Scalable Proactive Event-Driven Decision Making

Alexander Artikis, Chris Baber, Pedro Bizarro, Carlos Canudas de Wit, Opher Etzion, Fabiana Fournier, Paul Goulart, Andrew Howes, John Lygeros, Georgios Paliouras, et al.

► To cite this version:

Alexander Artikis, Chris Baber, Pedro Bizarro, Carlos Canudas de Wit, Opher Etzion, et al.. Scalable Proactive Event-Driven Decision Making. IEEE Technology and Society Magazine, 2014, 33 (3), pp.35-41. 10.1109/MTS.2014.2345131 . hal-01416881

HAL Id: hal-01416881

<https://hal.science/hal-01416881>

Submitted on 15 Dec 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Scalable Proactive Event-Driven Decision-Making

Alexander Artikis¹, Chris Baber², Pedro Bizarro³, Carlos Canudas-de-Wit⁴, Opher Etzion⁵,
Fabiana Fournier⁶, Paul Goulart⁷, Andrew Howes², John Lygeros⁷, Georgios Paliouras¹,
Assaf Schuster⁸ and Izchak Sharfman⁸

¹NCSR Demokritos, Greece, ²University of Birmingham, UK, ³FeedZai, Portugal,

⁴CNRS, France, ⁵Yezreel Valley College, Israel, ⁶IBM Research Haifa, Israel,

⁷ETH Zurich, Switzerland, ⁸Technion-Israel Institute of Technology, Israel

1. INTRODUCTION

Rapid social, economic and political changes are leading organizations to shift their thinking from reactive to proactive in order to detect opportunities and threats that could affect their business [6]. Eliminating or mitigating an anticipated problem, or capitalizing on a forecast opportunity, can substantially improve our quality of life, and prevent environmental and economic damage. Changing traffic light policies and speed limits to avoid traffic congestions, for example, will reduce carbon emissions, optimize public transportation and increase the quality of life and productivity of commuters. Similarly, adding credit cards to watch lists as a result of forecasting fraud will reduce the cost inflicted by fraudulent activities on payment processing companies and merchants, and consequently lower credit card rates.

In energy management, there is a need for real-time optimization of power consumption in individual houses and buildings equipped with renewable energy sources. This requirement may be addressed by forecasting energy consumption and production, say for the next 30 minutes, and making decisions about load adjustments and/or rescheduling. In post-earthquake disaster management, loss forecasts can be vital in planning the actions to be taken immediately after an earthquake occurs.

To prevent problems and capitalize on opportunities before they even occur, we propose a methodology for proactive event-driven decision-making. Decisions are triggered by forecasting events instead of reacting to them once they happen. The motivation for proactive computing stems from social and economic factors, and is based on the fact that prevention is often more effective than the cure. The decisions are made in real-time and require on-the-fly processing of Big Data, that is, extremely large amounts of noisy data flooding in from various locations, as well as historical data.

Proactive applications have been developed for several years [19]. Consider e.g. proactive security systems and proactive routing in mobile ad-hoc networks. Proactive applications have been largely developed in an ad hoc manner. In contrast, we aim to develop a generic methodology for proactive event-driven computing.

We are moving from the information economy to the “intelligent economy”, where it is not only access to information but the ability to analyze and act upon it that enables sustainable management of communities, and promotes appropriate distribution of social, healthcare, and educational services [11]. Our methodology for proactive event-based decision-making, therefore, comprises the following steps. First, Big Data is continuously acquired and aggregated from various types of sensor. The aggregated data is analysed and fused in order to recognise, in real-time, events and situations of special

significance. To allow for timely recognition, communication volume is minimized by moving as little data as possible from one place to another. Second, the events recognised are correlated with historical information to forecast problems and opportunities that may actually take place in the near future. Third, the forecast events along with the recognised events are leveraged for real-time operational decision-making. Fourth, visual analytics tools prioritise and explain possible proactive actions, enabling human operators to reach and execute informed decisions.

