3-5-2015

# Moving Populations Event Recognition Under Re-Identification and Data Locality Constraints

Wolfgang Maaß

Tom Michels

# Moving Populations Event Recognition Under Re-Identification and Data Locality Constraints

Wolfgang Maaß, Tom Michels

Saarland University, Saarbrücken, Germany
{wolfgang.maass,tom.michels}@iss.uni-saarland.de

**Abstract.** For more than a decade tracking and tracing physical objects has been target of information systems within the realm of research on the Internet of Things. But application to human populations requires reconsideration of re-identification and data locality requirements due to ethical and legal constraints. For this domain, we propose a generic event recognition architecture (GERA) and evaluate its applicability for developing a sensor-based information system for recognizing moving population densities by obeying non-re-identification and data decentrality requirements. Empirical evaluations show that this information system provides mean structures for measuring event data and deriving predictions that are statistically equal to manually measured actual data. Finally, a general discussion on the integration of event recognition systems into business process environments is given.

**Keywords:** event recognition, predictions, population densities, sensor-based systems, Internet of things

## 1 Introduction

Integrating real-world activities with business processes is central to any Internet of Things application [1] and recently to Industry 4.0 [2]. Logistics has been a trailblazer for RFID-based systems that give unique identifiers to physical entities by means of small chips [3]. As part of the EPC Global Network architecture, incidents are measured by sensing systems and translated into business events that are used as input for business processes [4].

For various reasons, societies abstain from applying these approaches directly to humans, i.e. by tagging humans with electronic devices. In several countries, storing and evaluating video data in public and commercial spaces is subject to strong restrictions. The deployment of cameras with advanced image processing algorithms turned out to be a reasonable approach for capturing crowd densities or recognizing panic movements in big crowds, but poses difficulties with respect to real-time analysis, privacy concerns and large scale deployments [5], [6], [7]. Therefore passive, decentralized, and non-identifying measures are better indicated for measuring population densities. Measuring population densities is a widespread task when heterogeneous masses of people move. Examples are sports arenas, exhibitions, fairs, malls, music concerts,

and many others situations. For queue management [8], actual and perceived population densities are important factors for managing resources, such as in retail stores [9]. However, for many customers, waiting in a service is associated to negative experiences [10]. Researchers have argued, that managing waiting lines can generally be influenced with two different approaches: operation management or perception management [11]. In this paper we focus on the operation management approach, but extend the general idea, that this approach is only based on improving and ameliorating the effectiveness of the service provider dispatching speed.

A general sensor-based event recognition model is introduced that translates incidents in real-world into business events as input for business processes. This general model is applied to measuring moving population densities with spontaneous behavior. On technological level, a non-invasive event-capturing sensor-based system is presented that measures population densities and derives predictions in real-time without interaction with central services which makes it applicable for any environment. Limited computational resources of smart sensors require massive data reductions by pre-processing incident data effectively (cf. [12]). Due to the advanced pre-processing capabilities, real-time event recognition and prediction processing are supported. In fact, our approach resembles the ATLAS experiment for the Large Hardon Collider (LHC) at Cern, where 300.000 MByte/s have been reduced around 300 MByte/s by various filters [13].

Our research follows the problem solving paradigm and adopts a design science research approach [14, 15]. In this paper we focus on the building phase for constructing artifacts that enable information systems to track population frequencies without re-identification. On the technical level, we obey a decentrality constraint, i.e. data is processed close to its incidents. Our work contributes to research on the Internet of Things and in particular to decentralized control and data processing [16] and real-time information systems [17].

This paper is structured as follows. Section 2 discusses related research results Section 3 introduces generic architecture for event recognition. Section 4 presents an application of event recognition with moving populations. Section 5 evaluates and discusses the accuracy of the generated events and section 6 presents a summary and conclusion of our paper and discusses open issues.

