measured by a sensor – rather than the observational and performance characteristics of the individual sensors or the interpretation of the observed values. ID Sensors, i.e. RFID readers, would thus report the detection of raw RFID reads. There is no functionality to represent aggregates, such as entry/exit events. Figure 5.12 shows a PML Core example. The PML Core language has been adopted by a number of commercial RFID implementations, such as SUN's Java System RFID Software Programming Platform [87] and SAP's Auto-ID Infrastructure [16]. While we could in principle have used at least part of the PML Core vocabulary as a message format in our content-based routing messaging system, the decision to use the Elvin content-based router prohibited this option, since Elvin uses dynamically typed messages comprising key-value pairs.

Today, the RFID middleware in the EPC Network is subject to a standardization effort under a set of new names within the EPCglobal community (cf. Figure 5.13). The Reader Protocol [6] and Application-Level-Events (ALE) [34] interfaces specify the functionality that was previously provided by the Savant framework and the PML Core Language in the original EPC Network architecture. Both of the interfaces provide functionality to disseminate, filter, and aggregate the captured

```
<?xml version="1.0" encoding="UTF-8"?>
< pmlcore: Sensor>
                <pmluid:ID>urn:epc:1:4.16.36</pmluid:ID>
                <pmlcore:Observation>
                                <pmluid:ID>00000001</pmluid:ID>
                                < pmlcore: DateTime>2002-11-06T13:04:34-06:00</pmlcore: DateTime>
                                < pmlcore: Command > READ_PALLET_TAGS_ONLY < / pmlcore: Command > reader and > re
                                <pmlcore:Tag>
                                                < pmluid: ID > urn: epc: 1:2.24.400 < /pmluid: ID >
                                <pmlcore:Tag>
                                                < pmluid: ID > urn: epc: 1:2.24.401 < /pmluid: ID >
                                <pmlcore:Tag>
                                                < pmluid: ID > urn: epc: 1:2.24.402 < /pmluid: ID >
                                </pmlcore:Tag>
                                <pmlcore:Tag>
                                                 <pmluid:ID>urn:epc:1:2.24.403</pmluid:ID>
                                <pmlcore:Tag>
                                                 <pmluid:ID>urn:epc:1:2.24.404</pmluid:ID>
                                </pmlcore:Sensor>
```

Figure 5.12.: *PML Core example.* This sample XML instance represents the event that an RFID reader has detected 5 different RFID tags. Namespace definitions are omitted.

data. The reader protocol implementations are deployed on reader devices directly or on a so-called concentrator, which plays a role similar to the surrogate in our RFID platform. RFID middleware acts as a reader protocol host and implements the application level event specification, which is exposed towards the applications. The two specifications address many of the application requirements we listed in Section 5.1. The interfaces support filtering, aggregation, and buffering (cf. Figure 5.14). There is also synchronous and asynchronous data communication with applications. However, the ALE specification currently provides no mechanism for reading other data than the tag identifier from the tag or a mechanism for writing to a tag. Our RFID middleware differs from the design proposed in these two specifications in a number of ways. In our implementation, there is one common vocabulary and subscription language. In the EPC Network design, ALE subscriptions need to be translated into the appropriate reader protocol queries. Notifications by the readers are received by



Figure 5.13.: EPC Network Architecture as defined by EPCglobal [6]. (Source: EPCglobal)

the ALE implementations and need to be converted into the appropriate ALE notifications. Having a common vocabulary and an event router that permits subscription feedback to the readers alleviates the need to implement this mapping functionality. The ALE and Reader Protocol also do not rely on a general purpose subscription language with predefined operators, but specify a set of queries and subscriptions that invoke filtering and aggregation. The reader protocol also places more complexity within the reader devices. The readers need to handle multiple connections and provide buffering complexity.

On the other hand, the reader protocol and especially the reader management protocol [6] provide significantly more features that fa-



Figure 5.14.: Reader Protocol Overview. The specification supports filtering (via tag selectors in the data acquisition subsystem), aggregation (via the event subsystem), and buffering (via the output subsystem) [6]. (Source: EPCglobal)

cilitate the integration of RFID devices into an IT system monitoring environment than our implementation. This includes device monitoring features, but also the possibility to configure external triggers, such as motion sensors or programmable logic controllers.

There are a number of other commercial and non-commercial RFID middleware products available, among others [16, 87, 92, 96]. All of them address the application requirements for device and data management. These implementations also provide substantially more configuration and system monitoring functionality than our implementation. To our knowledge, none of them use a design that features a general purpose content based router with subscription feedback and a subscription language that relies only on predefined operators, such as string comparators. There is also little dedicated support that addresses the constraints of passive RFID, e.g. the limited available bandwidth and the problems associated with writing to a tag.

The requirements which an overall RFID network infrastructure should meet, and the corresponding system architectures, have recently also been studied by [6, 16, 23, 103, 109, 121, 126]. Our work differs from the above because it exclusively focuses on the data and device management in an RFID system architecture. We consequently do not cover how RFID data can be interpreted in a given business context and turned into the corresponding application events. This work focuses strictly on the constraints imposed by passive RFID and how these can be addressed in an appropriate design.

In Section 5.2, we mentioned the limited bandwidth available to RFID readers. To address this problem, we proposed a middleware design that selectively propagates the subscriptions to the RFID readers. The implication is that readers which would capture data that no application is interested in at any given point of time will refrain from operating and occupying the scarce bandwidth. However, there are a number of complementary approaches to address this reader collision problem, which was first documented in [29]. They include approaches that involve frequency planning and time-division multiplexing schemes [19, 32, 59, 125], but also methods such as shielding and triggering RFID readers with motion sensors [82].

In [15], Bohn et al. propose the concept of tag deployment schemes where passive RFID tags are deployed in vast quantities and in a highly redundant fashion over large areas or object surfaces. The authors argue that such a super-redundant distribution of tags enables novel RFID-based services and applications, including a new means of cooperation between mobile physical entities. Our work relates to their work because we also foresee the impending ubiquity of RFID tags. However, the focus of our work is on dealing with the constraints of RFID and managing large deployments of networked readers, whereas Bohn et al. concentrate on tag distribution patterns and the novel services super-redundant tag distributions enable.

## 5.7. Summary

This chapter analyzes the requirements RFID middleware solutions should meet in order to manage large deployments of readers and the amount of data these readers capture. We argue that the characteristics of passive RFID technology introduce constraints that are unique to the development of middleware for the RFID domain. These constraints include the limited communication bandwidth available to RFID readers, the occurrence of false negative reads, tag memory variations, and the heterogeneous reader landscape. To address these constraints and the application requirements for filtered and aggregated RFID data, we propose an RFID middleware design that uses a publish/subscribe system featuring full content-based routing and a subscription feedback mechanism to the event producers. To reduce the notification latency for those applications which require "real-time" updates, we also make the case for a separation of notification and persistent storage service. This chapter also features a description and evaluation of the RFID-Stack, a middleware implementation that is based on the above design considerations. The RFIDStack uses the general purpose content-based router Elvin [112] which features a subscription feedback mechanism and a subscription language with predefined operators. We illustrate that this approach is well-suited to addressing application needs and technology constraints. However, we also discuss characteristics of the content-based router in its current version that limit the scalability of this approach. The RFIDStack also features a virtual tag memory system that abstracts from the different memory structures on RFID tags and facilitates the writing to a tag from an application perspective.

## 6. Conclusion

In this thesis, we show that the proliferation of RFID tags and readers introduces a number of technical challenges. We contend that appropriate infrastructure support is required to maintain adequate performance levels, to protect the privacy of the individual, and to facilitate the management of large scale reader deployments.

In particular, this thesis provides three contributions:

- two algorithms to optimize throughput over the radio channel and to speed up the identification of large tag populations using Bayesian transmission control strategies.
- a method to address the privacy concerns associated with RFID by integrating fair information practices into the communication protocols between readers and tags.
- a system design and implementation which decouples RFID readers and applications and addresses the application requirements, but also the constraints of passive RFID technology, such as the limited communication bandwidth available.

