

A Real-Time Architecture for Proactive Decision Making in Manufacturing Enterprises

Alexandros Bousdekis¹, Nikos Papageorgiou¹, Babis Magoutas¹,
Dimitris Apostolou²(✉), and Gregoris Mentzas¹

¹ Information Management Unit, National Technical University of Athens,
9 Iroon Polytechniou Str. 157 80 Zografou, Athens, Greece
{albous, npapag, elbabmag, gmentzas}@mail.ntua.gr

² Department of Informatics, University of Piraeus, 80 Karaoli & Dimitriou Str.
185 34 Piraeus, Greece
dapost@unipi.gr

Abstract. We outline a new architecture for supporting proactive decision making in manufacturing enterprises. We argue that event monitoring and data processing technologies can be coupled with decision methods effectively providing capabilities for proactive decision-making. We present the main conceptual blocks of the architecture and their role in the realization of the proactive enterprise. We illustrate how the proposed architecture supports decision-making ahead of time on the basis of real-time observations and anticipation of future undesired events by presenting a practical condition-based maintenance scenario in the oil and gas industry. The presented approach provides the technological foundation and can be taken as a blueprint for the further development of a reference architecture for proactive applications.

Keywords: Proactivity · Decision-making · Event-driven computing · Condition-based maintenance

1 Introduction and Motivation

The emergence of the Internet of Things paves the way for enhancing the monitoring capabilities of enterprises by means of extensive use of physical and virtual sensors generating a multitude of data. The sensing enterprise concept can influence a wide range of industries. For example, manufacturing companies can utilize sensors to enable the identification of deviations from production plans as soon as they appear; logistics networks can identify delays about in the delivery time in real-time through sensor-generated events. The main driving concept in sensing enterprises is the use of multi-dimensional data captured through physical and virtual sensors generating events and providing added value information that enhances context awareness. Consequently, the large amount of data generated by sensors leads to a strong demand for data-driven, real-time systems capable of efficiently processing data in order to get meaningful insights about potential problems.

Event monitoring and data processing accompanied with enabling real-time systems are essential for managing problems in complex, dynamic systems. Advanced

monitoring capabilities should provide the basis for a new level of sensing performance that not only observes current problems, but also senses that the problem might appear, that is, by focusing on a proactive approach. Indeed, observing a delay is very useful information, but anticipating that there will be a delay is even more important from the business point of view. The capability to anticipate leads to the possibility to decide and act ahead of time, i.e., to be proactive in resolving problems before they appear or realizing opportunities before they become evident and be able to recover and support continuity.

Proactive, event-driven decision-making has been recently introduced in the literature as a conceptual model for deciding ahead of time about the optimal action and the optimal time for its implementation [1]. A proactive enterprise decision support architecture should integrate different sensor data, provide large-scale and real-time processing of sensor data and combine historical and domain knowledge with current data streams in order to facilitate proactive decision-making. In this paper we focus on proactive decision making in the manufacturing domain where the challenges associated with the provision of decision support based on predictions become significant, especially when dealing with maintenance where several factors should be considered such as costs of maintenance actions as a function of time, safety issues and degradation of equipment.

Despite the plethora of existing works for and prognosis in maintenance, most of them do not examine the integration with real-time, data-processing platforms and the automation of decisions by providing recommendations for maintenance actions, while the supported level of proactivity is typically low [2], [3]. Further, there is no support for switching easily between available decision methods or selecting a preferred method among the available ones since decision methods may address different challenges in terms of the availability of data and domain knowledge.

The integration of various decision methods in a real-time platform that would allow users to select appropriate methods based on the available data and the desired proactive decision support is the research objective for our work. Technically, the challenge is to develop a real-time architecture that would support the development of decision support applications enabling the business analyst to select decision methods and configure them so that they are operable for the problem at hand. In this paper we present such an architecture for decision making that enables the transition from sensing to proactive enterprise.

The rest of the paper is organized as follows. Section 2 discusses enabling technologies and works related to real-time architectures for enterprise decision making with an emphasis on supporting proactivity. Section 3 outlines the proposed architecture for realizing proactive enterprise decision making. Section 4 presents a scenario in which proactive decision making in maintenance is enabled with the proposed architecture. Section 5 discusses the main findings of our work and our future plans.