Proactive computing requires capabilities for forecasting, real-time decision-making and visual analytics. These capabilities are extremely important in a multitude of application domains. E.g. no system supports fraud forecasting. Furthermore, a typical fraud detection system may raise up to 9 false alarms for each true alarm. Without an appropriate explanation of why a specific transaction is considered fraudulent, the operator overseeing transactions will not be able to confirm the fraud and will have to either let it go through, or contact the end user, spending valuable time. We propose exposing the user to how proactive computing works through visual analytics.

To summarize, our methodology supports on-the-fly, low-latency processing of extremely large, geographically distributed, noisy event streams and historical data, for recognizing and forecasting opportunities and threats, making decisions to capitalize on the opportunities and mitigate the threats, and explaining, through user-interaction, the decisions to human operators in order to facilitate informed decision execution.

2. APPLICATIONS

Our approach to proactive event-driven decision-making is applicable to a wide range of application domains. In the context of the SPEEDD project (<http://speedd-project.eu/>), we will evaluate our methodology in two such domains: traffic management and credit card fraud management. Proactive traffic management is realized following the steps below.

- **Detect.** Vehicles along with their speed are detected in a road sector and/or at a specific point through a multi-technology sensor network. This information is used to recognize traffic flow and density patterns along the selected road. Traffic accidents are also recognized on the basis of acceleration/deceleration patterns, as well as violations of road safety (vehicles driving too close to each other, long vehicles driving too fast, etc).
- **Forecast.** Traffic flow, traffic density and subsequently duration of travel will be forecast for different time horizons in the future. The carbon print (CO₂) and energy consumption

will also be forecast for different waiting times (5, 10, 15, 20 minutes).

- **Decide.** Calculate within 30 seconds the optimal variable speed limits and duty cycles for the ramp metering lights.
- **Act.** Change the actual values of the variable speed limit panels and the operation of lights on the ramp metering course. Actions will be taken in a matter of seconds for ramp metering and a matter of minutes for the variable speed limits.

Traffic forecasting requires the analysis of massive data streams storming from various sensors, including fixed sensors installed in highways and mobile sensors such as smart phones and GPS traces, as well as large amounts of historical data.

In proactive credit card fraud management, the goal is to forecast fraudulent activity and make decisions in order to prevent the financial loss. In 2010, fraud in the Single Euro Payments Area (that includes 27 EU member states) was estimated at 1.26 billion Euros [12]. Fraud detection is a needle in the haystack problem as fraudulent transactions constitute at most 0.1% of the total number of transactions, while new fraud patterns appear on almost a weekly basis. Proactive credit card fraud management is realized as follows:

- **Detect.** Fraudulent activities, such the following, must be detected:
 - ‘Cloned card event’ — a credit card is being used simultaneously in different countries.
 - ‘Risky usage event’ — the card is being frequently used by a ‘risky’ merchant.
 - ‘Potential batch fraud event’ — many transactions from multiple cards are being used in the same point-of-sale in high amounts.
- **Forecast.** Recognize fraudulent activity that has started to take place but is not completed yet. E.g. forecast with a certain probability a ‘risky usage event’ when there are a few transactions by ‘risky’ merchants in some period of time.
- **Decide.** Decide to block or review the transaction in less than a second after the forecast.
- **Act.** Depending on the type of fraud, add the corresponding credit card to the black/gray/watch list.

Credit card fraud forecasting requires the analysis of very large, noisy transaction streams storming from all over the world, as well as massive amounts of historical data.

3. APPROACH

We propose a highly synergetic approach to proactive event-driven decision-making by combining the research areas of event processing, scalable data processing, optimization for decision-making, and decision support through visual analytics. The approach will be realized in a distributed system comprising the following components:

- Real-time event recognition and forecasting under uncertainty. Events of special significance are recognized and forecast, and then communicated to the decision-making component. To allow for timely recognition, communication volume is minimized by moving as little data as possible from one place to another.
- Real-time event-based decision-making under uncertainty. The forecast and recognized events are leveraged for real-time operational decision-making.
- Visual analytics for proactive decision support. Visualization techniques explain the decisions made and the

possible proactive actions, enabling human operators to reach and execute informed decisions.

Figure 1 illustrates our methodology in the context of proactive traffic management. The following sections present in detail the main aspects of the methodology.