In this final chapter, we revisit the arguments that led us to the above techniques and systems and summarize the technical approaches and their benefits. We also discuss limitations of the approaches presented and outline future work.

## 6.1. Speeding Up the Identification of Large Tag Populations

In this section, we summarize our approach to increasing the identification speed of large tag populations. We also discuss limitations and future work.

### 6.1.1. Contribution

As the number of objects which are equipped with RFID tags increases, it is becoming increasingly important to identify large tag populations quickly. In this thesis, we provide two transmission strategies that speed up the identification of RFID tag populations by improving throughput over the shared radio channel. The throughput increase results from a more accurate estimate of the number of tags transmitting. Our transmission strategies address in particular the characteristics of the RFID domain. These include the frequent use of a variant of slotted ALOHA, known as framed ALOHA, in RFID protocols and the unknown tag arrival and departure rates. The two approaches presented in this thesis address these characteristics by recursively estimating the number of tags transmitting, based on feedback from the reader for respective slot outcomes. The two Bayesian strategies we developed differ in that the first strategy only updates the estimate of the number of the tags at the end of a frame, while the second strategy updates the estimate after each slot. The evaluation shows that, due to the unknown arrival and departure rates, the transmission strategy that incorporates the feedback from the reader on a slot-by-slot basis performs significantly better than the strategy that waits until the end of a complete frame before the estimate is updated. This is due to the fact that our slot-by-slot scheme will cancel a frame early if the frame size is estimated to be non-optimal. The results of our experiments and simulations also show that both transmission strategies have a higher throughput than existing approaches. The simulations rely on an RFID simulation engine we developed that supports different pathloss, fading, capture, and tag mobility models.

### 6.1.2. Limitations and Future Work

The increased throughput that can be achieved with our transmission schemes comes at the expense of a significant amount of computations. Alternative transmission schemes make certain assumptions about the distribution of the number of tags present or simply assume a fixed multiplicity of conflict. This significantly reduces the resources required to estimate the number of tags transmitting and to choose the frame size accordingly. While our Bayesian approaches permit some computations to be made a-priori, there is a significant amount of computation remaining that needs to be carried out online. Our approaches have also not been validated experimentally with a UHF RFID system. While we did use an HF RFID system to evaluate our transmission schemes, we relied on simulations at UHF. Future work should thus aim to implement the above approaches in a UHF RFID reader, e.g. in a dock door scenario. The simulation engine could also be upgraded to include tag antenna directivity and multiple reader antennas.

In our evaluations, we assumed that a reader can operate independently in a given channel. In practice, there will be other readers operating in the close vicinity, which will possibly interfere. Future work might thus also consider the effect of such reader collisions on the performance of the transmission schemes. Furthermore, future investigations could include a transmission scheme that chooses an appropriate frame size for a number of readers which are synchronized in order to deal with the limitations of the listen-before-talk schemes introduced in some countries.

Lastly, our Bayesian models assume a feedback model where the reader can successfully distinguish between no, a single, a single but corrupted tag reply, and more than a single tag reply. In practice, it might be difficult to distinguish a corrupted single tag reply from a collision where more than a single tag replied. Future versions of our transmission schemes should thus include the possibility and ideally the likelihood of such a wrong classification.

## 6.2. Addressing RFID Privacy Concerns

In this section, we outline our method of making RFID communication more transparent to the individual. We discuss our contribution, list limitations of our approach, and present future work.

## 6.2.1. Contribution

While the invisible nature of RFID technology has many benefits from an automation perspective, it is also the cause of privacy concerns. The intended deployment of RFID tags on everyday items epitomizes for some the dangers of an Orwellian future: unnoticed by consumers, embedded microchips in our personal devices, clothes, and groceries can covertly be triggered to reply with their ID and other information, potentially allowing for a fine-grained yet invisible surveillance mechanism that pervades large parts of our lives. In this thesis, we present an approach to address this problem that integrates a subset of the widely accepted fair information principles into the communication protocols between RFID readers and tags. We contend that having RFID readers that explicitly declare the scope and purpose of the tag data collection, as well as disclosing the identity of their operators, will allow both consumers and regulators to better assess and control the impact of everyday RFID encounters. Our analysis shows how the fair information principles of collection limitation, purpose specification, use limitation, openness, and accountability can be incorporated in today's RFID communication protocols without significant performance penalties. We also present the prototype of a watchdog tag that allows consumer interest groups and privacy-concerned individuals to judge whether a particular RFID reader deployment complies with the corresponding regulations by displaying and logging the information regarding the data collection broadcast over the radio channel.

### 6.2.2. Limitations and Future Work

Even with our proposed protocol extensions, unauthorized read attempts by readers not conforming to our specification will still be possible. While consumers carrying a watchdog tag might be able to actively jam the tag-to-reader communication, the average consumer would still need to resort to explicitly disabling her tags in order to completely prevent misuse.

In its current form, the prototype implementation of the watchdog tag also relies on a wireless LAN connection to transmit the reader policy. Future versions of such a watchdog tag prototype could use a semipassive tag that decodes the privacy policy information transmitted by the reader over the RFID communication channel. The incorporation of the fair information practices in the RFID communication protocol could then be realized not by extending the actual command set of the protocol, but by embedding the reader policy ID and purpose of data collection in existing commands which carry user-defined data, e.g. a "write" command.

Future RFID tags could also implement the resurrecting duckling model proposed by Stajano [116], where tags would only respond to a "mother" reader, but ignore requests from all others. Instead of killing tags at checkout, stores would transfer their "mother" rights to the customer's reader, thus allowing for safe post-sales RFID usage. Additionally, such "mother" readers could inhibit replies from "its" tags for non-desired purposes and intentions to unknown readers by programming the tags accordingly.

While technical approaches to addressing the threats to privacy resulting from the proliferation of RFID systems, such as the one presented in this thesis, can help to protect the rights of the individual, their contribution to changing the current negative consumer attitude towards RFID should not be overestimated. In order to positively influence risk perception, education about the true capabilities of the technology and the development of a trust relationship between data collectors and consumers are also required. In the long term, the availability of suitable applications that let the consumer benefit from the proliferation of RFID technology might play the most important role in addressing consumer fears.

## 6.3. Managing RFID Systems

In this section, we provide a summary of our middleware design and implementation. We also discuss limitations and future work.

### 6.3.1. Contribution

The proliferation of RFID readers and tags also introduces the need for RFID middleware solutions designed to manage large deployments and the data captured. In this thesis, we argue that the characteristics of passive RFID technology introduce constraints that are unique to middleware for the RFID domain. These constraints include the limited communication bandwidth available to RFID readers, the occurrence of false negative reads, tag memory variations, and the heterogeneous reader landscape. To address these constraints and the application requirements for filtered and aggregated RFID data, we propose an RFID middleware design that uses a publish/subscribe system, which features full content-based routing and a subscription feedback mechanism to the event producers. To reduce the notification latency for those applications which require "real-time" updates, we also make the case for a separation of notification and persistent storage services. This thesis also features a description and evaluation of the RFIDStack, a middleware implementation that is based on the above design considerations. The RFIDStack uses the general purpose content-based router Elvin [112] which features a subscription feedback mechanism and a subscription language with predefined operators. We illustrate that this approach is well-suited to addressing the application's needs, e.g. to receive filtered and aggregated RFID data, and the constraints of the technology, such as the limited communication bandwidth. The RFID-Stack also features a virtual tag memory system that abstracts from the different memory structures on RFID tags and facilitates writing to a tag from an application perspective.

### 6.3.2. Limitations and Future Work

One reason why we chose the Elvin message router [112] for our implementation was its quenching functionality. It provides us with the subscription feedback mechanism, which is essential for the efficient use of the limited communication bandwidth. While the readers do not need to maintain connections to the different applications, the quenching functionality provides a mechanism for the readers to determine whether the subscribers are interested in their notifications. Such a quench filter is expressed as a list of attribute names that have to be present in the subscription. However, Elvin does not allow quench filters that specify both attribute name and the corresponding attribute value. The result is that each quench update contains a large proportion of the entire subscription database, due to the notification structure of RFID events and to the fact that the readers have to filter the quench update locally to check whether there is a subscription that is relevant.

