

## Introduction

RFID and telemetry sensor data add tremendous value to supply chain logistics. RFID data provides an account of where objects have been and when they have been there. Telemetry sensors measure environmental attributes such as temperature, light, humidity, pressure, vibration, and sound. Combining these data, not only do we have an account of object traversal through a supply chain, but we also have an account of the environment within which objects exist.

Potential applications range from tracking a set of objects for expedited recall or delivery, to insuring the quality of environment-sensitive goods such as food or medicine, to detecting breaches in shipping containers. Two particular applications we focus on are Track&Trace and Cold-Chain. In Track&Trace one wishes to track the current location of a particular object, or trace its traversal in the supply chain. In Cold-Chain one wishes to insure the quality of temperature-sensitive goods as they traverse the supply chain.

In order to build such applications atop RFID and telemetry sensor data, we need a standard way to describe these data. We need efficient ways to construct and gather data.

Furthermore, we need a sensible way to build applications on distributed data and machines. These applications must not only be distributed themselves, but also must allow for dynamic reconfiguration so that new models and applications may be seamlessly deployed.

We proceed to describe the physical hardware of sensors and some related issues for applications in Chapter 1. In Chapter 2, we propose a data structure to represent and process streaming RFID and sensor data, in addition to ancillary relational data. In Chapter 3, we describe efficient ways to sample RFID and sensor data and to construct these low-level data. We next describe a system architecture in Chapter 4 for building systems on distributed data and machines, and we explore how to optimize this architecture and also allow for dynamic reconfiguration to modify existing applications and to deploy new ones. Lastly, in Chapter 5 we describe our particular implementation of two model applications: Track&Trace and Cold-Chain.