## 2      Related work

A variety of vision-based approaches for capturing movement data in public areas exist [18-20]. Drawing conclusions on populations based on image-processing algorithms has been research for years, but is mostly being used to draw conclusions on the denseness of a crowd or recognize unnatural and panic movements [5-7]. Counting the number of people based on image-processing algorithms has not been widely research and mainly has the disadvantages, that is might be considered as invasive by customers and that the counting algorithm has to be calibrated for each new environment [21].

Another method for capturing movement data is based on GPS-sensors in smartphones. This approach has been successfully deployed during the Olympic games in London in 2012 [22]. It allowed security personnel to determine where conglomerations of people exceeded critical levels. However, due to the lack of GPS signals indoors and privacy concerns, this approach is questionable for the purpose of this paper.

A network for indoor and outdoor air quality monitoring based on sensors was presented in [23]. Advanced processing of multiple-input single-output neural networks has been implemented at the smart sensors to obtain temperature and humidity compensated gas concentration levels. Others presented smart sensor platforms focussing on the advantages and disadvantages of the different wireless communication and networking challenges (e.g., [24]).

In general, integration of physical real-worlds events with digitized business processes is the target of any ERP system approach [25]. With further developments of low-cost optical and wireless sensor technologies, such as RFID [26], QR tags, and reacTIVision[1], tight integration of objects and digital representations are supported for practical applications. Highly structured business processes with a focus on physical products show strong efficiency gains by applying RFID technologies implementing the EPC Global Network Architecture [4]. Event recognition is a central part of this architecture handled by the EPCIS Capturing Application that, in turn, is based on applications for filtering and collecting raw tag reads (ibid.). Example events are tracking and tracing events for products along the supply chain.

In contrast to products, event recognition of humans underlie severe legal and ethical restrictions in many European countries. What can be applied to physical products is generally unacceptable for humans. It is an ongoing debate colored by cultural backgrounds what is perceived as being acceptable and lawful, e.g. differences between the UK and Germany. Therefore a discussion of requirements is crucial for any human event recognition application. For a more specific conceptualization, two dimensions for initially distinguish human event recognition applications are introduced: (1) re-identification and (2) data locality. Re-identification defines the possibility to derive the identity of a recognized entity. In RFID applications re-identification is central because each EPC provided by a RFID tag provides a unique identifier for a physical entity [4]. For humans, re-identification is perceived as an ethical threat [27]. Therefore, event recognition of human behavior might have strong limitations on re-identification. In public spaces, re-identification is generally forbidden and typically requires criminal prosecution for overriding. Data locality defines the location where various forms of data are stored that is related to event recognition. This dimension is delimited by decentrally storing data on capturing devices only and storing all data on central systems with any combination and middle layers inbetween. Again, typical RFID applications use a central storage approach by which raw data is transferred to central databases due to little storage capacity on capturing devices [12]. In contrast, applications applying a decentralized storage approach capture raw and processed data on capturing devices and only transfer qualified data to

---

[1] http://reactivision.sourceforge.net

more central applications. Initial examples for the latter approach can be found in the medical field that enable storage of raw data near capturing devices and data exchanges are governed by decentralized permission management [28].

Based on re-identification and data locality, requirements for measuring moving population event recognition can be derived. In this paper, we will focus on an event recognition system for moving population densities under the requirement of non-re-identification of individuals and by applying a decentralized storage and processing approach.

## 3 Generic Event Recognition Architecture

In the following, a generic event capturing model is presented that generalizes on object-centric approaches, such as the EPC Global Network Architecture [4].

### 3.1 System view

The generic event recognition architecture (GERA) consists of three modules: (1) incident detector, (2) filtering, and (3) event producer. The event producer contains a set of sub-modules responsible for analytics of event related data, aggregation of data, and higher-order processing, such as prediction modules. The incident detector is the initial starting or triggering point of the system. At this point, raw data on incidents i.e. movements are being detected and captured through multiple sensing elements. In case of continuous movements, a data stream emerges. Subsequently, the respective data is forwarded to the filtering module. Since our approach is considering dynamic environments with non-identifiable moving individuals, the filtering module is playing a key role in reducing noise and, thus, increasing the signal-to-noise ratio.