3.1 EVENT RECOGNITION & FORECASTING

Systems for symbolic event recognition [4] (event pattern matching) identify composite events of interest — collections of events that satisfy some pattern. The ‘definition’ of a composite event imposes temporal, logical and, possibly, spatial constraints on its sub-events, that is, events coming from sensors or other composite events. Consider e.g. the recognition of a traffic incident in a road segment given the speed of the vehicles passing that segment.

Typically, event recognition systems operate on top of stream processing platforms [1]. This way, complex events are defined by means of expressive event recognition languages, and efficiently detected using the optimized data processing of stream processing platforms.

Event recognition systems have to deal with various types of uncertainty, such as incomplete data streams, erroneous data and imperfect composite event definitions [5]. E.g. in traffic management fixed sensors are often out of order, inappropriately calibrated or inaccurate. To address this requirement, we will develop a framework for real-time event recognition able to deal with the inherent uncertainty of Big Data. The framework will exhibit a declarative, formal (probabilistic) semantics. To achieve this task, we will build upon existing frameworks combining probabilistic reasoning, such as Markov Logic Networks [10], and symbolic methods. The starting point will be probabilistic extensions of the Event Calculus [2], [5] — a logic programming language for representing and reasoning about events and their effects. Probabilistic Event Calculi facilitate the integration of domain knowledge, such as traffic models, and deal with uncertainty both in the input data and the composite event definitions. To minimize the performance overhead of uncertainty reasoning, we will place emphasis on distributed probabilistic reasoning techniques (see Section 3.2).

To allow for proactive decision-making, we will develop a framework for event forecasting able to deal with the volume and lack of veracity of Big Data. The framework will indicate the probability of a forecast event, as well as the probability of when an event will happen; a probability distribution over the expected event occurrence time will be provided. The basis of this framework will be ‘forward’ event recognition algorithms that are capable of recognizing incrementally composite events, but incapable of dealing with the lack of veracity [8].

The manual development of composite event definitions is a tedious, time-consuming and error-prone process. Machine learning techniques may be used for the acquisition of domain knowledge: constructing and/or refining composite event definitions (expressing e.g. traffic congestion) in dynamic and evolving environments. A common technique for learning the structure of composite event definitions in a supervised manner involves the use of Inductive Logic Programming (ILP) (e.g. [17]). ILP constructs theories that capture exceptional cases in data streams. This is particularly helpful in highly imbalanced

streams such as those of credit card fraud. On the other hand, ILP does not handle numerical reasoning, such as comparing the time-points of events emitted by vehicles, which is quintessential in the representation of composite event definitions. In the case of partial supervision, ILP is used in combination with abduction in order to learn an event definition. This combination of techniques, however, does not scale to Big Data.

In addition to learning the structure of a composite event definition, the confidence values/weights attached to the definition can be learned from data. Usually the tasks of structure learning and weight learning are separated; that is, first the structure of an event definition is learnt and then the weights of the definition are estimated. Separating the two learning tasks in this way, however, may lead to suboptimal results, as the first optimization step (structure learning) needs to make assumptions about the weight values, which have not yet been optimized.

To address these issues and avoid the error-prone process of manual composite event definition construction, our methodology will consist of incremental learning techniques for successfully combining abduction with induction in Big Data. Furthermore, we will develop techniques for the simultaneous optimization of the numerical parameters of a composite event definition (weights and numerical temporal constraints) and its structure.

3.2 SCALABLE PROCESSING

The high velocity of incoming events poses challenges both in terms of computational resources and in terms of communication resources. Computational scalability issues are addressed by distributing event recognition tasks among multiple nodes (see e.g. [16]), while communication scalability issues are addressed by algorithms that perform as much of the processing as possible on the nodes where events are generated, thus reducing the amount of data that is transferred between nodes (see e.g. [15], [9]).

In traffic management, for example, a common task is counting the number of vehicles traversing on a set of paths, where some of the paths may have shared locations (consider e.g. paths {A, B, C} and {A, D, B}). A simplistic approach that does not take uncertainty into consideration would use detectors at each of the points (A, B, C and D), and define two patterns consisting of the corresponding sequences. The system would detect these sequences using finite state automata. The volume and velocity of the events that are required to be processed, as well as the complexity of some of the automata, require distributing the automata processing task among multiple nodes.