2 Enabling Technologies and Related Work

In the context of the sensing enterprise, physical and virtual sensing devices such as sensors, actuators and controllers can detect state changes of objects or conditions and create events, which can then be processed by a system or service. From the point of view of communication, the use of a web-service communication paradigm allows

sensors to be easily integrated into a complex architecture. To this end, the Service-Oriented Architecture (SOA) paradigm strongly contributes to the development of monitoring and control infrastructures, enabling interconnectivity at an object level. Moreover, the Event Driven Architecture (EDA) provides an architectural computing paradigm that has the ability to react to changes by processing events [4], [5]. EDA can complement SOA because services can be activated by triggers fired on incoming events[6], [7]. Building on EDA, proactive event-driven computing is a new paradigm where a decision is neither made due to explicit requests nor as a response to events, but is triggered by real-time predictions of an event. Therefore, the decisions are taken under time constraints and require the exploitation of large amounts of historical and streaming data [6-8].

In the manufacturing domain, sensors have the capability of measuring a multitude of parameters frequently and collecting plenty of data. Analysis of Big Data, both historical and real-time, can facilitate predictions on the basis of which proactive maintenance decision making can be performed. The e-maintenance concept can significantly address these challenges [9-11]. DYNAMITE ‘Dynamic Decisions in Maintenance’ research project has examined e-maintenance [12], [13]. It developed the TELMA ‘TELeMAintenance platform’ which provides intelligent agents directly implemented at the shopfloor level into the PLCs and decision-making services in front of the degraded situation process performance, including assessment of the degraded process performance, prognostic of the future situation and decision to be taken to control the process in its optimal performance state. The WelCOM project developed an e-maintenance architecture exploiting the following key relevant technological factors: web-based maintenance services, wireless sensing and identification technologies, data and services integration and interoperability, as well as mobile and contextualized computing [14]. Within such a framework, the authors proposed a layered e-maintenance architecture, leveraging upon the strengths of smart and wireless components in order to upgrade the maintenance-services from the low level of operations to the higher levels of planning and decision making. For an overview of other e-maintenance platforms, both from academia and industry, please refer to [12], [13], [15], [16] and [17].

E-maintenance can be leveraged with EDA and proactive event-driven computing in order to enable proactive decision making about optimal maintenance actions and the optimal time for their implementation. To do this, e-maintenance should be extended in order to handle real-time, data-driven predictions and coupling them with domain knowledge and decision methods. Several decision methods, ranging from operational research to machine learning and statistical ones, have been proposed in the literature in support of proactive maintenance decision making. Nevertheless, decision methods have not been integrated in e-maintenance platforms yet and are rarely validated in an industrial environment.

3 A Conceptual Architecture for Decision Making in the Proactive Manufacturing Enterprise

In this section we outline the proposed architecture for proactive enterprise decision-making, its main conceptual blocks as well as the main functionality implemented by

each block. The role of the architecture depicted in Figure 1 is two-fold. *Configuration role*: to allow business analysts create decision method instances addressing anticipated problems and configure them by adding, removing or changing possible mitigating actions as well as other domain knowledge required by the underlying decision methods. Domain knowledge can include the list of alternative actions, their costs, their delays (corresponding to the time period from its implementation until it starts taking effect), the time-to-undesired event after their implementation as well as the next planned maintenance. *Processing role*: to support decision-making ahead of time on the basis of real-time observations and anticipation of future undesired events, by coupling decision methods to a real-time processing environment. The proposed architecture consists of a user interaction and a real-time processing layer, along with a *data layer* which houses a relational database engine where all information needed by the two other layers is stored and retrieved.




