

Enabling business process integration of IoT-events to the benefit of sustainable logistics

Markus Schief¹, Christian Kuhn¹, Philipp Rosch¹, Todor Stoitsev¹

¹ SAP Research, Bleichstraße 8, Darmstadt, Germany

*markus.schief@sap.com, christian.kuhn@sap.com, philiroesch@sap.com,
todor.stoitsev@sap.com*

Abstract: In this paper, we introduce a concept that focuses on innovative sustainability solutions in the logistics domain reflecting process-embedded events from the Internet of Things. We outline sustainability challenges in the target industry and derive generic software requirements. To address them a conceptual framework is presented in line with the two core functionalities. Firstly, we introduce event-based processing and quality methodologies to retrieve valuable information from the Internet of Things. The other subset of key concepts pertains to the resulting process-centric analysis and decision making in the context of user-process interaction. The developed concepts are derived from experiences applying recent research advances to sustainability in the logistics industry. Nevertheless, the presented functionalities are designed to be applicable to further scenarios as well.

Keywords: Internet of Things, Event Processing, User-Process Interaction, Sustainability, Logistics

1. Introduction

The logistics industry is challenged as the material flow network is getting more and more complex, multi-modal, and tightly interlinked with other processes such as manufacturing. While companies gradually outsource their logistic processes, logistics service providers take responsibility for the complete supply chain. Those companies orchestrate partners and processes and optimize the logistics in terms of cost, quality, and reliability. Complexity often increases with the number of involved partners. The more companies contribute to a logistics service provider network, the greater is the need for high-quality, up-to-date, and shared information (Kaplan et al., 2010). Thus, the efficient and

effective coordination of interactions is an important success factor (Baumoel and Winter, 2001).

Likewise, sustainability has become an important business criterion. The increasing demands of society, governments, and customers for sustainable logistics converge former disjunction ecological and economic goals to the direction of 'Green Logistics'. Governance, Risk, and Compliance (GRC) regulations require logistic firms to interact with authorities in various areas such as environmental safety.

The sustainable-oriented orchestration of logistic processes is a key competency. Respective people in charge are hence in the sweet spot of our scenario (Section 2). They align the design and execution of logistic processes with the specific business goals. First and foremost, the choice of logistic partners is a key challenge. Ideally, partner selection and capacity planning can be based on sustainability indicators (e.g. actual fleet consumption) that go beyond the traditional selection criteria such as price and service levels. With that respect, the calculation of eco-balances for each shipment is a prerequisite. They cannot only be used as compliance evidence but also as a business differentiator.

Performance of the relevant ecological indicators needs to be measured and analyzed. With that respect, we differentiate between direct and indirect information needs. Direct ones comprise data such as fuel consumption of transport vehicles or resulting carbon footprints (BSI, 2008). In contrast, indirect information can reveal root causes for variances of the mentioned parameters (Simons et al., 2003). For instance, traffic jams or detours due to ecological zones can explain consumption variances if a relation between both can be drawn. So, additional information can significantly support the analysis, if the information is put into context and hence provides value.

The various information sources require an information infrastructure. The Internet of Things (IoT) and the Internet of Services (IoS) can provide a lot of relevant data. With today's IoT-technologies, for instance, fuel consumption and Co2 emissions can be measured at the point of usage and communicated to information systems. In contrast, the IoS rather provides information such as Co2 emission regulations. Information from both infrastructures need to be mapped to conclude appropriate actions (see also (Maass and Filler, 2007)).

For completeness and appropriate analysis, the data measurement level is very important. It needs to be captured on two major logistic allocation bases: shipments and transport vehicles. Where-used lists can put both information streams into context. Dedicated shipments can then be mapped to the vehicles

involved in the transport. Finally, an eco-balance per shipment can be calculated. To make the provided information really matter, the quality of the processed data events must be very high. People receiving information need to have confidence in the data in order to avoid false conclusions. In addition to quality, traceability of data origin can increase trust and analysis possibilities. So, the information processing must be of high specificity, quality, and traceability.

Finally, the provided information can only be leveraged if it is aligned with a human-process interaction concept. The daily job of respective knowledge workers demands various skills and manifold knowledge. Beyond industry domain knowledge, process management skills, the analysis and synthesis of information, as well as decision capability belong to their core competencies. So, employees need to be supported in information analysis, decision simulation, and the implementation of changes. From a business point of view, the underlying revenue and cost framework should be included into this concept. If information can be mapped to meaningful associated costs and revenues, the business value becomes more transparent. For instance, activity-based costing can be enhanced to reflect toll-, Co₂-, and fuel-consumption costs based on the IoT-data. Potential trade-offs hence become transparent in a meaningful business context and proper decisions can be made.