Subsequently, the data is being aggregated to more abstract data types, for instance, population denseness per time period. An event producer operates in two modes: (1) accumulative mode or (2) pre-processing mode. In the accumulative mode, event producers add abstracted data to raw data while raw data is replaced by abstracted data in the pre-processing mode. The latter mode is important for event recognizing applications with data streams that exceed storage capacity (cf. [29]). For instance, the ATLAS experiment for the Large Hardon Collider (LHC) at CERN, generates 300.000 MByte/s by incident detectors which exceeds storage capabilities. Therefore this data stream is reduced to 300 MByte/s abstracted data by various filters [30].

Event producers provide events to other (business) processes that use events as triggers for various purposes, e.g. activation and deactivation of processes, personalization, and contextualization. More abstract, event producers act as bridges between activities in physical to digital environments. Actuation producers take a reverse orientation, i.e. from digital environments to physical environments.
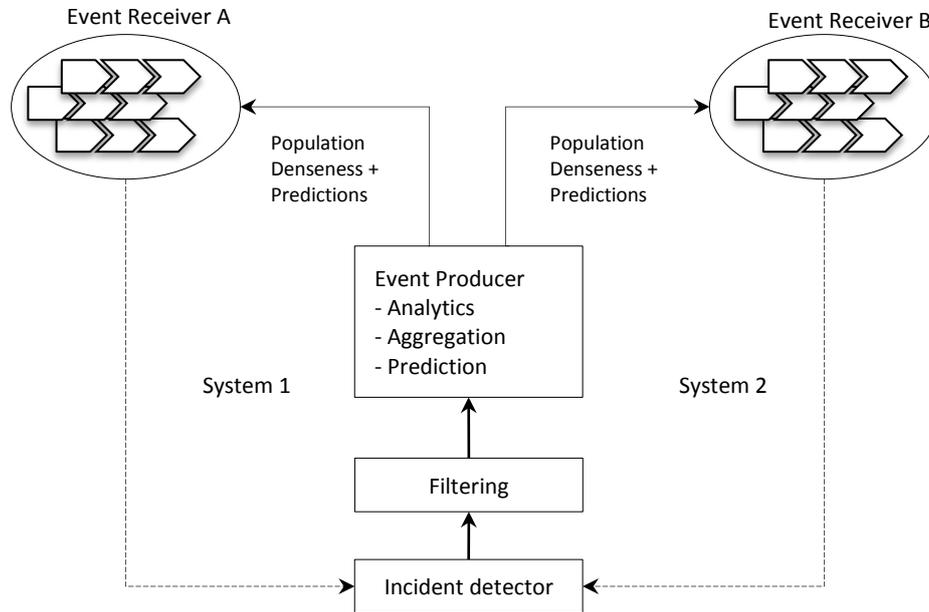
**Fig. 1.** System View

# 4 Application of event recognition with moving populations

We have tested the GERA model by applying it to the domain of student cafeterias. This domain fulfills the requirement of strongly restricted re-identification by event recognition. Furthermore university information systems are centralized, in large parts customized and historically grown systems. Extensions of these information systems are rarely possible and tend to be expensive. Therefore, decentralized data locality was defined as a requirement. According to GERA, domain-specific modules are described before details are given for human-computer interfaces used by students and members of the cafeteria workforce. Technical details are described at the end of this section.

## 4.1 System architecture according to GERA

In order to validate GERA by a feasibility study, an information system was developed and tested during two weeks at a University cafeteria. The incident detector consists of two distance sensors that were placed at a stair leading to a particular menu service of the student cafeteria. Distance sensors measure the depth of free space orthogonal to the movement direction. Maximum distance measures if no person is standing in front of a sensor. Both distance sensors were located next to each other along the movement direction. Thus, the second sensor detected a person slightly after the first sensor. Distances are only measured if incidents occur.