In addition, the Elvin subscription language does not support compound filters that analyze more than a single notification. The result is that aggregates cannot be computed within the message router. In our implementation, aggregates are either computed by a separate aggregation component or directly by the readers or their surrogates.

The desirable integration of RFID readers into an existing IT-service management concept that performs incident, change, release, and configuration management is only possible to a very limited extent with our implementation. There is currently only a simple heartbeat event that a system monitor can subscribe to in order to monitor the health of a reader.

The messaging system is also not suitable for wide-area deployments, possibly on an Internet scale. While it is sufficient in many scenarios to make the RFID data available only to a single organization and to operate within a local area network, there are scenarios that would benefit from wide-area deployments. Future work might thus investigate the possibility of combining the features of Elvin which are desirable for an RFID middleware implementation, such as the quenching functionality, with the features of other content-based routers that specifically address wide area scalability. Appendices

## A. RFID Simulator

This Appendix describes the components of the RFID Simulation engine and its configuration options.

### A.1. Components

In the following sections, we present the different components (or *entities* in the terminology of JiST) of the simulator and show their functionality and purpose inside the simulator. Table A.1 shows an overview of the components with their concrete Java class names and the interfaces they implement.

All entities are Java classes that extend an interface that specifies which methods are exposed to other entities for simulation time invocation. In other words, the interface specifies the simulation events that an entity is able to receive. However, the implementation can contain other public methods that can be invoked with standard Java invocation semantics through a "normal" object reference.

### A.1.1. Physical Layer

### **Class Field**

Field represents the simulated radio field and is responsible for modeling signal propagation among reader and tag radios as well as the mobility of nodes. The Field entity has references to all radios in the field and all messages sent between a reader and tags pass through Field, which decides based on location information, radio node information, and a signal attenuation model (path loss and fading), which of the radio nodes will receive a given message. Field manages the radio it contains in a hierarchical spatial data structure for efficient signal propagation calculations.

Field delivers messages (implementations of interface Message) by making up-calls to a receive() method of all radio entities that should

Class	Entity Interface	Description
Field	FieldInterface	Radio field entity
RFIDReaderRadio	RFIDRadioInterface	Reader radio entity
RFIDReaderMac	RFIDMacInterface	Reader MAC entity
RFIDReader	RFIDReaderInterface	Reader Logic entity
RFIDTagRadio	RFIDRadioInterface	Tag radio entity
RFIDTagMac	RFIDMacInterface	Tag MAC entity
RFIDTag	RFIDTagInterface	Tag logic entity
RFIDApplication	RFIDApplicationInterface	Application entity
RFIDApplicationObserver	RFIDApplicationObserverInterface	Observer entity
Inventory	RFIDApplicationInterface	Inventory entity

Table A.1.: The main components with corresponding Java classes and interfaces. Package names are omitted for clarity.

receive the message. It also receives down-calls from radio entities (method transmit()) that wish to transmit a message.

Supported fading models include none (Fading.None), Raleigh fading (Fading.Raleigh), and Rician fading (Fading.Rician). Supported path loss models include free-space (PathLoss.FreeSpace) and two-ray (PathLoss.TwoRay) path loss.

The Field entity was adopted directly from SWANS without any changes, since it delivers all the functionality needed for the radio field used in a UHF RFID protocol.

### A.1.2. Reader

### Class RFIDReaderRadio

RFIDReaderRadio implements the interface RFIDRadioInterface, which extends the interface RadioInterface provided by SWANS. It represents the physical layer of the reader, i.e. the antenna, and is responsible for transmitting and receiving messages from the Field entity, detecting bit-errors, managing captures (if enabled) and delivering successfully received messages to the upper MAC entity.

Since the RFID reader radio should be able to consider the capture effect, and the original radio entity RadioNoiseAdditive from SWANS

does not support this, we had to rewrite a great deal of the functionality in the **receive()** method. Other parts only needed minor adjustments or none at all.

Each RFIDReaderRadio has a set of radio properties which are used by the Field entity for signal propagation. These properties include radio mode (idle, sensing, receiving, transmitting, sleeping), transmit power, gain, data rate, sensitivity threshold, receive threshold, signal wave length and background noise level.

The reader radio supports four capture models: no capture, a *sto-chastic model* which receives the strongest of multiple signals with a certain probability, a *simple power model* which receives the strongest signal if its receiving power is greater than a certain factor of each other signal, and an *advanced power model* which receives the strongest signal if its power is greater than a certain factor of the sum of all other signals.

We implemented three bit-error models for receiving messages in tag and reader radios: no bit-errors, bit-errors occur independently and randomly with a fixed probability (denoted bit-error rate), and biterrors occur independently and randomly with a probability that depends on the signal-to-noise ratio of the receiving signal.

### Class RFIDReaderMac

The RFIDReaderMac entity represents the link or MAC layer of the simulated RFID reader. It is responsible for checking the CRC of incoming messages, determining if it is a recognized message by the reader (one that the reader can handle), and sending the upper logic layer an appropriate event containing the message content if that is the case. When the upper layer requests a message to be sent, the RFIDReaderMac entity takes care of correct link timings (i.e. sufficient wait times between messages and handling of time-outs when waiting for replies), checks that the radio entity is idle and calculates the transmission time of the message plus preamble based on the bit string of the message and the current symbol length of a data-1 and data-0.

RFIDReaderMac holds an entity reference to the RFIDReaderRadio and RFIDReader entities for sending and receiving simulation events.

### **Class RFIDReader**

Finally, the RFIDReader entity contains the logic of the simulated reader and its main purpose is to provide the interface to the reader's functionality to applications that want to use the reader. It provides methods for every reader command in the protocol specifications, e.g. sendQuery() for sending a *Query* command or sendACK() for sending an *ACK* command. When the RFID application wants to send a command, it calls the corresponding method of the RFIDReader entity, which then constructs an appropriate instance of Message with the specified parameters and sends it to the RFIDReaderMac entity.

Simulation applications communicate only with the RFIDReader entity, i.e. make method calls to the provided methods and receive upcalls from the RFIDReader entity on specific simulation events, e.g. the arrival of a new message. For an application, RFIDReader is the only access point to the simulation.

RFIDReader receives messages from the underlying RFIDReaderMac entity by calls to the various handleXXX() methods , but since the logic for handling incoming messages is implemented in a subclass of RFIDApplication implemented by the simulation programmer in the application layer, the RFIDReader entity does not handle any incoming messages directly, but rather unpacks the message contents and passes them up to the RFIDApplication entity it holds a reference to by sending it a corresponding simulation event.

## A.1.3. Tag

### Class RFIDTagRadio

The RFIDTagRadio entity, which, like RFIDReaderRadio, implements the interface RFIDRadioInterface, represents the physical layer of a tag and its functionality is roughly the same as the RFIDReaderRadio entity. Theoretically, RFIDTagRadio can support capture models, but since RFID tags do not need to account for capture (they talk to only one reader), the simulator disables capture for tag radios.

## Class RFIDTagMac

The RFIDTagMac entity is analogous to the RFIDReaderMac entity in the reader, and since they both share the same superclass RFIDMac, they also share most of the functionality. RFIDTagMac is responsible for checking the CRC of incoming reader commands, checking if a valid command was received and passing the command to the RFIDTag entity. When the tag wants to send a reply, RFIDTagMac makes sure that specified link timing intervals are maintained and sends the message to the RFIDTagRadio entity for transmitting over the air interface.

### Class RFIDTag

At the top of a simulated tag lies the RFIDTag entity which encapsulates the tag logic. The purpose of this entity is to implement all the Class1Gen2 protocol functionality in the tags. The entity keeps track of the different protocol flags, saves generated random numbers, randomly generates the slot counter on demand, has access to the tag memory banks and maintains the tag's state.