2. Scenario

Given the presented sustainability context in logistics, the following scenario briefly summarizes the challenges along with some specific business process steps. Therefore, we introduce the role of a business process orchestrator being in charge to evaluate logistic partners and to judge them upon their performance. In this context, a typical regular business process for the user follows four phases:

KNOW the knowledge worker would like to review several performance criteria of performed shipments in a time interval. Particularly, eco-balance related criteria such as Co₂ emissions and energy consumption per shipment are important. Therefore, various information is required that needs to be contextualized and presented to the user. While some of the information (e.g. fuel consumption, Co₂ emissions, etc.) can only be retrieved from transport vehicles a mapping to each shipment is required for shipment-based information presentation. As the user needs to rely on the information, the mapping and aggregation of the raw data requires high quality event capturing and processing.

ANALYZE assuming a recent regulation change, the shipments' eco-balances need to be analysed if they still adhere to the given limits. Likewise, in

case of deviations beyond thresholds, a drill-down to the root causes should be enabled. Therefore, the information on the shipment level needs to be traced to the involved transport vehicles. For instance, a root cause might be due to traffic jams or the extremely high Co2 emissions of a logistic partner's truck that has a damaged particle filter.

DECIDE Based on the analysis, conclusions can be derived. For instance, the knowledge worker can inform the owner of the identified vehicle and organize a replacement. Such decisions have an impact on logistic parameters such as capacity planning. Hence, decision support is required. For instance, different capacity scenarios can be simulated and best practice information can be used to benefit from similar previous cases.

ACT Assuming the knowledge worker follows the recommendation of the decision support tool to get in touch with the vehicle owner, the knowledge worker requires the contact details of the vehicle owner and some briefing material. The following replacement process needs to be tracked until the issue is solved. Then, the issue should not show-up anymore in future analyses.

LEARN the performed steps during the previous phases can be used to support similar occurrences in the future. Therefore, the lessons learned from this scenario need to be stored for future reference. For instance, the actual analysis steps and their relation to the final decision represent valuable information. Finally, if some cases occur on a continuous basis, business and event rules can be defined.

This scenario reveals an open gap to deliver business value in terms of sustainability. The eco-balance on a shipment level is one key benefit that offers opportunities in terms of cost savings, marketing, and compliance. While the eco-balance provides a business value on its own, it also contributes to eco-driven analyses. If defined eco-balance targets are not achieved, logistic processes and resources can be analyzed. For instance, deviations over time with respect to certain transport vehicles can be used as triggers for maintenance or modernization. The various analysis possibilities, finally, enable an eco-driven route-map planning, capacity utilization, and partner selection.

3. Underlying Technologies & Architecture

In this paper, we particularly focus on the event-based processing and the information-based user-process interaction. While these two functionalities are core components to achieve real-world integration to the benefit of sustainable logistics, further components and technologies are required to leverage the full potential of the concepts. As a reference, this section presents underlying

technologies and projects that build the foundation for our presented concept. Finally, we introduce the overall architecture design.

3.1 Underlying Technologies

Our concept relies on the foundation for the Internet of the Future -namely the Internet of Things and the Internet of Services. Both technologies provide tremendous potential for applications in various business areas and form the basis of transparent business networks. The rationale behind is that quicker, better, and more accurate information leads to better business decisions and, finally, increased competitiveness. While, real-world integration represents the original information source for competitive strengths, the continuing trend towards business networks is the organizational frame that needs to be integrated into information concepts. Companies more and more concentrate on their core competencies and engage in mutually beneficial partnerships with other companies.

In addition to the IoT, the IoS is a core infrastructure supporting our concept. The IoS extends today's Internet to become service-enabled. By that, it allows enterprises to focus on core competencies and to reach out to a global market. It uses semantic technologies that understand the meaning of information and facilitate the accessibility of content. Thus, data from various sources and different formats can easily be combined and processed. These services enable the dynamic creation of business networks while guaranteeing transparency on relevant information such as GRC compliance and costs. The innovative technological developments within the IoS drive the creation of new delivery channels for services and entirely new business models. The service creation is facilitated by an open platform and interface architecture as provided by the Enterprise Service-Oriented Architecture (ESOA). The IoS takes the ESOA approach to the next level by simplifying the design, provision, management, and the consumption of services.

Moreover, our concept is based on underlying technologies of related research activities. Within the IKT-2020 research program of the German Federal Ministry of Education and Research, the Innovation Alliance 'Digital Product Memory' (DPM) develops key technologies for the Internet of Things and Service in the cooperative projects SemProM, Aletheia, and ADiWa.

SemProM1 follows an item-oriented view. It increases object intelligence by collecting and rep-representing information of the lifecycle of an object and communicates with other objects or systems. Typical 'intelligent' objects in this context are products, goods, transport resources or devices. In the context of this

paper, the SemProM concepts and solutions are used for the integration of event sources – devices, transportation units, intelligent products – into the ADiWa event bus.

