Aspects of Modelling and Processing Complex Networks of Operations' Risk

Udo Inden
Cologne University of Applied Sciences (CUAS)
Cologne, Germany
udo.inden@fh-koeln.de

Despina T. Meridou,
Maria-Eleftheria Ch. Papadopoulou,
Angelos-Christos G. Anadiotis,
Iakovos S. Venieris[b]
School of Electrical and Computer Engineering,
National Technical University of Athens
Athens, Greece
{dmeridou, marelpap, aca}@icbnet.ece.ntua.gr
[b]venieris@cs.ntua.gr

Claus-Peter Rückemann
Leibniz Universität Hannover /
Westfälische Wilhelms-Universität,
Münster, Germany
ruckema@uni-muenster.de

***Abstract*— "Landscape of risk" (RL) is a metaphor to describe agglomerations of interdependent risk. The idea is to integrate the full scale, variety, velocity, variability and the related determinants of a complex operations' system into one computable model. The atomic elements of this network are managed nodes being exposed to risk, thus becoming source or target of unplanned events, of positive or negative impacts and propagation effects. Management is understood as continuing effort of operations' intelligence to realise and evaluate risk and to effectively act on it. The challenges are vast increases of the resolution of object and time and the accelerating change, of particularly technological innovation. These are reasons that RLs become more and dynamic, that models need to identify and capture interdependency across local and global levels and life-cycles, that learning needs to be directly integrated into the managerial workflows. Therefore, the RL concept allows for the integration of the "Big V" of data (volume, velocity, variability etc.) as well as for human and machine intelligence respectively learning. We discuss various problems and alternative models as well as architectures for processing complex landscapes and provide a first formal semantic model about the managerial handling of risk of for the management of unplanned events.**

Keywords-Integrated risk management, resolution of object and time, semantic models and technology, high-end computing

I. INTRODUCTION

A first and shorter version of this paper has been issued for the INFOCOMP Conference 2013 in Lisbon [1].

Landscapes of risk (risk landscape, RL) describe agglomerations of interdependent risk in business operations. *Risk* is one of the most general concepts of managerial decision making and capable of integrating a large variety of aspects into a coherent model of managerial acting. Of specific interest are the risk of occurrence of an event (*event risk*), positive or negative *impact* conditioned by this and strategies to learn from managing risk and impact. Figure 1 depicts basic views:

1) In the most general form, an RL is a network of interdependent nodes, each being target and source of unplanned events, i.e., of risk and impact (Figure 1-1). *Unplanned events*, discussed in Section V, are main issues of managing risk. The interdependency of nodes refers to *impact*, thus to conditioned probability as well as to learning. Figure 1-1 also differentiates autonomous (active) nodes disposing of managerial capacity as well as passive sub-nodes that are managed but may be a relevant resource.

2) Figure 1-2 describes an RL as complex supply chain, that is a distributed product-production-system (PPS) with the common goal to deliver material products or services to other businesses and finally to consumers. This network is designed according to the specifications of the PPS, as well as to the costs or availability of required resources. Impacts propagate along related dependencies.

To reduce complexity, supply-chains are typically defined on the level of the main nodes, the factories, i.e., the technological, organisational or other details may be known but are not managed on that level. However, the reasons of many failures that affect the overall efficiency of the supply