The tag protocol logic, i.e. the protocol state machine, is actually not implemented directly in RFIDTag, but is rather encapsulated in seven different subclasses of TagState, which are implemented as Singleton classes, each corresponding to a state the tag can be in. So the RFIDTag holds a reference to the TagState object that represents the current state, and all messages it receives are passed to this TagState object for state-sensitive handling of the messages. With this object-oriented design, it was straightforward to implement the rather complex state machine of the protocol, and if future changes to the state machine are required, they are very easy to incorporate.

## A.1.4. Application Layer

### Abstract Class RFIDApplication

RFIDApplication is the abstract superclass of all RFID application entities. It implements some functionality common to all applications, such as providing the reference to the RFIDReader entity and managing a set of application observers and notifying them if the application has finished. Additionally, it provides empty implementations of all methods in the RFIDApplicationInterface, so that applications can just overwrite the methods they are interested in and do not have to implement the complete interface.

**RFIDApplicationInterface** defines the different simulation events that an application can receive, i.e. handling of different tag replies, handling of no replies, handling of tag collisions and handling of transmission errors. It also defines methods to start, stop, and reset an



Figure A.1.: The components of the inventory application.

application, which are called from the simulator at specific points in time.

RFIDApplication is a subclass of RFIDApplicationObserver, any application can therefore act as an observer of another application.

### **Class Inventory**

Inventory is a concrete subclass of RFIDApplication and the only full-featured RFID application implemented in the scope of this thesis. The application just tries to identify all tags in the range of the reader and prints out a list of the tag EPCs when it is done, together with statistics on the inventory procedure, e.g. how many tags were inventoried or the achieved throughput.

Inventory adjusts the frame size dynamically with the help of a slotcount algorithm, which can be chosen in the simulator property file. The slot-count algorithm calculates the frame size based on continuous feedback from the reader regarding the occupancy of the slots during the inventory procedure.

Inventory may optionally use an Estimator application, which is used to estimate how many tags are present at the beginning of the inventory procedure and initialize the slot-count algorithm with the estimated value. Other options can be configured for Inventory in the simulator property file.

Figure A.1 shows the components of the inventory application.

### Abstract Class RFIDApplicationObserver

The abstract RFIDApplicationObserver entity is kind of a metaapplication that can observe other applications and receive events from
them if specific events occur. At the moment, the functionality is quite limited and a RFIDApplicationObserver is only notified when an observed application has finished running. An application entity can pass an ApplicationResult object to the observer, thus allowing some kind of return value from the application.

A concrete subclass of RFIDApplicationObserver is AppRepeater, which just runs an application multiple times successively, aggregates the application results and finally calculates the average of all results.

#### Interface SlotCountAlgorithm

This is not a separate entity but rather an object inside the Inventory entity (Fig. A.1). Its purpose is to calculate a frame size based on feedback from the Inventory application. After each slot during the inventory procedure it receives feedback whether no, exactly one, or more than one tag replied in the slot, or if a capture or error occurred.

Thanks to a plug-in architecture, new slot-count algorithms can easily be implemented and plugged into the **Inventory** entity. For this thesis, five concrete implementations of the interface **SlotCountAlgorithm** were implemented, evaluated and compared to each other.

# A.2. Simulator Properties

This is a reference of all available simulator properties that can be specified in the simulator properties file.

All property values that represent units of time (denoted below as *Time*) can be appended with a time unit following an underscore. Valid time units are **ns** (nanoseconds), **us** (microseconds), **ms** (milliseconds), **s** (seconds), **m** (minutes), **h** (hours), **d** (days), **tari** or **t** (value of tari), and **rtcal** (value of rtCal). If no unit is specified, the value is interpreted as simulation time units.

All property values that represent units of power (denoted below as *Power*) can be appended with a power unit following an underscore. Valid power units are dBm (decibels) or mW (milliwatts).

*Boolean* properties can take one value of true, yes, on, or 1 for *true* and any other value for *false*.

*Coordinate* properties represent a coordinate or a dimension in 2Dspace as a pair of real values separated by a comma, e.g. 4.5,5.0.

### A.2.1. General Simulator Properties

These are general properties of the simulator and its components. They can be freely adjusted

- sim.number\_of\_tags Integer Number of tags to simulate.
- sim.time\_unit *Time* Time unit that is used in logging output.
- sim.duration *Time* Maximum length of simulation.
- sim.bounds *Coordinate* Dimension of the simulation.
- sim.placement Class name Placement model for tag placement as full class name representing the model. Currently supported are rfidsim.util.PlacementRandom (random placement), rfidsim.util.PlacementGrid2D (placement on 2D grid), rfidsim.util.Placeme (manual placement read from file), rfidsim.util.PlacementCircular (circular placement).
- field.propagation\_delay *Time* The signal propagation delay in the radio field. Default is 60*ns*.
- field.fading\_model *Class name* The fading model to use as full class name. Supported are jist.swans.field.Fading\$None (no fading), jist.swans.field.Fading\$Rician (Rician fading), and jist.swans.field.Fading\$Rayleigh (Rayleigh fading).
- field.pathloss\_model Class name The path loss model to use as full class name. Supported are jist.swans.field.PathLoss\$FreeSpace (free-space path loss), jist.swans.field.PathLoss\$TwoRay (tworay path loss), and rfidsim.util.RFIDPathLoss (free-space path loss with higher exponent).
- field.propagation\_limit Power The minimum power of a signal that is still propagated.
- reader.name *String* The name of the reader.
- reader.location *Coordinate* The location of the reader in the simulation. Should be inside the bounds specified with sim.bounds. If not specified, the reader is placed with the chosen placement model.

- reader.radio.gain *Real* The gain of the reader antenna in decibels.
- reader.radio.threshold *Power* The receiving threshold of the reader antenna. Signals below this threshold can be detected, but not received.
- reader.radio.sensitivity Power The sensitivity threshold of the reader antenna. Signals below this threshold can no longer be detected.
- reader.radio.transmit\_power Power The transmitting power of the reader antenna.
- reader.radio.antenna\_radiation pattern *File path* The radiation pattern of the reader antenna.
- reader.radio.antenna\_direction *Real* The direction of the reader antenna.
- reader.delay *Time* The delay introduced by the logic layer of the reader.
- reader.mac.delay *Time* The delay introduced by the MAC layer of the reader.
- reader.biterror\_table *File path* The file to read the biterror table from. The table specifies bit-error values for different signal-to-noise ratios.
- reader.bit\_error\_rate *Real* The fixed bit-error rate the reader should apply to incoming messages.
- reader.powersignal\_interval *Time* The time interval at which the reader sends a virtual power beacon to the tags when the radio field is powered.
- reader.capture\_model Integer The capture model the reader should use: -1 for no capture, 0 for a simple stochastic model, 1 for a independent power-based model, or 2 for a additive power-based model.
- reader.capture\_threshold\_ratio *Real* The minimum power ratio for a successful capture for the power-based capture model.

- reader.capture\_percentage *Real* The probability that a collided slot can be resolved in the stochastic capture model.
- tag.epc Class name The type of EPC tags should use as a full class name. Implemented are rfidsim.epc.GID96 (EPCglobal 96bit General Identifier), rfidsim.epc.SimpleEPC16 (16bit integer ID), rfidsim.epc.SimpleEPC32 (32bit integer ID), rfidsim.epc.SimpleEP (64bit integer ID).
- tag.radio.transmit\_power Power The transmitting power of the tag antenna.
- tag.radio.gain *Real* The gain of the tag antenna in decibels.
- tag.radio.threshold *Power* The receiving threshold of the tag antenna. Signals below this threshold can be detected, but not received.
- tag.radio.sensitivity Power The sensitivity threshold of the tag antenna. Signals below this threshold can no longer be detected.
- tag.delay *Time* The delay introduced by the logic layer of the tag.
- tag.mac\_delay *Time* The delay introduced by the MAC layer of the tag.
- tag.data\_rate Integer The data rate of the tags in bits per second.
- tag.mobility.model *Class name* The mobility model of the tags as a full class name. Only jist.swans.field.Mobility\$Static and jist.swans.field\$ContinousTranslation are supported at the moment.
- tag.mobility.translation\_delta\_x *Real* The relative tag translation in the x-direction.
- tag.mobility.translation\_delta\_y *Real* The relative tag translation in the y-direction.
- tag.mobility.translation\_interval Integer The interval during which the tags move the above distance.