Aletheia2 provides a holistic view on the object. It collects all information regarding a specific item or item class (set of items with equal attributes) and presents this information to a knowledge worker in a context specific manner. So, the Aletheia results are used in our concept as knowledge platform for enabling the right business decisions.

ADiWa3, finally, has a process-oriented view on the intelligent objects. It focuses on the usage of the collected real-world information inside business processes. So, it combines the underlying technologies and enhances them with a business process management perspective. The respective architecture is also the foundation of our presented concepts.

The alliance's goal is to enable a future 'Digital Company' where business processes are based on the capabilities of the Future Internet. Key requirements for successful companies are transparency, causality, productivity, agility, and adaptively.

3.2 Architecture

The vision of ADiWa is to model, optimize, and control dynamic business processes with information of the physical world. Essential building blocks are the usage of all process-relevant information out of the IoT/IoS and the flexible composition as well as the dynamic and variable adaption of business processes. So, the two main entities used and linked are events and processes. Both subsets represent the core paradigm of ADiWa: event-driven business process management (ED-BPM). The goal is to enable information workers to know, analyze, decide, act, and learn (see also Section 2) to the benefit of optimized business processes. Referring to the initial scenario this interaction schema generically implies:

KNOW: Transparency – Transform, in near real-time, raw data in form of events into meaningful business insights

ANALYZE: Causality – Analyze the information in a specific business context, trace quality, and reveal root causes

DECIDE: Productivity – Simulate alternatives and make effective decisions to increase business performance

ACT: Agility – Implement actions within collaborative business networks and facilitates seamless communication

LEARN: Adaptively – Continuous assessment of lessons learned to derive

best practices and pro-active rulings

The ADiWa architecture (Schief et al., 2011) supports the information flow of the presented logistics scenario by different technology modules, which are integrated to a technology & application platform (see Figure 1). The main innovation potential is inside these modules as well as in the technical and functional combination of advanced technologies that are used in specific application domains. Core elements of the architecture are:

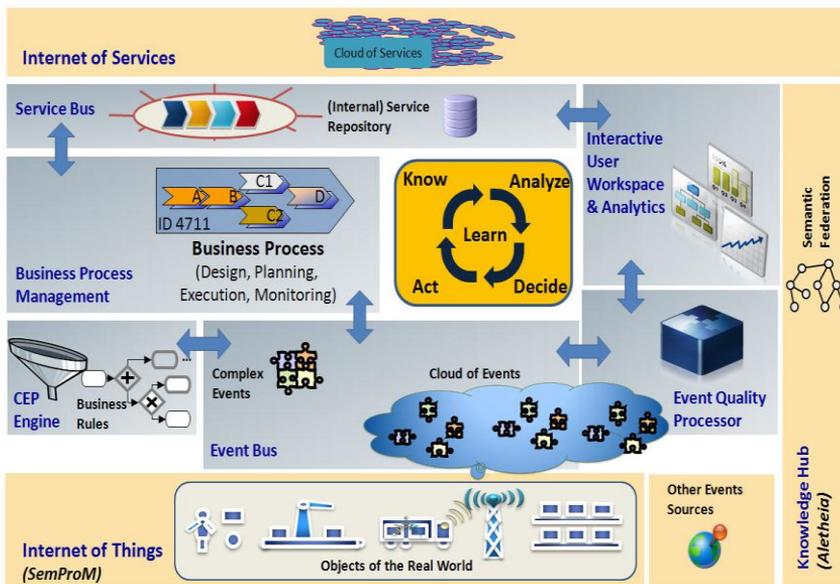


Figure 1: High-Level Architecture

Event Bus – the central information push backbone with dispatching mechanism and a Complex Event Processing (CEP) Engine. The Event Bus is the input channel of all IoT-based real-world events (e.g. Co2 emissions) into the platform. A special generalized event schema is offered in order to connect arbitrary event sources to the event bus.

Service Bus – the information pull backbone with an internal service repository. The service bus connects the different modules inside the platform and is used for process orchestration including external service calls via the IoS (e.g. GRC regulations).

Business Process Engine (BPE) – event-enabled modeling and execution environment for business processes. The BPE component is the central component for design-and run-time of processes (e.g. route and capacity planning for a certain shipment). It also contains a planning & simulation component as well as an event-enabled process monitoring.

Interactive User Workspace & Analytics – flexible dashboard approach and

Business Intelligence for overview and decision support within the specific application context. It also offers a list of recommended pro-active and re-active actions (e.g. capacity re-planning for previous shipments) to achieve the business goals.

Event quality processor-online validates for event-driven information. The data events are continuously processed and analyzed in “near real-time”, in order to provide quality and trace-ability (e.g. root causes for rise in Co2 emission) for fast and successful business decisions.