Measured distances of both sensors together with time stamps were forwarded to the filtering module where unrealistic and false values i.e. negative distances or distances exceeding the measurement range of the sensor were discarded. Distances values are subsequently being forwarded to the event producer.

The physical design with two sensors supports event recognition by an hierarchical finite state machine (HFSM) [31]. **Table 1** represents the possible states occurring in the state machine. Note that both, initial state *I* and goal state *G*, are represented by the same combination of sensor states, since *before/after* every movement detection cycle both sensors will not be covered.

**Table 1.** Step model

| State of HFSM | State of sensor 1 (S1) | State of sensor 2 (S2) |
|---|---|---|
| I/G | Not Covered | Not covered |
| 1 | Covered | Not covered |
| 2 | Covered | Covered |
| 3 | Not covered | Covered |

The HFSM is illustrated in **Fig. 2**. Labels on nodes represent the state of the HFSM. Labels on the edges illustrate incidents, representing a state change of a sensor, ie. from covered to non-covered or vice-versa. As an example, assuming the state machine is at state 1, and the incident "S2 covered" has been detected, the state of the HFSM moves from state 1 to 2.
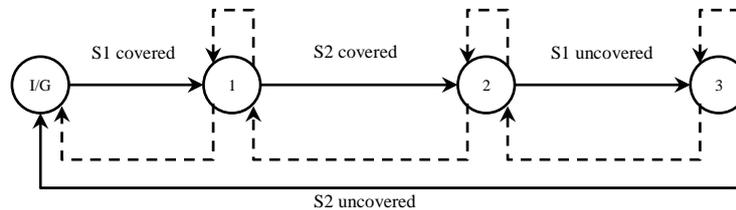


**Fig. 2.** State machine process

The dotted lines represent iterations over the same states or fallbacks to a lower state. In order to ensure clarity of **Fig. 2**, the incidence represented by those line have not been labeled. With the above-presented approach, movements in front of the incidence detectors are only considered as valid (and thus increase the measured population denseness) if the HFSM completes one cycle. Furthermore we are able to handle movements to and away from the incidence detectors (distance sensors), since those do not evoke a jump to another state of the HFSM.

The number of cycles that the state machine has been traversed during a predefined time period (in the study it was set to 10 minutes) represents the measured denseness of a moving population. At the end of each time period, all detected cycles are condensed to one event, representing the denseness of moving population. In the follow-

ing paragraphs, we refer to these events as *moving population denseness* (MPD) events. Each MPD event consists of a denseness value, a time stamp, and the length of the time period.

Based on past MPD events, the prediction module calculates a forecast for the next time period, i.e. for the next 10 minutes. Forecast values are referred to as *predicted moving population denseness* (pMPD) events. A forecasting algorithm with exponential smoothing and seasonal adjustment was selected because afflux to the cafeteria has a non-linear pattern for each day, starting low after opening, with a peak, and ending low before closing. Exponential smoothing supports quick reactions to recent changes due to decreasing weights over time. The forecasting algorithm has been implemented according to [32] and is sketched below.

$$I_t = \frac{A_t}{\bar{A}} \tag{1}$$

with $I_t$ being the seasonal index, $A_t$ the number of persons having passed the sensor at time $t$ and $\bar{A}$ the average number of persons having past the sensor in the last season.

$$S_t = \alpha \frac{A_t}{I_{t-L}} + (1 - \alpha) S_{t-1} \tag{2}$$

with $S_t$ being the smoothed value, $\alpha$ the smoothing factor, $I_t$-$L$ the seasonal index at time $t$ during the last season.

$$F_{t+1} = S_t * I_{t+1-L} \tag{3}$$

with Ft+1 being the forecast for period t+1. Regarding the presented setup above, the seasonality is given by a day, since the afflux of people visiting the cafeteria is depending on their weekly schedule and lectures..