A more realistic solution to the path counting task is to take into account the uncertainty in the detection of the locations of vehicles. Detectors may fail to detect some vehicles, may have false detections, and may report detections that are inherently uncertain (e.g. locating vehicles via a cellular network). Automata used for detecting patterns over deterministic events are unsuitable in this scenario. On the other hand, as discussed in the previous section, probabilistic models such as Markov Logic Networks are designed to handle uncertainty, and are therefore a natural choice for detecting events under uncertainty. Event recognition and forecasting with Markov Logic Networks is done by inference over probabilistic graphical models, which is fundamentally different than computations over state automata. Consequently, distributing these tasks among multiple nodes requires fundamentally different algorithms.

To address the Big Data issues of volume, velocity and lack of veracity, therefore, we will develop methods for distributing event recognition and forecasting tasks that incorporate probabilistic reasoning. This requires distributing on-line inference tasks among multiple nodes, as opposed to state automata used for recognition tasks over deterministic events. The proposed algorithms exploit the continuous nature of the recognition task by incrementally modifying the inference as new events arrive.

In addition to the computational scalability issues discussed above, the increasing number of distributed event-generating sources requires that inherently-limited network resources be employed efficiently. E.g. in traffic management some sensors may be deployed at locations where a high speed wide area network is not available, and will therefore be required to continuously transmit a high volume of sensor readings via a cellular network. Since communication efficiency reduces the volume of data sent to a data center for processing, it may also improve computational efficiency. Communication efficiency also helps in maintaining the privacy of the entities generating the events (e.g. terminals in credit card transactions).

Communication-efficient distributed detection has been an active research field in recent years. Proposed methods reduce communication by decomposing the recognition task into a set of local constraints on the data generated at the nodes. The constraints are such that as long as all of them are upheld, it is guaranteed that the event of interest has not occurred. Consequently, as long as all constraints are upheld, no communication is required. The event to be recognized is usually defined using a function over aggregate values derived at the nodes. In other words, event recognition is restricted to numerical reasoning.

To support the full range of functionality required by Big Data applications, we will develop distributed communication-efficient event recognition and forecasting algorithms. This includes events defined over aggregates as well as temporal, logical and spatial patterns over events as discussed in the previous section. Emphasis will be placed in handling functions that do not have a closed form, such as inference over probabilistic graphical models.

3.3 EVENT-DRIVEN DECISION-MAKING

In the proposed methodology, the forecast events along with the recognized events are leveraged for real-time operational decision-making. A body of tools for real-time proactive decision-making exploits the event forecasting models presented above, with an emphasis on optimization methods that intelligently handle forecast uncertainty using robust, stochastic or black-box methods.

In terms of real-time optimization techniques, the state-of-the-art is that optimization techniques are being activated mostly off-line and use a variety of optimization methods that fit different assumptions, e.g. robust (worst-case) optimization or stochastic optimization. In the field of robust optimization methods, the state-of-the-art focuses on tools for providing strong performance guarantees for convex optimization problems [3]. For real-time decision-making purposes, the use of robust optimization methods involving recourse, that is, modeling the notion that future decisions can be deferred until future information is available, is an area of intensive ongoing research. E.g. in the context of traffic management, ‘recourse’ decisions refer specifically to traffic management actions (such as alteration of speed limits and

restriction of on-ramp flows) computed as future responses to changes in traffic flows resulting from similar actions taken at an early time. The use of ‘robust’ or ‘worst-case’ models is most appropriate for those aspects of traffic management with hard limits, such as absolute limits on allowed flows or maximum closure time constraints.