Within this paper, we focus on a deeper discussion of essential aspects of this architecture, basically referring to the last two mentioned components. The basic assumption is that we get the required real-world events from the IoT into our system. Then, we present concepts to transform them into valuable information for business decisions. Core requirement for reasonable usage for managing business processes is reliability and quality of events. These properties are not self-evident, but need to be ensured by appropriate concepts and need to be transparent for the business user.

4. Event Processing and Quality

In this section, we start our deeper discussion with the aspects of the event processing and quality. The main targets are (I) to make effective use of the (potential) huge amount of raw events, and thus, to create knowledge and transparency instead of getting lost in an information overload, (II) to reveal quality and tracing information to assess the value of information and find root causes, and (III) to give decision makers a head start of information by reusing the information required for the tracing.

4.1 Filtering and Processing

The first (and more common) step towards the creation of transparency (KNOW) is the utilization of an Event Bus together with a CEP Engine. The Event Bus mainly filters and routes the events from event producers to event consumers. In our solution, the communication between these two parties follows the publish/subscribe paradigm; this allows a decoupled many-to-many communication. Moreover, we realize our Event Bus as a distributed message oriented middleware (similar to JMS (Sun Microsystems Inc., 2002)). Hence, the event transportation is both robust and scalable.

Clearly, routing and filtering alone will not lead to sophisticated and business-relevant information; we need to process and aggregate low-level IoT events to high-level business events (Schmidt and Schief, 2010). Thus, we

integrate a (rule-based) CEP Engine that allows both the complex multi-step processing of events into business events and complex structured events during the processing and as the result of the processing. Again, for scalability reasons our CEP Engine is realized as a distributed system.

4.2 Quality of Events

As mentioned above, filtering and processing events is important but it is only the first step for a sophisticated transparency of the actual business processes; it provides us the relevant information and gives us deep insight. However, we also have to take the quality of the information into account. As shown in (Leonardi et al., 2009), data in the logistics area is often incomplete and inconsistent. As a consequence, decision may be taken based on wrong or incomplete information. Hence, what we need once more is transparency.

The quality of data is a well addressed topic in the database community, especially in the field of information integration. Several approaches identify and categorize relevant dimensions (Wang and Strong, 1996) or propose methods for the management of quality information (Storey and Wang, 1998). All these solutions assume persistent and rather static datasets. Although there are first considerations of the quality of (flat) data stream items (Klein and Lehner, 2009), both the measurement and the efficient representation of the quality of more complex (especially hierarchical) items have not been sufficiently addressed so far.

We tackle this problem with a generic approach to handle quality of events; implementations of this approach may select quality criteria of interest. The basic idea is to track several quality criteria for each event or even for each attribute of an event. So far, we focus on timeliness, completeness, data volume, confidence, and accuracy of information provided by the events. With this quality information at hand, a decision maker can make the following exemplary considerations:

- Get the age of the current event or choose only events that are not older than 1 hour.
- Get the completeness of the current event or choose only events where at least 80% of the expected sub-events contributed to. Obviously, this requires fix event rates.
- Get the number of contributing events for the current event or choose only events containing average values of at least 10 measurements.
- Get the width of the 95% confidence interval of the current value or choose only events where the probability that the value deviates by at

most 5% is at least 95%.

- Get the accuracy of the current value or choose only events of event producers that guarantee an error of at most 10%.

We categorize these quality criteria based on their steadiness; this also reflects the way of how to determine and manage the respective quality information.