MPD and pMPD events are final outputs of the event producer. Both events are being sent to event receiver community *A* representing students that are likely to go to the restaurant in the short-term and event receiver community *B* representing the workforce of the cafeteria (cf. Fig. 3).
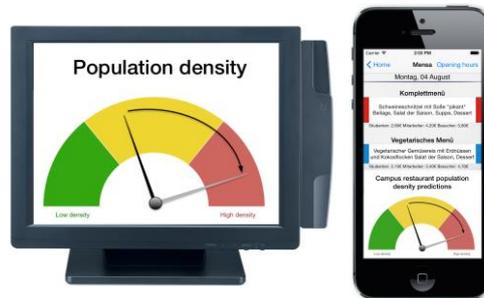


**Fig. 3.** Actual and forecasted population denseness representation for the vent receivers

### 4.2 Technical Realization

Two infrared analog distance sensors by sharp with a maximal measurement distance of 150 cm have been connected through a PhidgetInterfaceKit[2] to a Raspberry Pi[3] Model B. We further installed a Raspbian[4] Linux distribution on the Raspberry Pi and handled all pre-processing activities including filtering, analytics, aggregation, and event procedures implemented in *Java*. The complete setup is powered by a portable battery pack ensuring a very flexible and wireless deployment. The technical setup and installation of the smart sensor on the handrail is presented in **Fig. 4**.



**Fig. 4.** Study setup at the University cafeteria

## 5 Evaluation

### 5.1 MPD events

The precision of the sensory system for measuring moving population densities was evaluated by comparing system-defined MPD values and manually counted values during one day. The normalized root-mean-squared error (NRMSE) [33] between measured values by the system and manually counted number of persons is quite high (36%) for the respective day. Furthermore, the population denseness is consistently underestimated. The reason is based on the fact, that the sensor does not recognize more than one person at a time. Thus, if two individuals are walking next to each other and pass the sensor at the same time, the sensor is not able to detect both individuals. Due to the relatively constant under-estimation of the sensor, inflating the

---

[2] PhidgetInterfaceKit is an I/O Board that converts analog inputs from sensors to digital outputs. (http://www.phidgets.com/products.php?category=0&product_id=1018_2)

[3] Raspberry Pi Model B (http://www.raspberrypi.org/product/model-b/)

[4] Raspbian is a free operating system based on Debian optimized for the Raspberry Pi hardware. (http://www.raspbian.org)

estimated number by a constant factor $k$ can significantly improve the accuracy. As an example, inflating the computed number of persons by a factor $k = 1.5$ would decrease the NRSME from 36% down to 14% (cf. **Fig. 5**).

For a more detailed analysis of the correlation between both samples, a two-sided paired difference *t*-test was applied evaluating the null hypothesis that the mean value of measured and actual denseness of moving population are statistically equivalent. A correlation between the two samples being at r = 86.7% supports the null hypothesis. By assuming a higher significance level of 5% as typical for exploratory studies, the null hypothesis cannot be rejected (two-sided paired difference t-test, .055 p-value). Thus, the mean values of measured and actual denseness of moving population cannot be considered as being unequal for our sample. Hence, MPD events are statistically significant approximations for the longitudinal curve of actual values.
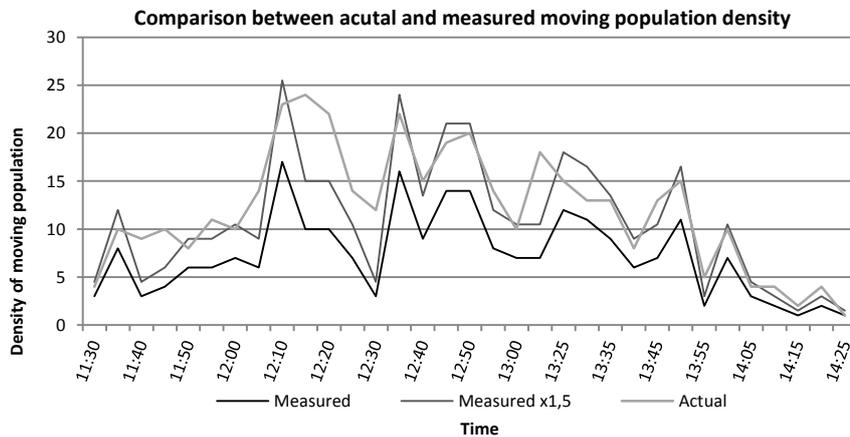