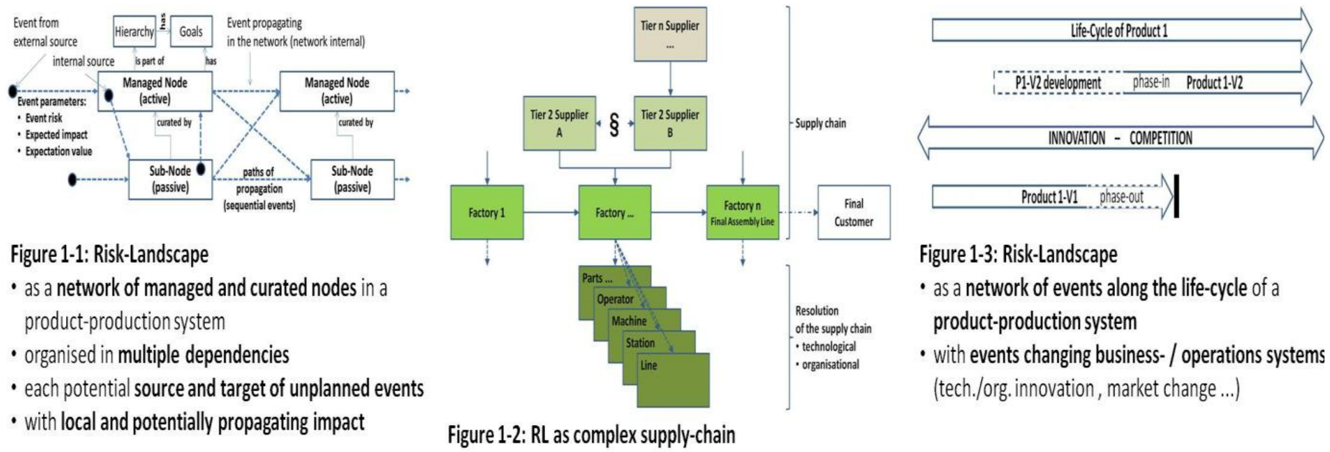

Figure 1-1: Risk-Landscape
- as a **network of managed and curated nodes** in a product-production system
- organised in **multiple dependencies**
- each potential **source and target of unplanned events**
- with **local and potentially propagating impact**

Figure 1-2: RL as complex supply-chain

Figure 1-3: Risk-Landscape
- as a **network of events along the life-cycle** of a **product-production system**
- with **events changing business- / operations systems** (tech./org. innovation , market change ...)

Figure 1. Three Basic Views at a Landscape of Risk.

chain actually lay in these details. Under conditions of close-to real time management, any detail matters, i.e., can become a target or source of unplanned events (see Section III-B).

In this respect, the former black box of the factory is resolved into production lines and possibly deeper into stations and machines in the production lines and, below, to teams or individual (named) operators. The Internet of Things also allows handling individualised resources like components or parts to be assembled. The final product of a car producer may not be the car, but the service it offers, i.e. the chain needs to include the details of a car-sharing system.

Not all of these "things" or operations are treated as active nodes (or only on demand). Nevertheless, the enormous heterogeneity of detail is the reason to increase the abstraction of models by risk-oriented concepts and by employing semantic models (Sections V, VI and VII).

3) Figure 1-3 presents the sequence of phases like the design and implementation / ramp-up, re-design needed to adopt a new technology, to respond to competition or, finally, to phase-out or terminate the product and the technology. In an RL-model, these phases translate into changing properties of nodes or deleting some of them in the network.

Learning is a major force in most of these phases: In the ramp-up phase the challenge is to overcome lacks of maturity of the design of the product or service and of the related production system (see examples in Sections I-C7 and II-B). In matured operations, it may become necessary to learn about options offered by new technologies or about the impact of competitors on the position in the market. RL-models will capture such changes or threads by adapting parameters that control risk or by adapting the structure of the network. So, the introduction of new additive production technologies will remove large segments of the suppliers' network and the related risk from the supply chain and from the RL-model.

In contrast to conventional models, RLs don't differentiate planning or simulation in operations or in life-cycle context on principle. It is always a network of interdependent risk and in the analysis of the vulnerability of an operations' system, even the same model that can be used.

The paper is organised as follows. Section II describes structural and dynamic aspects of RLs and Section III industrial research projects motivating the concept of RL. On that base, Section IV delivers a managerial framework and analyses forces that drive the problem of volume, velocity, variety, and variability in an RL.

Section V deals with knowledge models, continued by an overview of use of semantic technologies as well as Bayesian methods to gather and grow operations knowledge in Section VI, closing with a formal risk management ontology. Section XII drafts selected architectures to compute RLs, Section VIII gives an overview of future work.

II. On the Concept Of Landscapes of Risk

The idea is to integrate the full scale, variety, velocity, variability and the major related determinants into one computable model. The atomic elements of an RL are 'managed nodes' and 'unplanned events'. Nodes are the source or the target of unplanned events. Positive or negative impacts, as evaluated in the light of managerial goals, propagate in the network. Passive nodes are curated by active ones. RLs need continuous efforts to realise and evaluate risk as well as to act on it, while, on the other side, being challenged by an increasing speed of change and resolution of detail. This section describes structural and dynamic aspects of RLs.