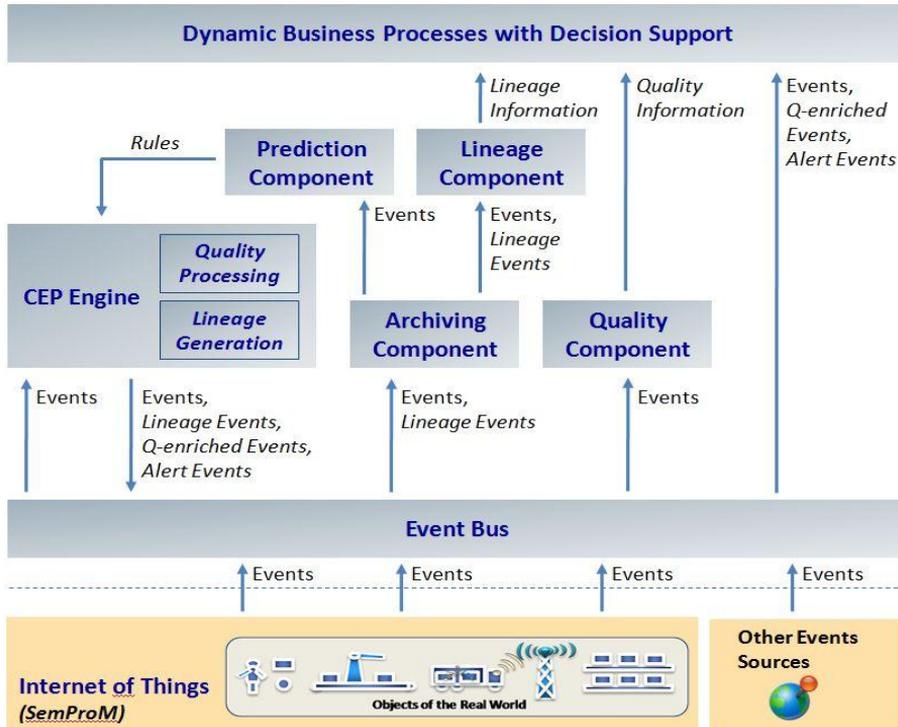


Figure 2: Architecture Detail for Event and Quality Processing

4.2.1 Dynamic Criteria

Timeliness, completeness and data volume are specific to the individual events and, thus, are highly dynamic. The timeliness is given by the timestamp that is part of each event; for the completeness and the data volume our CEP Engine annotates composite events with the number of contributing events.

4.2.2 Roughly Static Criteria

Confidence is considered roughly static; it depends on the variance of the information among the events. Our quality component subscribes to the information of interest and incrementally maintains its variance.

4.2.3 Static Criteria

Accuracy and the event rate (needed for the completeness) are considered static as they are defined by the event producers. This information is part of the

advertisement events sent by the event producers to register for the event-based system. Again, our quality component subscribes to the Event Bus and tracks the quality of the event producers.

Figure 2 gives an architectural overview of our solution; it is a more detailed view on the lower part of Figure 1. At the bottom, there are the objects of the real world and other event sources sending events to the Event Bus. The CEP Engine subscribes for events and processes the received events to complex events based on the registered rules. As stated above, we extend the CEP Engine for quality processing (see Quality Processing extension in the CEP engine in Figure 2); here, events are annotated with the number of contributing events and the (dynamic) quality information of these contributing events is aggregated. The results of this quality processing extension are quality-enriched events (Q-enriched Events in the figure). Additionally, our Quality Component subscribes to the Event Bus and tracks the static and the roughly static quality information. Dynamic business processes and decision support systems can then be provided with quality information.

Clearly, this quality information significantly enhances decision support; now, users can better assess the value of the provided information. Consequently, it is another fragment in the KNOW phase. Additionally, quality information provides valuable information in the ANALYZE step, e.g. by finding reasons for wrong decisions or unexpected values.

Moreover, the disclosure of the quality information additionally gives us the opportunity to optimize the production, the transport, and the processing of events. Users can identify weak points in this chain – event producers of low quality, slow event transportation or inaccurate event processing – and react appropriately to resolve this issue.

4.3 Lineage Information

To improve the support of the decision maker we go one step further: We additionally provide lineage information of composite events. This information allows users to reveal root causes in the ANALYZE phase.

As for the data quality, there exists a lot of research in the database area; here, the focus is the determination of the data provenance in data warehouse databases (Benjelloun et al., 2008; Bhagwat et al., 2004; Cui and Widom, 2003). But again, these solutions assume persistent and rather static datasets. Moreover, most of the solutions focus on relational operators.

In our solution, we add another function to the CEP Engine, the Lineage Generation (see Figure 2). This extension generates events solely containing

lineage information. For each (relevant) newly created composite event a lineage event containing the processing node, the contributing events, the applied rule, and the resulting event is generated and published; the contributing events are identified via uniquely as-signed IDs. Additionally, we employ an Archiving Component that subscribes to the lineage events and also stores the contributing events, see Figure 2. Finally, the Lineage Component can (relatively straightforward) compute the lineage information out of the information from the Archiving Component and provide it to the interested user or application.

4.4 Event Prediction

Having the Archiving Component in place, we can exploit the benefits of the stored information for another valuable feature: We offer the feature of the event prediction, i.e. & the pro-active detection of (probably) occurring complex events. The general idea of the event prediction is to send an alert event if the probability of the occurrence of a composite event exceeds a given threshold. This great feature gives users an head start of information (KNOW) and allows them to dynamically initiate appropriate actions at an early stage (DECIDE and ACT).

We realize the event prediction as follows: At first, the user specifies the complex events of interest together with the desired threshold for the notification. Based on this, our Prediction Component (see Figure 2) analyzes the events stored in the Archiving Component. It considers sub-events (which may also be complex) and computes the frequencies of how often these sub-events lead to the relevant complex events; here, also the time between combined sub-events is considered. Next, the Prediction Component selects all the sub-events whose frequency (or probability) of leading to a complex event of interest exceeds the specified threshold. Finally, new rules are created for those sub-events and registered at the CEP Engine. Note, with this creation and registration of rules the event prediction also contributes to the LEARN phase. Having the rules integrated, the CEP Engine generates alert events if one of the identified sub-events occurs, and so pro-actively notifies the probable occurrence of a complex event to the user.