**Fig. 5.** Measurement accuracy

### 5.2    pMPD events

Taking measured values for moving population densities as proxies for actual values, we compare measured values with predicted values for pMPD. Data collection was performed over four days. It is not surprising that the NRSME is lower for measured and actual data because predictions are based on measured values. NRSME between measured and predicted values settled at about 21% (cf. **Fig. 6**).

As before, a two-sided paired difference t-test was applied evaluating the null hypothesis that the mean value of measured and predicted denseness of moving population are statistically equivalent. 144 time periods were collected during 4 days resulting in a correlation of 61.6% (.000 p-value) between measured and predicted samples. Again, with a significance level of 5%, the null hypothesis cannot be rejected (two-sided paired difference t-test, p=.268), ie. the mean values of measured and predicted denseness of moving population cannot be considered as being unequal with respect

to our sample (cf. Fig. 6 for one day). In more than 26% of all cases, we find support for the null hypothesis. Hence, there is not sufficient evidence to reject the null hypothesis. This is taken as a robust result for an approximation with a standard forecast algorithm without considering deep context information.
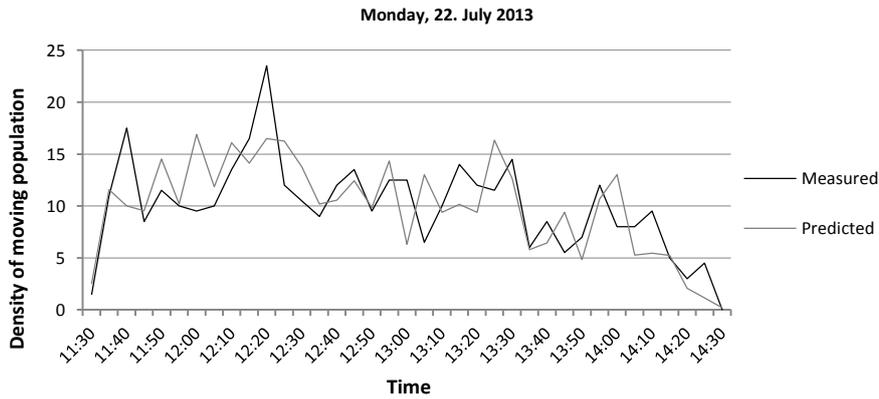
**Monday, 22. July 2013**



**Fig. 6.** Predicted vs. measured moving population densities (pMPD vs. MPD)

## 5.3 Discussion

By this experimental setup it was found that movement population densities are effectively measured by a distance sensor system that obeys non-re-identification and decentralization requirements. It was a crucial requirement that customers of the student cafeteria could not be tagged or traced. Additionally, budget restrictions did not allow integration into existing, central information systems. Thus a decentralized approach becomes a necessity.

Two hypotheses were tested empirically. First, it was found that longitudinal event values measured (MPD values) by the sensor system provide a data basis that is not significantly different from actual values of movement population densities. Higher levels of absolute correspondence are achieved by scalar weights. Therefore, measured event data has proven to be appropriate proxies for actual event data on movement population densities.

Additionally, it was found that predicted event data provides similar characteristics on the mean structure, ie. is not statistically dissimilar from measured event values. This means that the prognosis based on a forecast algorithm with exponential smoothing with seasonal adjustment provides a statistically appropriate approximation of measured values. Furthermore it is concluded based on results of the first part that predicted event values are also a good approximation of actual events.

# 6 Summary, limitations, and open issues

In this paper, we presented an event recognition system based on sensors that is governed by non-re-identification and decentralized data locality requirements. It was found that measured values strongly correlate with actual event values and that predicted values are strongly correlated with measured values. This setup support pre-processing of data streams on decentralized sensor systems with low-cost components. Compared to other approaches (e.g., [34] and [7]), the effort of deploying our sensor system is considered as relatively low and effective.