A. On the Structures of Landscapes of Operations' Risk

Reference [2] states that "economics define investment as the act of incurring a cost in the expectation of future rewards". In business, risk is directly associated with success or failure of investment. Irrespective of the investor, the bottom line is that returns at least enable financial sustainability. This turns into rules, e.g., to maximize profit or minimize risk. If the environment changes, businesses need to adapt, i.e., profit is required to finance adaptation. Notably, change emanates from investment into innovation that implies risk, but proves its value as source of future income and, thus, triggers propagation in the markets, i.e., needs for courage to further innovate or adapt.

Figure 2 starts from investments in a business. To earn returns, functional domains with corporate (strategic, legal, financial affairs) or operational (engineering, production, purchasing) responsibility and related goals are allocated to managers (actors). Operations again are structured by processes and flows of data for planning, control and implementation actions (work). Typical passive nodes are potentially critical resources without inherent decision-making capacity.

Acting includes decisions (choices) and, thus, any actor has a managerial role in the reach of his responsibility, which is focused by related goals. This fits to a definition of Goshal and Bruch [3] of management as the "art of doing and getting done" in the reach of an area of responsibility, for given goals and related risk.

1) Managerial Responsibility: The concept of managed nodes implies actors who do not just take responsibility for the decisions they make, but also for their ability to make decisions in their specific position organisation (Figure 2), that is primarily for the access to relevant data and to information that provide context. Typically, there are downstream flows of decisions and information that provide context as well as upstream flows that enable to track and evaluate the progress of work and its impact. Horizontal exchange enables the coordination of work on the same level of the hierarchy.

2) Strategic decisions: The corporate level mainly handles risk related to corporate integrity, business model and strategy or financial sustainability. Capital bound in operations is the bridge between strategic and operational action. Within this framework and considering main parameters of the environment (e.g., issues in markets, life-cycle of products), operations' strategies specify the implementation of the business strategy into a consistent operations' strategy including objectives, accoutrements and collaborative workflows between operations' domains, or decisions on insourcing and outsourcing, choices of technology, etc.

3) Tactical decisions chose operations' policies (best practices) for given strategies or in sales with pricing or discount schemes. On this base, plans and schedules are elaborated to synchronize and optimize activities and flows of orders and materials, the provision of capacity and the rostering of staff as well as take precautions for known contingencies. Finally, these plans have to be implemented, i.e., executed, monitored and, in case of unplanned events, to be maintained or recovered ("firefighting").

4) Work-level decisions: In wording of Goshal and Bruch, firefighting (or "educated improvisation") can be defined as the capability and capacity of accomplishing a goal, in spite of unplanned events, taking, though, the dependencies in the RL into account. The rationale is to maintain or recover active plans with minimal interventions. The speed of propagation of impact and the time left to effectively act are constraints to decision making. Thus, the lowest level of management runs in an 'exception mode'. Virtually, all responsibility concentrates in this node of action and the horizontal coordination with peer-nodes along the lines of propagation.

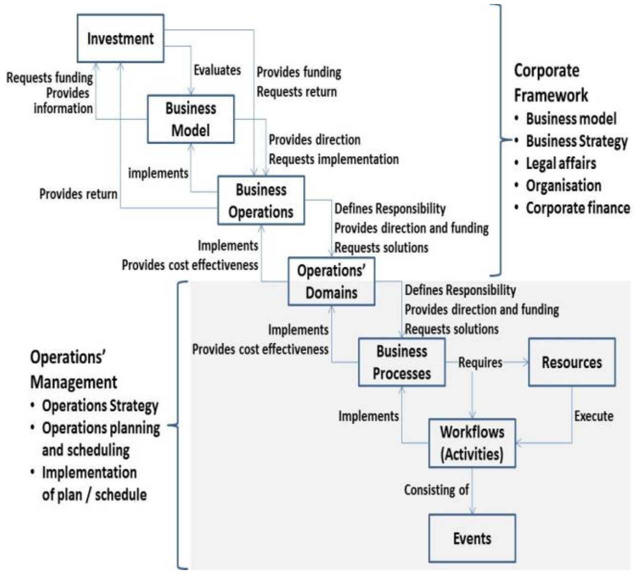

Figure 2. Locating Operations' Management in a Business Model.