- tag.biterror\_table *File path* The file to read the biterror table from. The table specifies bit-error values for different signal-to-noise ratios.
- tag.bit\_error\_rate *Real* The fixed bit-error rate the tags should apply to incoming messages.
- tag.memory.read\_access\_time *Time* The read access time for one memory read access.
- tag.memory.write\_access\_time *Time* The write access time for one memory write access.
- tag.memory.bandwidth *Integer* The bandwidth or data rate of tag memory when reading data.
- tag.radio.backscatter\_ratio *Real* The fraction of power of an incoming signal a tag uses for backscattering its reply.
- env.temperature *Real* The temperature of the environment in Kelvin.
- env.temperature\_factor *Real* The temperature factor used by the calculation of thermal noise.
- env.ambient\_noise Power The background noise that is present in the environment.

# A.2.2. Protocol Properties

These are properties specified by the protocol specifications. Generally they should not be changed unless you know what you are doing, i.e. you know the protocol specifications and are aware of the consequences of changing these properties.

- reader.tari Time Value of Tari, the reference time interval for a data-0 symbol period in reader-to-tag signaling. Preferred Tari values are  $6.25\mu s$ ,  $12.5\mu s$ , and  $25\mu s$ .
- reader.data\_1\_symbol\_length *Time* Length of a data-1 symbol period. Should be specified with unit tari and must lie between 1.5*Tari* and 2.0*Tari*.

- reader.rt\_cal *Time* The length of the reader-to-tag calibration symbol. Is set to the length of a data-0 symbol plus the length of a data-1 symbol automatically and should not be changed.
- reader.tr\_cal *Time* The length of the tag-to-reader calibration symbol. Should be specified with unit rtcal and must lie between 1.1*RTcal* and 3.0*RTcal*.
- reader.no\_tag\_waittime *Time* The time the reader waits when no tags replied. Should be specified with unit rtcal.
- reader.time\_between\_commands *Time* The time the reader waits between successive reader commands. Should be specified with unit rtcal.
- reader.carrier\_frequency Real The operation frequency of the reader in Hz. Set to 900Mhz per default.

# Bibliography

- Resolution on Radio Frequency Identification. 25th International Conference of Data Protection and Privacy Commissioners, November 2003. Available from: www.privacyconference2003. org/commissioners.asp.
- [2] Norman Abramson. The ALOHA System-Another alternative for computer communications. In Proc. 1970 Fall Joint Computer Conference, volume 37, pages 281–285, 1970.
- [3] AEW&C WINGS. IFF Questions and Answers. Available from: www.dean-boys.com/extras/iff/iffqa.html.
- [4] Rakesh Agrawal, Jerry Kiernan, Ramakrishnan Srikant, and Yirong Xu. Implementing P3P Using Database Technology. In Proceedings of the IEEE 19th International Conference on Data Engineering, pages 595–606, Bangalor, India, March 2003. Computer Society, IEEE Press.
- [5] K. Alexander, T. Gilliam, K. Gramling, C. Grubelic, H. Kleinberger, S. Leng, D. Moogimane, and C. Sheedy. Applying Auto-ID to Reduce Losses Associated with Shrink. Technical Report IBM-AUTOID-BC-003, Auto-ID Center, November 2002.
- [6] Architecture Review Committee. The EPCglobal Architecture Framework. EPCglobal, July 2005. Available from: www. epcglobalinc.org.
- [7] David Arnold, Bill Segall, Julian Boot, Andy Bond, Melfyn Lloyd, and Simon Kaplan. Discourse with disposable computers: How and why you will talk to your tomatoes. In *Proceedings of the Usenix Workshop on Embedded Systems*, pages 9–22, Cambridge, MA, AZ, Mar 1999.
- [8] Auto-ID Center. Draft protocol specification for a 900 MHz Class
  0 Radio Frequency Identification Tag, May 2003. Available from:
  www.epcglobalinc.org.

- [9] Auto-ID Center. The New Network, Feb 2006. Available from: http://archive.epcglobalinc.org/new\_media/ brochures/ENGLISH\_AUTO\_ID\_CENTRE.pdf.
- [10] Jean Bacon, Ken Moody, John Bates, Richard Hayton, Chaoying Ma, Andrew McNeil, Oliver Seidel, and Mark Spiteri. Generic Support for Distributed Applications. *IEEE Computer*, 68– 76(3):33, March 2000.
- [11] Ribon Barr. An efficient, unifying approach to simulation using virtual machines. PhD thesis, Cornell University, May 2004.
- [12] Rimon Barr. SWANS Scalable Wireless Ad hoc Network Simulator, User Guide, 2004. Available from: http://jist.ece. cornell.edu/docs.
- [13] Rimon Barr, Zygmunt J. Haas, and Robbert van Renesse. JiST: An efficient approach to simulation using virtual machines. Software - Practise and Experience, 35(6):539–576, May 2005.
- [14] Dimitri Bertsekas and Robert Gallager. Data Networks. Prentice-Hall, Inc., 1987.
- [15] Jürgen Bohn and Friedemann Mattern. Super-Distributed RFID Tag Infrastructures. In Proceedings of the 2nd European Symposium on Ambient Intelligence (EAZI 2004), number 3295 in Lecture Notes in Computer Science (LNCS), pages 1–12, Eindhoven, The Netherlands, November 2004. Springer-Verlag.
- [16] C. Bornhövd, T. Lin, S. Haller, and J. Schaper. Integrating Automatic Data Acquisition with Business Processes - Experiences with SAP's Auto-ID Infrastructure. In *Proceedings of the 30st* international conference on very large data bases (VLDB), pages 1182–1188, Toronto, Canada, 2004. VLDB Endowment.
- [17] J. Brusey, C. Floerkemeier, M. Harrison, and M. Fletcher. Reasoning about Uncertainty in Location Identification with RFID. In Workshop on Reasoning with Uncertainty in Robotics at IJCAI-2003, Acapulco, Mexico, 2003.
- [18] John I. Capetanakis. Tree Algorithms for Packet Broadcast Channels. *IEEE Transactions on Information Theory*, IT-25(5):505– 515, September 1979.

- [19] Bogdan Carbunar, Murali Krishna Ramanathan, Mehmet Koyuturk, Christoph Hoffmann, and Ananth Grama. Redundant Reader Elimination in RFID Systems. In Second Annual IEEE Communications Society Conference on Sensor and Ad Hoc Communications and Networks, IEEE SECON 2005, pages 176–184, Santa Clara, CA, USA, Sep 2005.
- [20] Antonio Carzaniga. Architectures for an Event Notification Service Scalable to Wide-area Networks. PhD thesis, Politecnico di Milano, Milano, Italy, December 1998.
- [21] Antonio Carzaniga, David S. Rosenblum, and Alexander L. Wolf. Achieving scalability and expressiveness in an internet-scale event notification service. In *Proceedings of the Nineteenth Annual* ACM Symposium on Principles of Distributed Computing, pages 219–227, Portland, Oregon, July 2000.
- [22] G. Chappell, D. Durdan, G. Gilbert, L. Ginsberg, J. Smith, and J. Tobolski. Auto-ID on Delivery: The Value of Auto-ID Technology in the Retail Supply Chain. Technical Report ACN-AUTOID-BC-004, Auto-ID Center, June 2002.
- [23] S. Chawathe, V. Krishnamurthyy, S. Ramachandrany, and S. Sarma. Managing RFID Data. In Proceedings of the 30st international conference on very large data bases (VLDB), pages 1189–1195, Toronto, Canada, 2004. VLDB Endowment.
- [24] Peter H. Cole, Kamran Eshraghian, and Ashim K. Roy. US Patent Nr. 4,364,043 – Efficient Object Identification, Sep 1981.
- [25] Lorrie Cranor, Marc Langheinrich, Massimo Marchiori, and Joseph Reagle. The Platform for Privacy Preferences 1.0 (P3P1.0) Specification. W3C Candidate Recommendation, December 2000. Available from: www.w3.org/TR/P3P/.
- [26] D.E.N. Davie, M.J. Withers, and R.P. Claydon. Passive Coded Transponder Using an Acoustic-Surface-Wave Delay line. *Elec*tronics Letters, 11(8):163–164, April 1975.
- [27] Economist. The Best Thing since the Bar-Code, Feb 2003.
- [28] EM Microelectronic-Marin SA. Multi Frequency Contactless Identification Device (EM 4022), 2002. Avail-

able from: www.emmicroelectronic.com/webfiles/Product/ RFID/ds/EM4022\_DS.pdf.