Fig. 1. Proactive Decision Making Architecture

The *user interaction layer* occupies the top level of the architecture and includes a web-based application that supports the configuration of the architecture, by allowing business analysts to select the most appropriate decision method for mitigating predicted

undesired events, on the basis of functional and non-functional requirements, as well as to embed domain knowledge with the aim to define and configure the various parameters of the decision method instances. Examples of decision methods incorporated in the aforementioned decision configurator dashboard include Markov Decision Process (MDP) [1] and Cost Optimization [18], while more details are provided in [3]. Decision method instances are specific instances of decision methods, corresponding to specific equipment or other subject of a predicted undesired event. As can be seen in Figure 1, another functionality exposed by the decision configurator to the business analyst is related to the visualization of feedback received by the recipients of the recommendation with respect to the implementation of the recommended actions. This feedback aims to support business analysts in the process of refining the recommendation generation process on the basis of such feedback.

The *real-time processing layer* fulfills the processing role of the proposed proactive decision making architecture and is based on the Observe, Orient, Decide, Act (OODA) model of situational awareness [19]. This model sees decision-making occurring in a recurring cycle of unfolding interaction with the environment, oriented via cues inherent in tradition, experience and analysis. These cues inform hypotheses about the current and emerging situation that, in turn, drive actions that test hypotheses. The real-time processing layer deals with the continuous processing of sensor data by applying OODA principles and subsequent sending of notifications/recommendations to relevant people or systems. The sensor data are collected from sensors and they are injected into the real-time processing layer, where they are handled by the OODA information-processing pipeline as follows:

- (Observe) The *sensing service*, deals with data acquisition, transformation (including cleaning) and publishing. It is responsible for sensing relevant sources and transforming data in a format useful for further analysis. The sensing service is followed by the *data enrichment service* that enables semantic enrichment of real-time streams (events) with background knowledge.
- (Orient) The orient phase includes services for anticipation management, which enable the generation of real-time, data-driven predictions of future undesired events through the *predictive analytics service*. Predictions are triggered on the basis of unusual situations discovered on the basis of complex enriched events identified by the *Complex Event Processing (CEP) service*.
- (Decide) The decide phase includes services enabling anticipation-driven decision-making, in the sense that the predictions of undesired events generated by the services of the Orient phase are taken into account. More specifically, based on a “prediction” event, which predicts the probability distribution for the occurrence of a future undesired event as a function of time, the *context-aware decision management service* generates proactive recommendations of actions that mitigate or eliminate this event along with the recommended activation time. The recommended actions and action activation times are calculated by enacting the decision method instances that are defined and configured through the decision configurator. The generated proactive recommendations are further propagated within the OODA information-processing pipeline through the *publish recommendation event service*, until they reach the relevant enterprise

stakeholders. Prediction parameters and contextual elements needed by the context-aware decision making service are made available by auxiliary services responsible for *extracting* (i) *prediction parameters* from received prediction events and (ii) *contextual parameters* that are important for the enterprise decision problems considered, from enriched sensed data carrying out contextual information, respectively.

- (Act) The act phase includes the business improvement analyzer, a component which deals with the visualization of anticipation-driven recommendations, the monitoring of their success, as well as the definition and monitoring of KPIs and corresponding adaptation of the whole OODA cycle, closing the feedback loop and leading to the continual proactive business optimization.

The real-time processing layer of the architecture has been implemented as a Storm¹ topology. Storm is a distributed data processing system whose processing is based on elements organised in a topology and called spouts and bolts. Spouts, which are the entry points into the real-time processing layer, poll relevant data sources such as sensors and distribute the data further in the topology. Bolts, which are the processing elements, implement the OODA information-processing services described above. Bolts are interconnected with an internal pub/sub mechanism and communicate through messages called tuples. The Decision Configurator Dashboard of the User Interaction Layer has been implemented as a Python web-application developed using the web2py² framework. Web2py is an open-source web framework (released under the LGPL version 3 license) for agile development of secure database-driven web applications, written also in Python. Finally, the Data layer of the architecture includes a Database Abstraction Layer (DAL) that generates SQL statements, transparently to the developer, for many databases engines such as SQLite, MySQL, PostgreSQL, MSSQL, FireBird, Oracle, IBM DB2, Informix and Ingres.