5. User-Process Interaction

After the discussion of the event and quality processing aspects of our approach, we now focus on our second main topic concerning the interactive user workspace and analytics.

5.1 Supporting Knowledge-Intensive Tasks

The excitement about the process optimization capabilities provided by conventional business process management systems has been seized through the increasing awareness that business processes are not limited to operational activities but involve also ad-hoc, knowledge-intensive tasks, which cannot be automated in a rigid manner (Riss et al., 2005).

Knowledge-intensive tasks are a blind spot for business process management systems, as these tasks are executed in an unsupervised, highly individual manner. Hence, individual experience is not disseminated and task execution largely depends on implicit business domain expertise.

For enabling efficient user-process interaction we propose a framework, realizing: (I) situation-specific and personalized task execution support through context capturing and proactive user assistance; (II) generalization of individual task execution activities on enterprise level towards establishing best-practices. As a core concept we suggest activity schemes: a structure capturing a probabilistic task execution model. The framework and the key concept of activity schemes are discussed in detail in (Schmidt et al., 2010). In the following we outline the major concepts for enabling efficient knowledge work support in the context of agile business processes in the logistics domain through increased process transparency, information reuse, and collaboration support.

5.2 Abstraction through Taxonomies

A great deal of the complexity in logistic processes emerges through the outsourcing and involvement of multiple service providers. While all stakeholders have the same ultimate goal (to deliver the merchandise) each of them is focused only on a subset of the overall process chain and requires only part of the overall information. The partners may come to an agreement and implement IoT-enabled processes where information is captured at an advanced detail level. However, even then different users may use different tools to view and analyze the information. Moreover, the involved users will have different information demands -e.g. a logistic service provider employee at an export hub who plans the packaging of a shipment may want to see the available room in a transportation vehicle, whereas a business process orchestrator (as introduced in Section 2) may need to see additionally warehouse energy consumption at intermediate locations to estimate the overall process costs and ecological footprint in various alternative cases. Referring to the introductory scenario (cf. Section 2), the participants in the overall logistic process need to work in

different contexts and to know different facets of the overall process information. What it all boils down to, is that a certain degree of abstraction is necessary, where available information can be interconnected into logical constructs in a loosely coupled manner and filled with process instance data on-demand and according to the individual information needs and business domain expertise. Our framework enables this loose coupling and on-demand instantiation through a set of taxonomies, which are realized as OWL ontology (Schmidt et al., 2011):

- Activity taxonomy: A hierarchy of verbs standing for interactive activities (relates to user-system interaction)
- Resource ontology: Classified resources which occur in knowledge work and their relations
- Application taxonomy: Classes of software applications

Connections between these taxonomies describe which activities can be performed on which resources by which applications e.g. create contract can be performed by using a template document as a resource in a word processing application.

5.3 Knowledge Actions

Referring to the scenario from Section 2, the ANAL-YSIS performed by a business process orchestrator is generally considered as knowledge-intensive work. Apart from the drill-down into fine-grained event information (cf. Section 4.3), we aim at supporting this analysis in a more generic way, by capturing user context and intent and providing proactive user assistance based on the discussed taxonomies. For formalizing the connections between the taxonomies we follow the knowledge work description of (Reich, 1992) and (Pyoeriae, 2005) as a semiotic activity of information creation, transformation, and consumption, often executed using computer desktop applications, which deliver a restricted set of functionalities. For this understanding of knowledge work, (Hadrich, 2008) identifies a set of eight reoccurring knowledge work actions as fragments of work. Each knowledge action is an abstraction from the actual execution process. These knowledge actions are: authoring, coauthoring, training, acquisition, update, feedback, expert search, and invitation. A knowledge action is independent from the specific task context, e.g. authoring always involves the creation or transformation of symbols. We suggest that knowledge-intensive task execution can be described by means of knowledge actions where a knowledge action is considered as a sequence of interactive activities on resources using applications.

5.4 Activity Schemes

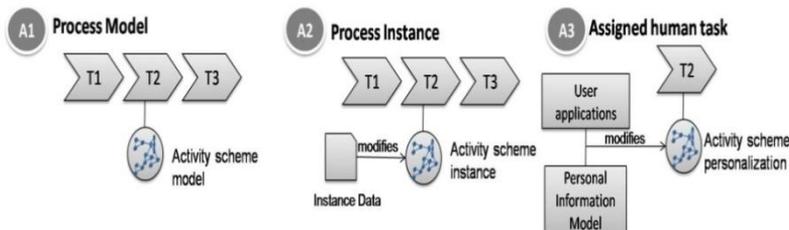
Again referring to the scenario from Section 2, the DECIDE step raises similar challenges as the ANAL-YSIS step, i.e. the user needs proactive assistance, where relevant information and possible further actions are suggested based on the current user context. On the other hand, LEARNING needs to be supported through dissemination and exchange of task execution knowledge on enterprise level. To embed comprehensive user support through personalization and generalization of ad-hoc task execution knowledge into business processes, we further introduce the concept of activity schemes.