- [29] Dan Engels. The Reader Collision Problem. Technical Report MIT-AUTOID-WH-007, Auto-ID Center, 2001.
- [30] EPCglobal. EPC Tag Data Specification 1.1, November 2003. Available from: www.epcglobalinc.org.
- [31] EPCglobal. EPC Tag Data Specification 1.1, November 2003.
- [32] EPCglobal. Class 1 Generation 2 UHF Air Interface Protocol Standard Version 1.0.9, 2005. Available from: www. epcglobalinc.org.
- [33] EPCglobal. Object Naming Service (ONS) Version 1.0. www.epcglobalinc.org, october 2005.
- [34] EPCglobal. The Application Level Events (ALE) Specification, Version 1.0, Sep 2005. Available from: www.epcglobalinc.org.
- [35] ERM TG34. Electromagnetic compatibility and Radio spectrum Matters (ERM);Radio Frequency Identification Equipment operating in the band 865 MHz to 868 MHz with power levels up to 2 W. Technical Report EN 302 208-1 V1.1.1, European Telecommunications Standards Institute (ETSI), 2004. Available from: www.etsi.org.
- [36] Patrick Th. Eugster, Pascal A. Felber, Rachid Guerraoui, and Anne-Marie Kermarrec. The Many Faces of Publish/Subscribe. ACM Computing Surveys, 35(2):114–131, June 2003.
- [37] European Commission. Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data, November 1995. Available from: http://europa.eu.int/eur-lex/en/lif/dat/ 1995/en\_395L0046.html.
- [38] European Radiocommunications Office (ERO). Relating to the use of Short Range Devices (SRD). Technical Report ERC/REC 70-03 Version of 17 November 2005, ERO, Nov 2005. Available from: www.ero.dk.

- [39] George Falcke. The New RFID Standard in Europe. Available from: www.rfidjournal.com.
- [40] Klaus Finkenzeller. *RFID Handbook: Radio-Frequency Identifi*cation Fundamentals and Applications. John Wiley & Sons, 2000.
- [41] E. Fleisch and M. Dierkes. Ubiquitous Computing: Why Auto-ID is the Logical Next Step in Enterprise Automation. Technical Report STG-AUTOID-WH-004, Auto-ID Center, 2003.
- [42] E. Fleisch and C. Tellkamp. Inventory Inaccuracy and Supply Chain Performance: A Simulation Study of a Retail Supply Chain. Int. J. of Production Economics, 95(3):373–385, 2005.
- [43] Richard Ribon Fletcher. Low-Cost Electromagnetic Tagging: Design and Implementation. PhD thesis, Massachusetts Institute of Technology, September 2002.
- [44] Christian Floerkemeier. EPC-Technologie vom Auto-ID Center zu EPCglobal. In Elgar Fleisch and Friedemann Mattern, editors, Das Internet der Dinge – Ubiquitous Computing und RFID in der Praxis. Springer-Verlag, 2005.
- [45] Christian Floerkemeier. Transmission control scheme for fast RFID object identification. In Proceedings of the Pervasive Wireless Networking Workshop at IEEE PERCOM 2006, Pisa, Italy, 2006.
- [46] Christian Floerkemeier, Dipan Anarkat, Mark Harrison, and Ted Osinski. PML Core Specification 1.0. Technical Report STG-AUTOID-WH005, Auto-ID Center, Sep 2003.
- [47] Christian Floerkemeier and Matthias Lampe. Issues with RFID Usage in ubiquitous computing applications. In Alois Ferscha and Friedemann Mattern, editors, Second International Conference, PERVASIVE 2004, volume 3001 of Lecture Notes in Computer Science (LNCS), pages 188–193, Linz/Vienna, Austria, April 2004. Springer-Verlag.
- [48] Christian Floerkemeier and Matthias Lampe. RFID middleware design – addressing application requirements and RFID constraints. In Proceedings of SOC'2005 (Smart Objects Conference), pages 219–224, Grenoble, France, October 2005.

- [49] Christian Floerkemeier, Matthias Lampe, and Thomas Schoch. The Smart Box Concept for Ubiquitous Computing Environments. In Proceedings of SOC'2003 (Smart Objects Conference), Grenoble, France, May 2003.
- [50] Christian Floerkemeier and Friedemann Mattern. Smart Playing Cards – Enhancing the Gaming Experience with RFID. In Thomas Strang, Vinny Cahill, and Aaron Quigley, editors, Proceedings of the third Pergames Workshop at Pervasive 2006, pages 79–88, Dublin, Ireland, may 2006.
- [51] Christian Floerkemeier, Roland Schneider, and Marc Langheinrich. Scanning with a Purpose – Supporting the Fair Information Principles in RFID protocols. In Hitomi Murakami, Hideyuki Nakashima, Hideyuki Tokuda, and Michiaki Yasumura, editors, Ubiquitious Computing Systems. Revised Selected Papers from the 2nd International Symposium on Ubiquitous Computing Systems (UCS 2004), November 8-9, 2004, Tokyo, Japan, volume 3598 of Lecture Notes in Computer Science (LNCS), pages 214– 231, Berlin, Germany, June 2005. Springer-Verlag.
- [52] Christian Floerkemeier and Matthias Wille. Comparison of Transmission Schemes for Framed ALOHA based RFID Protocols. In Proceedings of the RFID and the Extended Network Workshop at SAINT 2006, pages 92–95, Phoenix, AZ, USA, jan 2006.
- [53] Nathan Freedman. RAYTAG, An Electronic remote Data Readout System. In Carnham Conference on Electronic Crime Countermeasures, April 1973.
- [54] Jean-Francois Frigon and Victor C. M. Leung. A Pseudo-Bayesian ALOHA algorithm with mixed priorities. Wireless Networks, 7(1):55-63, 2001.
- [55] Simson L. Garfinkel. Adopting Fair Information Practices in Low-Cost RFID Systems. Privacy Workshop at the International Conference on Ubiquitous Computing 2002 (Ubicomp2002), September 2002. Available from: www.simson.net/clips/academic/ 2002\_Ubicomp\_RFID.pdf.
- [56] A.V. Gershman and A. Fano. A wireless world: The Internet sheds its chains. Available from: www.accenture.com.