5) Supra-network: Organisations depend on other organisations, e.g., across supply chains or service systems. Networks are not limited by organisational borders. In a wider scope, external nodes have to be considered. In a landscape of risk, such "external nodes" can represent a complete organisation, nodes in this organisation, etc. Structures and dynamics of this supra-network comply with Figure 1.

6) Dimensions of dependencies: Most actions have different contexts with different and potentially conflicting goals and dependencies. These contexts can be structured as a set of dimensions of managerial acting. For example, differences in the place may imply different legal frameworks. Most relevant dimensions are *organisation* (e.g., ownership, responsibility, hierarchy), *space* (the place of operations), '*time*' (aspects of synchronisation an performance) and '*technology*' (ways of performing activities). Further ones may refer to ratings in terms of solvency or product quality. The interdependency of dimensions is based on interactions between degrees of freedom in the domains. For instance, the introduction of containers offered additional degrees of freedom in global distribution, with further impact in other dimensions.

7) Structure determines function is a fundamental paradigm in many sciences. Propagation of impact follows dependencies and degrees of freedom in the different dimensions. For example, "not invented here syndrome" is about a problem that may propagate, e.g., along technical dependencies, but cannot be handled because of lacking responsibility, e.g., an organisational failure. Considering a dimension of different goals, the same event may simultaneously have negative and positive impact. An example may be a traveller arriving earlier at his destination by taking advantage from a delayed train.

B. On the Dynamics of Risk Landscapes

Events and their propagation transform the picture into a movie. In fact, propagation is the motivation behind risk landscapes. The identification and control of paths of propagation are major issues of the design of operations' systems. As an example, failing to avoid non-invented-here behaviour can convert paths into highways of propagation.

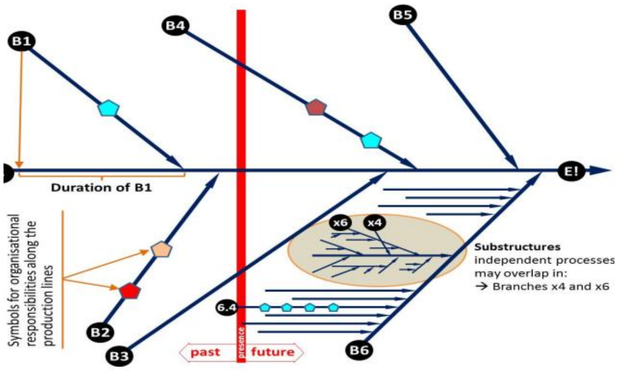

Figure 3. "Fishbone" Diagram for Cause-Effect Analysis (fictitious).

1) Paths of propagation: Figure 3 shows a version of a '*Fishbone Diagram*' [4] of a fictitious factory (for the construction of technical dependencies, the reader should refer to [81]). This diagram is an intuitive way of visualising basic cause-effect relations. A timeline is added, in order to show technological dependencies of a fictitious factory and, depicting the desired synchronization along processes in the work-breakdown. Parallel jobs are organised in different bones and sequential ones along the hierarchy of bones, while products move from left to right. The plant hosts six production lines with details shown for line B6. Each bone implies allocations of resources (materials, tools, and operators). Pentagons indicate different responsibilities in the process. Strategic change will have impact on these allocations, requiring taking measures to free or increase working capital.

A failure to assemble cable brackets in the fuselage of an aircraft can become the reason of a serious interruption of production, making the assembly of kilometres of cables in the next station impossible. Additionally, scheduled operations on the succeeding fuselages are temporarily stopped. Such defects also depend on structural decisions; in the case of the B787 'Dreamliner', Boeing finally decided to buy the suppliers that proved to be unable to solve the problems. The financial losses were tremendous [82].

2) Unpredictability and Non-linearity: A real RL is complex and exhibits an hardly predictable non-linear behaviour [5] that emerges from the number, the variety and the interconnectedness of acting nodes. The number of possible interactions equals to the power set of nodes, i.e., when external dependencies are to be considered, also small firms can generate landscapes of a very high complexity.

Non-linear behaviour appears in spite of well considered operations' standards. A major focus of risk management is to act on exceptions that may become critical in terms of the goals. From this point of view, standards of *firefighting*, so-called policies or best practices, are relevant for managing risk landscapes.