Stochastic optimization focuses on optimizing an expected value criterion subject to probabilistic constraints. Aside from the need to parameterize policies in the recourse sense discussed above, an additional difficulty relates to the interpretation of constraints. Due to the probabilistic nature of the uncertainty that enters the optimization, the hard, worst-case constraints used in robust optimization often turn out to be infeasible. One then has to resort to soft interpretations, such as chance constraints ensuring that the probability of meeting the constraint is above a certain threshold, integrated chance constraints ensuring that the expected value of a constraint function is above a certain threshold, or interpretations based on distributional robustness and conditional value-at-risk. In the context of traffic management problems, stochastic optimization methods are most appropriate when handling performance constraints that are ‘soft’ in a probabilistic sense, that is, the traffic management system is tasked with respecting the constraint with a high likelihood, or respecting it most of the time. Such constraints include expected transit time constraints and mean traffic flow targets. Stochastic optimization methods are also most useful for problems in which large amounts of historical data can be accessed to provide example ‘scenarios’ for modeling purposes. In traffic management, historical data relating to traffic inflows and congestion supply exactly these scenarios.

Our methodology for proactive event-driven decision making will advance the state-of-the-art in each of the preceding areas in two distinct ways. The first is to determine which aspects of the application under consideration should be treated in each way. The second, and more challenging, task is to develop real-time proactive planning tools for traffic and credit card fraud management using these optimization methods within an event-based planning framework. These methods will then be employed at a variety of levels of autonomy, ranging from simple decision support functions for human operators to fully autonomous decision-making.

3.4 VISUAL ANALYTICS

While the aim of our methodology is to automate much of the decision-making process, key points will require people to make choices and the system realizing the methodology will require human monitoring. E.g. in traffic management, determining the trade-off between minimizing average journey times and setting acceptable thresholds on maximum wait times requires human monitoring. Other tasks such as communicating traffic state, advising road users and road planners, require operators to maintain a good mental model of the dynamics of the road system, and also of the decision-making system itself (see the previous section). The effectiveness of human decisions will be enhanced to the extent that the dynamics of the entire system can be made transparent.

We will address these issues through visualization technologies that are tuned to what is known about human decision-making processes. We will build on work in online information foraging for decision-making [20] and in the time signature of the human cognitive architecture to drive new designs for visualization. Subtle changes in the time costs of making comparisons can lead

to macroscopic changes in decision strategy [13] and, indeed, we contend it is this regularity that provides the key opportunity for visualization technologies. For example, it is known that requiring users to mouse-over icons in order to reveal decision critical information reduces the amount of information that users retrieve, despite the fact that it only adds hundreds of milliseconds to the interaction. More interestingly, mouse-over designs can shift users from using non-compensatory to more compensatory strategies. Conversely, presenting too much information all at once leads to visual ‘crowding’ and the potential for feature swap, e.g. numerical transposition errors, and therefore error.

Visualization technologies work not simply because they are visual, but because, by enhancing the efficiency with which people can compare results, visualization can fundamentally modify the processes by which decisions are made. In the proposed system for proactive decision-making, visualization design will emphasize comparison, as others have done, but will do so as directed by recent theory in the cognitive sciences [21]. We also need to push beyond the individual. While much research on visualization has focused on understanding the performance of individuals engaged in diagnosis tasks, we contend that there is considerable potential for new insights for the design of collaborative visualization technologies. Visual Analytics is not simply the visualization of the output from analysis processes, but the creation of insight in the decision-makers working with these visualizations, that is, the analysts are active participants in constructing the manner in which these data are to be processed, creating and revising associations between parts of the dataset by manipulating the graphical user interface [7].

To develop visual analytics for decision support in Big Data applications, we will apply concepts and principles from Ecological Interface Design [18]. ‘Ecological Interfaces’ are designed to visualize the manner in which physical components of the system map onto the (more abstract) functions that the system performs. So, they are views of the process which are not simply maps of how physical components connect to each other but are abstractions which show how types of physical components affect particular functions. The purpose of such designs is to improve operator decision-making and diagnosis when dealing with faults relating to those specific functions. For our system, this means that the visualization will not only display the model’s input and output, but also the relationships between elements in the decision space. One element of Ecological Interface Design is simply the reflection of the constraints in the work domain through constraints in the user interface. In this way, the ‘ecologies’ of the work domain, of the environment and of the organization become reflected in the user interface through the definition and management of these constraints. Added to these ecological constraints are constraints from the analyst/modeler, such as expectations and mental models.