- [57] Bill C. Hardgrave, Matthew Waller, and Robert Miller. Does RFID Reduce Out of Stocks? A Preliminary Analysis, November 2005. Available from: http://itrc.uark.edu.
- [58] Peter Hawkes. Anti-Collision and Transponder Selection Methods for Grouped Vicinity Cards and RFID tags. In *Proceedings of IEE Colloquium on RFID Technology*, 1999.
- [59] Junius Ho, Daniel W. Engels, and Sanjay E. Sarma. HiQ: A Hierarchical Q-Learning Algorithm to Solve the Reader Collision Problem. In Proceedings of the RFID and the Extended Network Workshop at SAINT 2006, pages 88–91, Phoenix, AZ, USA, 2006.
- [60] Thomas J. Hutton. US Patent Nr. 3,964,024 Inductive coupling RFID, Sep 1975.
- [61] Sozo Inoue and Hiroto Yasuura. RFID privacy using usercontrollable uniqueness. MIT RFID Privacy Workshop, Cambridge, MA, USA, November 15, 2004. Available from: www. rfidprivacy.org/papers/sozo\_inoue.pdf.
- [62] International Organization for Standardization. ISO/IEC 15693: Identification cards – Contactless integrated circuit(s) cards – Vicinity cards, Oct 2003.
- [63] International Organization for Standardization. ISO/IEC 18000: Information technology automatic identification and data capture techniques - Radio frequency identification for item management air interface, 2003.
- [64] International Organization for Standardization. Information technology – Radio frequency identification for item management – Part 3: Parameters for air interface communications at 13,56 MHz, 2004.
- [65] International Organization for Standardization. Information technology – Radio frequency identification for item management – Part 6: Parameters for air interface communications at 860 MHz to 960 MHz, 2004.
- [66] Ari Juels and Ravikanth Pappu. Squealing euros: Privacy protection in RFID-enabled banknotes. In Rebecca N. Wright, editor, Proceedings of the 7th International Conference on Financial

Cryptography (FC 2003), volume 2742 of Lecture Notes in Computer Science, pages 103–121, Guadeloupe, French West Indies, January 2003. Springer-Verlag.

- [67] Ari Juels and Ronald L. Rivest. The blocker tag: Selective blocking of RFID tags for consumer privacy. In 10th Annual ACM CCS 2003, May 2003.
- [68] Günter Karjoth and Paul A. Moskowitz. Disabling RFID tags with visible confirmation: clipped tags are silenced. In Vijay Atluir, Sabrina De Capitani di Vimercati, and Roger Dingledine, editors, WPES, pages 27–30. ACM, 2005.
- [69] Udo Karthaus and Martin Fischer. Fully Integrated Passive UHF RFID Transponder IC With 16.7-μW Minimum RF Input Power. *IEEE Journal of Solid-State Circuits*, 38(10):1602–1608, Oct 2003.
- [70] A.R. Koelle, S.W. Depp, and R.W. Freyman. Short-Range Radio Telemetry for Electronic Identification Using Modulated RF Backscatter. In *Proceedings of IEEE*, pages 1260–1261, 1975.
- [71] Albert Krohn, Tobias Zimmer, Michael Beigl, and Christian Decker. Collaborative Sensing in a Retail Store Using Synchronous Distributed Jam Signalling. In Hans W. Gellersen, Roy Want, and Albrecht Schmidt, editors, *Third International Conference, PERVASIVE 2005*, number 3468 in Lecture Notes in Computer Science (LNCS), pages 237–254, Munich, May 2005. Springer-Verlag.
- [72] S.S. Lam and L. Kleinrock. Packet switching in a multi-access broadcast channel: Dynamic control procedures. *IEEE Trans. Commun.*, COM-23, 1975.
- [73] Matthias Lampe, Christian Floerkemeier, and Stephan Haller. Einführung in die RFID-Technologie. In Elgar Fleisch and Friedemann Mattern, editors, Das Internet der Dinge – Ubiquitous Computing und RFID in der Praxis, pages 69–86. Springer-Verlag, 2005.
- [74] Jeremy Landt. Shrouds of Time The history of RFID. AIM Publication, 2001. Available from: www.aimglobal.org.

- [75] Marc Langheinrich. A privacy awareness system for ubiquitous computing environments. In G. Borriello and L.E. Holmquist, editors, 4th International Conference on Ubiquitous Computing (UbiComp2002), number 2498 in Lecture Notes in Computer Science (LNCS), pages 237–245, Atlanta, GA, September 2002. Springer.
- [76] Marc Langheinrich. Personal Privacy in Ubiquitous Computing
   Tools and System Support. PhD thesis, ETH Zurich, Zurich, Switzerland, May 2005.
- [77] Marc Langheinrich. RFID and Privacy. In Milan Petkovic and Willem Jonker, editors, Security, Privacy and Trust in Modern Data Management, chapter RFID and Privacy. Springer-Verlag, 2006.
- [78] Ching Law, Kayi Lee, and Kai-Yeung Siu. Efficient Memoryless Protocol for Tag Identification. Technical Report MIT-AUTOID-TR-003, Auto-ID Center, 2000.
- [79] H. Lee, B. Peleg, P. Rajwat, S. Sarma, and B. Subirana. Assessing the Value of RFID Technology and the EPC Standard for Manufacturers. EPCGlobal White Paper, 2005.
- [80] Christine Legner and Frederic Thiesse. RFID-Based Facility Maintenance at Frankfurt Airport. *IEEE Pervasive Computing*, 5(1):34–39, Jan-Mar 2006.
- [81] Kin Seong Leong, Mun Leng Ng, and Peter H. Cole. The Reader Collision Problem in RFID Systems. In Proceedings of IEEE 2005 International Symposium on Microwave, Antenna, Propagation and EMC Technologies for Wireless Communications (MAPE 2005), Beijing, China, 2005.
- [82] Kin Seong Leong, Mun Leng Ng, Alfio Grasso, and Peter H. Cole. Synchronization of RFID Readers for Dense RFID Reader Environments. In *Proceedings of the RFID and the Extended Network* Workshop at SAINT 2006, pages 48–51, Phoenix, AZ, USA, jan 2006.
- [83] Masoud Mansouri-Samani and Morris Sloman. GEM: a generalized event monitoring language for distributed systems. *Distrib*uted Systems Engineering, 4(2):96–108, 1997.

- [84] Mantara. Elvin Router Product Datasheet, 2005. Available from: www.mantara.com.
- [85] Mantara Software. Elvin Subscription Language Reference, 4.0.0 edition, Dec 2003. Available from: www.mantara.com.
- [86] Duncan McFarlane. The Impact of Product Identity on Industrial Control - Part 1: - See More, Do More. Technical Report CAM-AUTOID-WH012, Auto-ID Center, Feb 2003. Available from: www.ifm.eng.cam.ac.uk/automation/ publications/w\_papers/cam-autoid-wh012.pdf.
- [87] Sun Microsystems. Java System RFID Software 3.0 Developer Guide. www.sun.com, Feb 2006.
- [88] Jin Mitsugi. UHF Band RFID Readability and Fading Measurements in Practical Propagation Environment. Auto-ID Lab Whitepaper Series Edition 1, Sep 2005.
- [89] Jin Mitsugi and Hisakazu Hada. Experimental Study on UHF passive RFID Readability Degradation. In *Proceedings of the RFID* and the Extended Network Workshop at SAINT 2006, pages 52– 55, Phoenix, AZ, USA, jan 2006.
- [90] Donald W. Neild. US Patent Nr. 3,290,675 Inductive coupling RFID, Sep 1975.
- [91] Oat Systems and MIT Auto-ID Center. The Savant Version 0.1. Technical Report MIT-AUTOID-TM-003, Auto-ID Center, 2002.
- [92] OATsystems. OAT C4 Architecture. www.oatsystems.com, 2006.
- [93] Organisation for Economic Co-operation and Development (OECD). Recommendation of the Council Concerning Guidelines Governing the Protection of Privacy and Transborder Flows of Personal Data, September 1980. Available from: www.privacy. gov.au/publications/oecdgls.pdf.
- [94] Philips Semiconductor. I-CODE1 System Design Guide, May 2002. Available from: www.semiconductors.philips.com/ acrobat\_download/other/identification/SL048611.pdf.
- [95] Philips Semiconductor. I-CODE1 Label ICs Protocol Air Interface Datasheet, January 2005. Available from:

www.semiconductors.philips.com/acrobat\_download/ other/identification/sl040616.pdf.