4 Envisaged Scenario

In this section we present a practical application of the proposed architecture for proactive event-driven decision-making, in the oil and gas industry. We describe the practical role and use of the proposed architecture focusing on how it can support decision-making ahead of time on the basis of real-time observations and real-time data-driven predictions of future undesired events, through an indicative scenario of proactive Condition-Based Maintenance (CBM). The practical application is illustrated through a decision configurator dashboard that receives prediction events and enables the embodiment of domain knowledge given by an expert in order to provide recommendations for the proactive enterprise.

CBM in the oil and gas industry employs various monitoring means to detect deterioration and failure in some critical drilling equipment. Equipment failure situations can be forecasted based on observations of events related to this equipment or the surrounding environment; e.g monitoring engine temperature indicators, monitoring electric indicators (measuring change in the engine's electric properties) and perform-

¹ <https://storm.apache.org>

² <http://web2py.com/books/default/reference/29/web2py>

ing oil analysis [20]. In reality, several different patterns will imply various failure distributions. In this scenario, we focus on the gearbox drilling equipment and consider as indicators the rotation speed of the drilling machine's main shaft in Rounds Per Minute (RPM), along with the lube oil temperature of the drilling machine's gearbox [21].

The OODA model - on which the real-time processing layer of the proposed architecture is based - deals first with real-time smart sensing of RPM and lube oil temperature of the drilling machine's gearbox (Observe). In the Orient phase, a prognostic model is developed in order to estimate Remaining Useful Life (RUL) and the probability distribution of the occurrence of a gearbox breakdown. This prognostic model is triggered by detecting in real-time abnormal friction losses on the basis of observed data. The friction losses detection deals with complex patterns of oil temperature and RPM events characterized by an abnormal oil temperature rise (10% above normal) measured over 30% of the drilling period when drilling RPM exceeds a threshold. This pattern, learned at the offline phase, is a strong indication that a gearbox equipment failure starts to occur. Decide phase deals with online provision of proactive recommendations of maintenance actions (take the equipment down for full maintenance, perform lubrication of metal parts, shift drilling to lower pressure mode) and suggested activation time that maximize the utility for the manufacturing enterprise. Finally, the Act phase defines and monitors related KPIs such as downtime and cost.

The decision configurator dashboard addresses the Decide phase and enables the expert to insert domain knowledge that is needed for the provision of recommendations as it is included in the user interaction layer of the proposed architecture. The Decide phase receives a prediction event from the Orient phase and utilizes domain knowledge to provide optimal solutions for maintenance. We have selected MDP method as the most suitable for the current scenario as it is a method that can provide recommendations about the optimal action and the optimal time of applying it and therefore, it covers the user requirements [22]. The user is able to create a new decision making instance for the use of MDP method in order to provide recommendations based on the gearbox breakdown prediction. Then, the user inserts the list of actions (take the equipment down for full maintenance, perform lubrication of metal parts, shift drilling to lower pressure mode) accompanied with the cost as a function of time for each action, the delay of the action (corresponding to the time period from its implementation until it starts taking effect) and the time-to-breakdown after the implementation of each action. Moreover, the user specifies the time of next planned maintenance.

The cost of each action can be either fixed (e.g. 1000 euros) or variable as a function of time (e.g. 800 euros / day due to production loss). The duration of the delay for each action increases in proportion to the complexity of the action (e.g. full maintenance requires a longer delay in comparison to lubrication). The time-to-breakdown after the implementation of each action is related to the extent of the maintenance action. For example, full maintenance transforms the equipment to good-as-new, while lower pressure and lubrication are actions for imperfect maintenance [1], [23]. Finally, the time of next planned maintenance corresponds to the end of decision epoch parameter of MDP. So, a recommendation about a maintenance action should

belong to the time period between now and the end of decision epoch in order to be valid [1]. This means that, since any defective part of equipment will be identified during the planned maintenance, the predicted gearbox breakdown will not be valid anymore after that time.

The business added value of proactive event-driven decision-making in this scenario is huge. With a typical day rate for a modern oil rig being around USD 500 000, reducing undesired downtime, with its associated high cost (one hour of saved downtime is typically worth USD 20 000) is of outmost importance in the oil drilling industry. Therefore, we expect that the proposed architecture, which supports the provision of proactive recommendations about optimal decisions on the basis of utility, cost and other factors, will allow enterprises in the oil and gas industry to gain a strong competitive advantage based on reduced downtimes and optimized performance.