On the measurement side we found high approximation values for heterogeneous moving populations. Nonetheless, the physical layout of the environment is assumed to be crucial for the MPD and pMPD calculations. For instance, if the average number of individuals moving together through the distance sensing system underlies large variations, the scalar weight $k$ needs to be replaced by functions in particular using machine learning approaches.

Furthermore, we indicated how the event recognition system connects real-world events with digital events and how this provides information to two communities, ie. a student community and cafeteria workforce community. Currently we investigate how both communities are actually affected by this dynamic information about actual and predicted events and corresponding situations.

Another issue for future work is the extension of seasonality. For the cafeteria domain context information can be included. For instance, we assume that the type of meal has significant effects on movement population densities with respect to absolute and relative values. Due the different class schedules of students during the week, the number of students being on campus around lunchtime varies over the different weekdays. With a longer deployment period of our setup, the seasonality index (currently based on a daily basis), could be calculated week-based. This change would indirectly consider the varying number of students on campus during the different weekdays and thus further ameliorate the accuracy of our predictions.

In our current study, we limited our study to one staircase only. Extension to all entry points of the cafeteria will allow us to analyze correlations and non-parametric analyses between movement population densities of different entry points.

In summary, we found evidences that the proposed sensor system approach with non-re-identification and decentralized data locality requirements provide an effective means for connecting real-world situations with business processes. Further research will target measuring the impact of this approach on, for instance, queue management, reduction of waiting times, decision making in retail stores, hotels or similar domains.

# References

1.      Atzori, L., Iera, A., Morabito, G.: The internet of things: A survey. Computer networks 54, 2787-2805 (2010)

2. Kagermann, H., Wahlster, W., Helbig, J. (eds.): Recommendations for implementing the strategic initiative INDUSTRIE 4.0. acatech e.V., Frankfurt/Main (2013)

3. Bose, I., Pal, R.: Auto-ID: managing anything, anywhere, anytime in the supply chain. Communications of the ACM 48, 100-106 (2005)

4. Traub, K.: The EPCglobal Architecture Framework (EPCglobal Final Version). (2005)

5. Ihaddadene, N., Djeraba, C.: Real-time crowd motion analysis. Pattern Recognition, 2008. ICPR 2008. 19th International Conference on, pp. 1-4. IEEE (2008)

6. Marana, A., Velastin, S.A., Costa, L.d.F., Lotufo, R.: Automatic estimation of crowd density using texture. Safety Science 28, 165-175 (1998)

7. Davies, A.C., Yin, J.H., Velastin, S.A.: Crowd monitoring using image processing. Electronics & Communication Engineering Journal 7, 37-47 (1995)

8. Fitzsimmons, J.A., Fitzsimmons, M.J.: Service management: Operations, strategy, and information technology. (2004)

9. Turley, L.W., Milliman, R.E.: Atmospheric effects on shopping behavior: a review of the experimental evidence. Journal of Business Research 49, 193-211 (2000)

10. Taylor, S.: Waiting for service: the relationship between delays and evaluations of service. The Journal of Marketing 56-69 (1994)

11. Katz, K.L., Larson, B.M., Larson, R.C.: Prescription for the waiting-in-line blues: entertain, enlighten, and engage. Sloan Management Review 32, 44-53 (1991)

12. Melski, A., Thoroe, L., Schumann, M.: Managing RFID data in supply chains. International Journal of Internet Protocol Technology 2, 176-189 (2007)

13. Aad, G., Abat, E., Abdallah, J., Abdelalim, A., Abdesselam, A., Abdinov, O., Abi, B., Abolins, M., Abramowicz, H., Acerbi, E.: The ATLAS experiment at the CERN large hadron collider. Journal of Instrumentation 3, S08003 (2008)