An activity scheme is a directed graph of nodes which stand for knowledge actions. Knowledge actions can be included several times in a single activity scheme, e.g. if information search for different topics is necessary. Activity schemes are formal constructs, which can be attached to formal human task definitions in business process models in order to provide support for context elicitation and proactive user assistance for these tasks. For example, a logistic service provider employee can be delegated a task to coordinate the repackaging of a shipment at an export hub. For this purpose the employee may need to check the specification of the provided transportation vehicle to calculate the optimal package dimensions and to potentially include additional stock in the shipment in order to optimize other shipments. In this case, the activity scheme approach would not require setting strict parameters for the task. The task may be indeed performed by various service providers in various application environments and by users with different skills and working preferences. Instead, an abstracted activity scheme will be attached to the human task “repackage shipment” in the overall logistic process model. This is displayed in step A1 in Figure 3 where T1 to T3 are task nodes in a business process model. The activity scheme incorporates possible knowledge actions (acquisition, analyze, communicate), linked to relevant resources (transportation vehicle specification documents, legal documents) and respective applications (Web Pages, Excel, Calculator) as defined in the provided taxonomies. The latter can be extended iteratively in a domain specific manner. In step A2 (cf. Figure 3), the human task is delegated as part of the business process execution and an activity scheme is instantiated and enriched with process instance information. In step A3, the activity scheme is further enriched according to the actual user work context and personal information model, described through the taxonomies and captured in a local semantic storage. Once activated, the scheme provides proactive user support by suggesting relevant resources,

applications, persons and possible related actions. This is the personalization process.

After the employee completes the delegated human task within the process instance, the personal in:

Activity Scheme Personalization



Activity Scheme Generalization

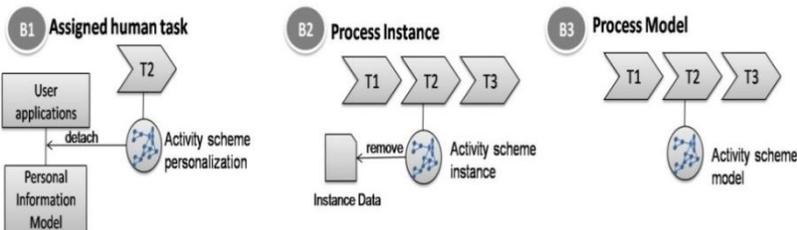


Figure 3: Activity Schemes Personalization and Generalization

Formation is detached from the activity scheme (cf. step B1 in Figure 3) by replacing personal information model entities with the appropriate taxonomy abstractions (e.g. MS Word with Word Processing Application, Truck/Trailer Specification Document with Means of Transport Specification etc.). Then in step B2, process instance data is removed from the activity scheme. The generalized scheme is finally associated with the task model in the business process model (step B3) and can be reused later on for best-practice recommendation and assistance in further process instances.

6. Outlook

In this paper, we provide a brief overview of the motivation, progress, and goals of our project. In line with the overall concept, the two main functionalities (Sections 4 and 5) deliver a holistic solution framework that covers the various challenges in the field of IoT-process integration to the benefit of sustainable logistics. The described event-processing and -quality functionalities increase the needed transparency, the value of information, and the possibility to track and trace root causes for analysis purposes. The user-process interaction tools

improve the flexibility, analysis, and decision making of knowledge workers. All in all, the business process performance can be enhanced, significantly.

Most importantly, we aim to close the gap between technical possibilities and their business usage and, thus, to demonstrate the business value of our overall concept. In this respect a deep and thorough evaluation of the concept and the functionalities by means of the presented logistics scenario, is our benchmark.

For synergy and consistency reasons, we are basing our work on well-established standards in related research areas. Leveraged by our research activities, we intend to raise the research field of Internet of Things and sustainability one step higher and, hence, to provide a fertile ground for future research projects.

Beyond our concept, we see a need for further research in various areas. As far as concerns the logistics domain, today's challenges need to be further incorporated into an IoT-based concept for sustainable logistics. For instance, the increasing global network in line with the internationalization and the collaboration among a multitude of heterogeneous partners needs to be addressed. In that respect, process supporting collaboration hubs are a promising research field. In addition, common data standards are one central prerequisite that must be achieved. For instance, collaborative data feeds can be addressed by advanced Business Intelligence capabilities. Likewise, IT-support for multi-modality along with a mobile integration into the concept is a direction that can leverage the business benefits. Moreover, the addressed topics need to be extended to further areas in the value chain such as production (cf. (Ameling et al., 2010)) in order to realize an end-to-end eco-balance. Sustainable production based on the IoT can significantly benefit from the presented concepts for logistics. In addition, further synergies can be realized base on the interplay of homogeneous concepts in production and logistics.