4. DISCUSSION

Passively waiting until a plan is missed is an expensive way to solve the problem and increasingly risky, particularly when prevention and problem optimization can be designed into the process [14]. We proposed a methodology for proactive event-driven decision-making in order to eliminate or mitigate anticipated problems, and capitalize on forecast opportunities. By facilitating proactive decision-making, we expect to open up a range of new opportunities for services that will help people in their everyday lives. Indeed, there is an ever increasing need for

knowing how to forecast that something will happen and when it will happen (e.g. a heart attack, or an act of violence), and knowing what to do before it happens. Therefore, our methodology is expected to have a significant impact on time-critical and often life-critical situations, where it is vital to prevent problems and capitalize on forecast opportunities.

5. ACKNOWLEDGMENTS

This work is part of the EU-funded SPEEDD project (FP7-ICT 619435). We would like to thank the anonymous reviewers for their helpful comments.

6. REFERENCES

- [1] A. Artikis et al. Heterogeneous Stream Processing and Crowdsourcing for Urban Traffic Management. EDBT, 2014.
- [2] A. Artikis, M. Sergot and G. Paliouras. Run-Time Composite Event Recognition, DEBS, 2012.
- [3] A. Ben-Tal, S. Boyd. and A. Nemirovski. Extending Scope of Robust Optimization: Comprehensive Robust Counterparts of Uncertain Problems. Mathematical Programming. Vol. 107, 1, 2006.
- [4] A. Margara and G. Cugola. Processing flows of information: From data stream to complex event processing. ACM Computing Surveys. 2012.
- [5] A. Skarlatidis et al. A Probabilistic Logic Programming Event Calculus. Theory and Practice of Logic Programming. 2014.
- [6] B. Burton et al. Predicts 2011: Pattern-Based Strategy Technologies and Business Practices Gain Momentum . Gartner Report #G00208030. 2010.
- [7] C. Chabot. Demystifying Visual Analytics. IEEE Computer Graphics and Applications. Vol. 29, 2009.
- [8] C. Dousson, and P. Le Maigat. Chronicle Recognition Improvement Using Temporal Focusing and Hierarchization. IJCAI. 2007.
- [9] D. Keren et al. Geometric Monitoring of Heterogeneous Streams. IEEE TKDE. 2014
- [10] D. Lowd. and P. Domingos. Markov Logic: An Interface Layer for Artificial Intelligence. Morgan & Claypool Publishers. 2009.
- [11] D. Vesset et al. Worldwide Decision Management Software 2010-2014 Forecast: A Fast-Growing Opportunity to Drive the Intelligent Economy. IDC report #226244. 2010.
- [12] European Central Bank: Report on Card Fraud. <http://www.ecb.int/pub/pdf/other/cardfraudreport201207en.pdf>
- [13] G. L. Lohse and E. J. Johnson. A Comparison of Two Process Tracing Methods for Choice Tasks. Organizational Behavior and Human Decision Processes. Vol. 68, 1, 1996.
- [14] M. Fleming. Business Navigation Systems Combine CEP with BPM. IDC Report. 2009.
- [15] N. Giatrakos et al. Distributed Geometric Query Monitoring using Prediction Models. ACM TODS. 2014.
- [16] N. P. Schultz-Moller, M. Migliavacca and P. Pietzuch. Distributed complex event processing with query rewriting. DEBS. 2009.
- [17] O. Ray. Nonmonotonic abductive inductive learning. Journal of Applied Logic. 2009.
- [18] P. Duez and K. J. Vicente. Ecological interface design and computer network management: the effects of network size and fault frequency. International Journal of Human-Computer Studies. Vol. 63, 2005.
- [19] R. Want, T. Pering and D. Tennenhouse. Comparing Autonomic and Proactive Computing. IBM Systems journal. 2003.
- [20] S. Lelis and A. Howes. A Bayesian Model of How People Search Online Consumer Reviews. Proc. CogSci. 2008.
- [21] Z. Pousman, J. Stasko and M. Mateas. Casual Information Visualization: Depictions of Data in Everyday Life. IEEE Transactions on Visualization and Computer Graphics. Vol. 13, 2007.

Figure 1. Proactive Traffic Management. Sensors and actuators are labeled with ‘N’ and ‘M’ on the motorway.