5 Conclusions and Future Work

We outlined a visionary approach for a new architecture supporting proactive decision making in enterprises. The main novelty of our approach lays on the integration of state-of-the-art decision methods into an event-driven real time environment which can handle big data generated by a multitude of enterprise sensors. Although the concept of proactive decision-making is not completely new, there are still many challenges associated with its application in large scale, big data-based enterprise environments. A major challenge is the treatment of anticipation as a first class citizen: supporting the whole life-cycle of the anticipation, from sensing and generating anticipations till validating the reactions based on them, through the Observe-Orient-Decide-Act loop.

Our work has several implications for both practitioners and researchers. Practitioners need to design and implement physical (such as smart sensors and actuators, location-aware sensors, cyber-physical systems) and virtual sensors (such as agents in customer transaction and relationship systems) in virtually every aspect of their enterprise that has an impact on the end result. Moreover, practitioners should be ready to select and apply decision methods that will leverage business performance through the proactive realization of actions in anticipation of predicated enterprise challenges and opportunities. There is a need for new business methodologies that will enable recognition of possible opportunities for application decision methods in the enterprise environment, design of actions for responding to such situations as well as benchmarks for comparison real performance measurements and identified KPI's.

Researchers will be able to build on top of the proposed real-time platform new algorithms to model the data that overcome the deficiency of traditional analytic methods by fixing the granularity (resolution) of the analytic problem ahead of time. Additionally, new methods for semantic enrichment of the data with the background knowledge from ontologies describing the data domain and its contextual environment will improve the capabilities of the platform to react intelligently. Finally, new situational awareness methods for inferring users' situational state and approaches to model situational probabilistic influences on user needs as well as new data-driven

and knowledge-based recommender algorithms based on streaming data can be researched so that business users are provided with the most relevant support based on rich landscape of sensed data.

Regarding future work, we aim to follow a multi-aspect approach for validating the main blocks of the proposed architecture which are currently under development in the ProaSense project. We will pursue validation in diverse enterprise settings with different technical constraints and user requirements so that the impact is leveraged. Validation will be performed on a technical level, covering system-related metrics such as performance, and on a business level, covering the benefits for end-users from the leveraged business decisions. Specifically for the maintenance perspective, validation will be focused based on performance in terms of decreased maintenance costs, decreased equipment deterioration and better quality of products that are leaving the assembly line. On the other hand, domain experts will validate the results based on factors which are usually hard to measure such as increase in safety, decrease of environmental impact and the accessibility and adaptability of the system in order to support the needs and expectations of beneficiaries. Validation of the approach will be performed in the context of the ongoing EU project ProaSense (<http://www.proasense.eu/>) in two main use cases: proactive manufacturing in the area of production of lighting equipment, and proactive maintenance within the oil and gas sector.

Acknowledgments. This work is partly funded by the European Commission project FP7 STREP ProaSense “The Proactive Sensing Enterprise” (612329) <http://www.proasense.eu/>.

References

1. Engel, Y., Etzion, O., Feldman, Z.: A basic model for proactive event-driven computing. In: 6th ACM Conf. on Distributed Event-Based Systems, pp. 107–118. ACM (2012)
2. Peng, Y., Dong, M., Zuo, M.J.: Current status of machine prognostics in condition-based maintenance: a review. *J. Advanced Manuf. Technology* **50**(1–4), 297–313 (2010)
3. Bousdekis, A., Magoutas, B., Apostolou, D., Mentzas, G.: A Proactive Decision Making Framework for Condition Based Maintenance. *Industrial Management & Data Systems* **115**(7), 1225–1250 (2015)
4. Luckham, D.: *Power of events*. Reading: Addison-Wesley (2002)
5. Dunkel, J., Fernández, A., Ortiz, R., Ossowski, S.: Event-driven architecture for decision support in traffic management systems. *Expert Systems with Applications* **38**(6), 6530–6539 (2011)
6. Engel, Y., Etzion, O.: Towards proactive event-driven computing. In: Proceedings of the 5th ACM International Conference on Distributed Event-Based System, pp. 125–136. ACM (2011)
7. Fournier, F., Kofman, A., Skarbovsky, I., Skarlatidis, A.: Extending event-driven architecture for proactive systems. In: *Event Processing, Forecasting and Decision-Making in the Big Data Era (EPForDM), EDBT 2015 Workshop* (2015)
8. Feldman, Z., Fournier, F., Franklin, R., Metzger, A.: Proactive event processing in action: a case study on the proactive management of transport processes. In: Proceedings of the Seventh ACM International Conference on Distributed Event-Based Systems (DEBS 2013), pp. 97–106 (2013)