14. Hevner, A.R., March, S.T., Park, J., Ram, S.: Design Science in Information Systems Research (1). MIS Quarterly 28, 75-105 (2004)

15. March, S., Smith, G.: Design and natural science research on information technology. Decision Support Systems 15, 251-266 (1995)

16. Fleisch, E., Dierkes, M.: Ubiquitous Computing aus betriebswirtschaftlicher Sicht. Wirtschaftsinformatik 45, 611-620 (2003)

17. Cundius, C., Alt, R.: Real-Time or Near Real-Time?-Towards a Real-Time Assessment Model. International Conference on Information Systems (ICIS), Mailand (2013)

18. Junior, S.J.: Crowd analysis using computer vision techniques. IEEE Signal Processing Magazine 27, 66-77 (2010)

19. Ali, S., Shah, M.: Floor fields for tracking in high density crowd scenes. Computer Vision–ECCV 2008, pp. 1-14. Springer (2008)

20. Chan, A.B., Liang, Z., Vasconcelos, N.: Privacy preserving crowd monitoring: Counting people without people models or tracking. Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on, pp. 1-7. IEEE (2008)

21. Ryan, D., Denman, S., Fookes, C., Sridharan, S.: Crowd counting using multiple local features. Digital Image Computing: Techniques and Applications, 2009. DICTA'09., pp. 81-88. IEEE (2009)

22. Wirz, M., Schläpfer, P., Kjærgaard, M.B., Roggen, D., Feese, S., Tröster, G.: Towards an online detection of pedestrian flocks in urban canyons by smoothed spatio-temporal clustering of GPS trajectories. Proceedings of the 3rd ACM SIGSPATIAL international workshop on location-based social networks, pp. 17-24. ACM (2011)

23. Postolache, O.A., Pereira, J.M.D., Girão, P.S.: Smart sensors network for air quality monitoring applications. Instrumentation and Measurement, IEEE Transactions on 58, 3253-3262 (2009)

24. Ramamurthy, H., Prabhu, B., Gadh, R., Madni, A.M.: Wireless industrial monitoring and control using a smart sensor platform. Sensors Journal, IEEE 7, 611-618 (2007)

25. Becker, J., Kugeler, M., Rosemann, M.: Process Management: a guide for the design of business processes: with 83 figures and 34 tables. Springer (2003)

26. Want, R.: RFID, A Key to Automating Everything. Scientific American 290, 56-65 (2004)

27. Podesta, J., Pritzker, P., Moniz, E., Holdren, J., Zients, J.: Big data: seizing opportunities, preserving values. Executive Office of the President. Report. Retrieved on May 1, 2013 (2014)

28. Almeida, J.S., Deus, H.F., Maass, W.: S3DB core: a framework for RDF generation and management in bioinformatics infrastructures. BMC bioinformatics 11, 387 (2010)

29. Cheong, T., Kim, Y.: RFID data management and RFID information value chain support with RFID middleware platform implementation. On the Move to Meaningful Internet Systems 2005: CoopIS, DOA, and ODBASE, pp. 557-575. Springer (2005)

30. Dos Anjos, A., Armstrong, S., Baines, J.T.M., al., e.: Muon Identification with the Event Filter of the ATLAS Experiment at CERN LHC's. (2005)

31. Girault, A., Lee, B., Lee, E.A.: Hierarchical finite state machines with multiple concurrency models. Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on 18, 742-760 (1999)

32. Fitzsimmons, J.A., Fitzsimmons, M.J.: Service management: operations, strategy, and information technology. McGraw-Hill New York (2006)

33. Signoretto, M., Van de Plas, R., De Moor, B., Suykens, J.A.: Tensor versus matrix completion: a comparison with application to spectral data. Signal Processing Letters, IEEE 18, 403-406 (2011)

34. Ryan, D., Denman, S., Sridharan, S., Fookes, C.: Scene invariant crowd counting and crowd occupancy analysis. Video Analytics for Business Intelligence, pp. 161-198. Springer (2012)