- [96] B. S. Prabhu, Xiaoyong Su, Harish Ramamurthy, Chi-Cheng Chu, and Rajit Gadh. WinRFID – A Middleware for the enablement of Radio Frequency Identification (RFID) based Applications. In Rajeev Shorey and Chan Mun Choon, editors, *Mobile, Wireless and Sensor Networks: Technology, Applications and Future Di*rections. Wiley, 2005.
- [97] Privacy Rights Clearinghouse. Position Statement on the Use of RFID on Consumer Products. Available from: www. privacyrights.org/ar/rfidposition.htm.
- [98] Theodore S. Rappaport. Wireless Communications. Prentice Hall PTR, second edition, 2002.
- [99] RFID Journal. Wal-Mart Draws Line in the Sand, June 2003. Available from: www.rfidjournal.com/article/view/462/1/ 1/.
- [100] Ronald L. Rivest. Network Control by Bayesian Broadcast. IEEE Transactions on Information Theory, IT-33(3):323–328, May 1987.
- [101] Fred S. Roberts. Applied Combinatorics. Prentics-Hall, 1984.
- [102] L.G. Roberts. ALOHA packet system with and without slots and capture. *Computer Communication Revue*, 5(2):28–42, 1975.
- [103] Kay Römer, Thomas Schoch, Friedemann Mattern, and Thomas Dübendorfer. Smart Identification Frameworks for Ubiquitous Computing Applications. Wireless Networks, 10(6):689–700, December 2004.
- [104] Sanjay Sarma. Integrating RFID. ACM Queue, 2(7):50–57, 2004.
- [105] Sanjay Sarma, David L. Brock, and Kevin Ashton. The Networked Physical World – Proposals for Engineering The Next Generation of Computing, Commerce & Automatic Identification. Technical Report MIT-AUTOID-WH-001, MIT Auto-ID Center, 2000.
- [106] Sanjay E. Sarma. Towards the Five-Cent Tag. Technical Report MIT-AUTOID-WH-006, MIT Auto-ID Center, 2001.

- [107] Sanjay E. Sarma, Stephen A. Weis, and Daniel W. Engels. RFID Systems and Security and Privacy Implications. In Workshop on Cryptographic Hardware and Embedded Systems, pages 454–470. Lecture Notes in Computer Science (LNCS), 2002.
- [108] Stefan Schlegel. RFID Framework. Master's thesis, ETH Zurich, Zurich, Switzerland, Mar 2004.
- [109] Thomas Schoch. *Cooperating smart objects.* PhD thesis, ETH Zurich, Zurich, Switzerland, January 2005.
- [110] Frits C. Schoute. Control of ALOHA Signalling in a Mobile Radio Trunking System. In International Conference on Radio Spectrum Conservation Techniques, pages 38–42. IEE, 1980.
- [111] Frits C. Schoute. Dynamic Frame Length ALOHA. *IEEE Trans*actions on Communications, COM-31(4):565–568, Apr 1983.
- [112] Bill Segall, David Arnold, Julian Boot, Michael Henderson, and Ted Phelps. Content Based Routing with Elvin4. In *Proceedings* AUUG2K, Canberra, Australia, June 2000.
- [113] Softronica. RIDEL5000 Long Range ICODE Reader/Encoder, July 2000. Available from: www.softronica.org.
- [114] S. Spiekermann and O. Berthold. Maintaining privacy in RFIDenabled environments: Proposal for a disable model. In P. Robinson, H. Vogt, and Waleed Wagealla, editors, *Privacy, Security,* and Trust within the Context of Pervasive Computing, volume 780 of The International Series in Engineering and Computer Science. Springer-Verlag, 2004.
- [115] Thorsten Staake, Frederic Thiesse, and Elgar Fleisch. Extending the EPC network: the potential of RFID in anti-counterfeiting. In SAC '05: Proceedings of the 2005 ACM Symposium on Applied Computing, pages 1607–1612, New York, NY, USA, 2005. ACM Press.
- [116] Frank Stajano. Security for ubiquitous computing. John Wiley & Sons, Ltd, 2002.
- [117] Fred Sterzer. An Electronic License Plate for Motor Vehicles. In RCA Review, volume 35, pages 167–175. RCA, June 1974.

- [118] Harry Stockman. Communication by Means of Reflected Power. In Proceedings of the IRE, pages 1196–1204, 1948.
- [119] Peter Sutton, Rhys Arkins, and Bill Segall. Supporting Disconnectedness - Transparent Information Delivery for Mobile and Invisible Computing. CCGrid 2001 IEEE International Symposium on Cluster Computing and the Grid, Brisbane, Australia, May 2001.
- [120] Symbol. Two RF Inputs Make a Better RFID Tag, Jan 2005. Available from: www.symbol.com/assets/files/2RFInputs\_ WP.pdf.
- [121] Frederic Thiesse. Architektur und Integration von RFID-Systemen. In Elgar Fleisch and Friedemann Mattern, editors, Das Internet der Dinge – Ubiquitous Computing und RFID in der Praxis, pages 101–118. Springer-Verlag, 2005.
- [122] Tibco. Tibco Rendezvous, March 2006. Available from: www.tibco.com/resources/software/messaging/ rendezvous\_ds.pdf.
- [123] H. Vogt. Efficient Object Identification with Passive RFID Tags. In F. Mattern and M. Naghshineh, editors, *First International Conference*, *PERVASIVE 2002*, volume 2414 of *Lecture Notes in Computer Science (LNCS)*, pages 98–113, Zurich, Switzerland, August 2002. Springer-Verlag.
- [124] W3C. XML Schema. www.w3.org/XML/Schema, Apr 2000.
- [125] J. Waldrop, D. W. Engels, and S. E. Sarma. Colorwave: an anticollision algorithm for the reader collision problem. In *Proceed*ings of the IEEE International Conference on Communications, volume 2, pages 1206–1210, Anchorage, Alaska, USA, 2002.
- [126] Fusheng Wang and Peiya Liu. Temporal management of RFID data. In Proceedings of the 31st international conference on very large data bases (VLDB), pages 1128–1139. VLDB Endowment, 2005.
- [127] Stephen A. Weis, Sanjay E. Sarma, Ronald L. Rivest, and Daniel W. Engels. Security and Privacy Aspects of Low-Cost Radio Frequency Identification Systems. In *First Annual Conference on Security in Pervasive Computing*, 2003.

- [128] M. Weiser, R. Gold, and J.S. Brown. The origins of ubiquitous computing research at PARC in the late 1980s. In *IBM Systems Journal*, pages 693-696, 1999. Available from: www.research. ibm.com/journal/sj/384/weiser.html.
- [129] Jeffrey E. Wieselthier, Anthony Ephremides, and Larry A. Michaels. An Exact Analysis and Performance Evaluation of Framed ALOHA with Capture. *IEEE Transactions on Communications*, COM-37(2):125–137, 1989.
- [130] Wikipedia. Theremin. Available from: http://en.wikipedia. org/wiki/Theremin.
- [131] Matthias Wille. Evaluation and Optimization of RFID Transmission Control Strategies. Master's thesis, ETH Zurich, March 2005.
- [132] Bin Zhen, Mamoru Kobayashi, and Masashi Shimizu. Framed ALOHA for Multiple RFID Objects Identification. *IEICE Trans*action on Communications, E88-B(3):991-999, March 2005.

# Curriculum Vitae

# Christian Floerkemeier

#### Personal Data

Date of Birth	February 14, 1975
Birthplace	Braunschweig
Citizenship	$\operatorname{German}$

### Education

1981 - 1985	Bredde-Schule, Witten, Germany
1985 - 1991	Schiller-Gymnasium, Witten, Germany
1991 - 1992	Oregon Episcopal School, Portland, Oregon, USA
1992 - 1994	Schiller-Gymnasium, Witten, Germany
June $1994$	Abitur
1995 - 1999	Cambridge University, United Kingdom
June 1999	Graduation with a BA and Master of Engineering in Elec- trical and Information Science
2001-2006	Ph.D. Student at the Department of Computer Science, ETH Zurich, Switzerland
Civil Service	
1994 - 1995	St. Marien Hospital, Witten, Germany

## Employment

1999-2001	Head of software development, Ubiworks, Amsterdam, The
	Netherlands
2001-2006	Research Assistant at the Department of Computer Science, ETH Zurich, Switzerland