9. Muller, A., Suhner, M.C., Iung, B.: Formalisation of a new prognosis model for supporting proactive maintenance implementation on industrial system. *Reliability Engineering & System Safety* **93**(2), 234–253 (2008)
10. Lee, J., Ni, J., Djurdjanovic, D., Qiu, H., Liao, H.: Intelligent prognostics tools and e-Maintenance. *Computers in Industry, Special Issue on e-Maintenance* **57**(6), 476–489 (2006)
11. Muller, A., Crespo Marquez, A., Iung, B.: On the concept of e-maintenance: review and current research. *Reliability Engineering & System Safety* **93**(8), 1165–1187 (2008)
12. Levrat, E., Iung, B.: TELMA: a full e-maintenance platform. In: *Proceedings of the Second World Congress on Engineering Asset Management (WCEAM 2007)* (2007)
13. Irigaray, A.A., Gilabert, E., Jantunen, E., Adgar, A.: Ubiquitous computing for dynamic condition-based maintenance. *Journal of Quality in Maintenance Engineering* **15**(2), 151–166 (2009)
14. Pistofidis, P., Emmanouilidis, C., Koulamas, C., Karampatzakis, D., Papatthanassiou, N.: A layered e-maintenance architecture powered by smart wireless monitoring components. In: *Proceedings of the 2012 International Conference on Industrial Technology (ICIT 2012)*, pp. 390–395. IEEE (2012)
15. Iung, B., Levrat, E., Marquez, A.C., Erbe, H.: Conceptual framework for e-Maintenance: Illustration by e-Maintenance technologies and platforms. *Annual Reviews in Control* **33**(2), 220–229 (2009)
16. Campos, J., Jantunen, E., Prakash, O.: A web and mobile device architecture for mobile e-maintenance. *The International Journal of Advanced Manufacturing Technology* **45**(1–2), 71–80 (2009)
17. Macchi, M., Crespo Márquez, A., Holgado, M., Fumagalli, L., Barberá Martínez, L.: Value-driven engineering of E-maintenance platforms. *Journal of Manufacturing Technology Management* **25**(4), 568–598 (2014)
18. Elwany, A.H., Gebrael, N.Z.: Sensor-driven prognostic models for equipment replacement and spare parts inventory. *IIE Transactions* **40**(7), 629–639 (2008)
19. Boyd, J.R.: *The Essence of Winning and Losing*. Unpublished lecture notes (1996)
20. Jagadish, H.V., Gehrke, J., Labrinidis, A., Papakonstantinou, Y., Patel, J.M., Ramakrishnan, R., Shahabi, C.: Big data and its technical challenges. *Communications of the ACM* **57**(7), 86–94 (2014)
21. Magoutas, B., Stojanovic, N., Bousdekis, A., Apostolou, D., Mentzas, G., Stojanovic, L.: Anticipation-driven architecture for proactive enterprise decision making. In: *CAiSE 2014*, pp. 121–128 (2014)
22. Bousdekis, A., Magoutas, B., Apostolou, D., Mentzas, G.: Supporting the selection of prognostic-based decision support methods in manufacturing. In: *Proceedings of the 17th International Conference on Enterprise Information Systems (ICEIS 2015)*, pp. 487–494 (2015)
23. Jardine, A.K., Lin, D., Banjevic, D.: A review on machinery diagnostics and prognostics implementing condition-based maintenance. *Mechanical Systems and Signal Processing* **20**(7), 1483–1510 (2006)