Further to the technological research fields, we see open gaps with respect to research in business fields. The shown possibilities, in terms of information usage for decision making, need to be incorporated into the business domain. For instance, an enhanced cost, revenue, and risk model that comprises IoT-based sustainability information has to be defined. By that, an organization's incentive structures can reflect the respective goals. Finally, compliance and security topics need to be addressed for any changes that are triggered in business processes. Extended flexibility and real-world response may not result in any compliance violations such as stated by the Sarbanes Oxley Act.

Acknowledgements

The project was funded by means of the German Federal Ministry of Education and Research under the promotional reference 01IA08006. The authors take the responsibility for the contents.

References

- Ameling, M., Schief, M., Nietzold, F., & Kuhn, C. (2010). Demand-side energy management systems for manufacturing. *In GI Jahrestagung*, 1, 467–472.
- Baumoel, U., & Winter, R. (2001). Intentions Value Network. *ICEIS* (2), 1075–1080.
- Benjelloun, O., Das Sarma, A., Halevy, A., Theobald, M., & Widom, J. (2008). Databases with uncertainty and lineage. *The VLDB Journal*, 17(2):243–264.
- Bhagwat, D., Chiticariu, L., Tan, W. C., & Vijayvargiya, G. (2005). An Annotation Management System for Relational Databases. *In VLDB'04: Proceedings of the Thirtieth international conference on Very large data bases*, 900–911.
- PAS 2050: Specification for the assessment of the life cycle greenhouse gas emission of goods and services. *BSI*, 2008.
- Cui, Y., & Widom, J. (2003). Lineage Tracing for General Data Warehouse Transformations. *The VLDB Journal*, 12(1):41–58.
- Hädrich, T. (2007). Situation-oriented Provision of Knowledge Services. In *Information Systems*. University of Halle-Wittenberg, Halle (Saale) .
- Kaplan, R., Norton, D., & Rugelsjoen, B. (2010). Managing Alliances with the Balanced Scorecard. *Harvard Business Review January–February*, pp. 114–120.
- Klein, A., & Lehner, W. (2009). Representing Data Quality in Sensor Data Streaming Environments. *J. Data and Information Quality*, 1(2):1–28.
- Leonardi, J., Browne, M., & Allen, J. etc. *Data Needs and Data Review for Green Logistics Research*. University of Westminster, Department of Transport Studies.

Maass, W., & A. Filler. (2007). Tip'n Tell: Product-Centered Mobile Reasoning Support for Tangible Shopping. *Proceedings of MSWFB 2007: Making Semantics Work for Business, part of 1st European Semantic Technology Conference*, Vienna, Austria.

Pyoeriae, P. (1997). The concept of knowledge work revisited. *Journal of Knowledge Management*, 9(3):116–127.

Reich, R. B. (1992). *The work of nations -Preparing Ourselves for 21st Century Capitalism*, Vintage.

Riss, U., Rickayzen, A., Maus, H., & van der Aalst, W. M. P. (2005). Challenges for Business Process and Task Management. *Journal of Universal Knowledge Management*, 0(2):77–100.

Schief, M., Kuhn, C., Zimmermann, B., Roesch, P., & Waterfeld, W. etc. (2011). The ADiWa Project -on the Way to Just-in-time Process Dynamics based on Events from the Internet of Things. In *Proc. ICEIS (4)*, pp.371-377.

Schmidt, B., Paulheim, H., Stoitsev, T., & Mäuser, M. uhlh Towards a Formalization of Individual Work Execution at Computer Workplaces. In Forthcoming: Lecture Notes in Artificial Intelligence.

Schmidt, B., & Schief, M. (2010). Towards Agile Business Processes Based on the Internet of Things. *Advanced Manufacturing and Sustainable Logistics*.

Schmidt, B., Stoitsev, T., & Muehlhaeuser, M. (2010). Activity-Centric Support for Weakly-Structured Business Processes. In *Proceedings of the 2010 ACM SIGCHI Symposium on Engineering Interactive Computing Systems*.

Simons, D., Mason, R., & Cardiff, U. (2003). Lean and green: doing more with less. *ECR Journal*, 3(1):84–91.

Storey, V. C., & Wang, R. Y. (1998). Modelling Quality Requirements in Conceptual Database Design. In Conf. on *Information Quality*, pages 64–87.

Sun Microsystems Inc. Java(TM) Message Service Specification Final Release 1.1. Sun.

Wang, R. Y., & Strong, D. M. (1996). Beyond Accuracy: What Data Quality Means to Data Consumers. *Journal of Management Information Systems*, 12(4), 5–33.