Using the Pareto's Power Law, 20% of unplanned events typically produce other unplanned events with 80% probability, while 4% of unplanned events cause 64% of trouble that may prove to be disruptive and translate into sizeable non-linear effects. Though the concept of the risk landscape is an abstraction, the actual complexity has to be captured. Lacking a comprehensive model, this can only be shown exemplary. Table 1 depicts parameters of events that structure these examples.

C. On Parameters of the Dynamics of Risk Landscapes

The dynamics of the RLs refer to the frequency and impact of change that may occur in different forms.

TABLE I. DESCRIPTIVE PARAMETERS OF UNPLANNED EVENTS

Type (aspect)	**Specifications**		
Type	technological		commercial
Scale	size of business		reach of impact
Decision hierarchy	strategic		tactical/implementing
Competition	slow		dynamic
Interruption	low frequency		high frequency
Disruption	disruptive change		operations' disruption
Knowability (simplified)	known	butterfly	black swan

1) The *type of change* here shall distinguish the technological and the commercial sides of operations. Both aspects, however, are interdependent: adapting to technological change may be inevitable, but the commercial impact finally decides about the sustainability of operations and business.

2) The dynamics depends on the *scale* of the related business (large firms/groups versus small/medium enterprises), simply because size implies a higher number of interdependent nodes and, in direction, a higher resolution of object and time (see Section III-B). However, the most important aspect of dynamics is the scale of *power*; the introduction of standard containers had high-power impact. With even more energy, the growing shareconomy, 3-dimensional printing (3DP), Internet technologies, business platforms and data driven businesses today change the rules in actually all industries worldwide (see Section III-C).

3) The hierarchy of decision making refers to the reach of events. For example, a new technology, like 3-dimensional printing (3DP), has impact on corporate decisions (see Figure 2). However, with the acceleration of the rate of change, the efficiency of hierarchies disappears. Finally, in firefighting scenarios, well managed firms enable low level actors with local knowledge to make decisions, even if they have strategic impact [93].

4) Competition boosts complexity by linking players into feedback circles: Challenges issued by competition need answers that may become an issue for competitors in the future. These cycles run in respective markets fast and resourcefully and imply a competition for innovation and related internal and external financial resources. Not at least it comes with significant risk. The loops can become even more complex. The "cooperate to compete" (coopetition) strategy exploits the competitive intelligence to make a cake together and, then, to compete for its pieces [6].

5) The Interruption of the progress of ongoing operation is the "lowest" level of negative impact of an event. It may occur if a critical resource is missing. For example, in 2010, the eruption of the Eyjafjallajökull volcano in Iceland temporarily stalled airline operations. Non-polluted air was the critical resource, analogously to the missing bracket in the scenario described before. However, an interruption can also leverage local noise to the level of strategic issues.

6) Disruptive change forces to discontinue a way to operate (the 3DP example) or to change or close business models. Interruption is a failure that can be solved, but may spiral into disruption (CargoLifter case). The cause may be a loss of a critical resources, external events, innovation or behavioural change. In the example of a shareconomy, customers deprive operations of valid business models.

7) Knowability refers to the chance to predict a development, that is, to recognise a potential event and its relevance. It is directly related to intelligence and acts as a prerequisite in the case of answering or driving creative destruction. Knowability implies abilities to estimate the expectation value of an event defined as impact multiplied by the event risk (see Sections III-A and IV-A). The question is what could be or can be made known early enough to reasonably act on a risk. 3D-printing is an example for issues that can be known. In this context, Bayesian inferencing can be applied in order to systematically improve such knowledge (see Section IV-B). Two further aspects that need to be taken into account, the Butterfly Effects and the Black Swans, are described below.

Butterfly Effects emerge from the non-linear behaviour of systems; small, hardly discernible causes lead to a large impact in hardly traceable ways. The analysis of large sets of data can support learning and a deeper understanding of the behaviour, i.e., the identification of positive feed-backs that fuel non-linear behaviour and propagation of impact. They are relevant in case of measuring the *criticality* of operations. In organisational context, a phrase in a contract can make a difference.

In contrast, *Black Swans* [7] come from places behind capabilities of imagination, at least for the vast majority of actors. Examples of strategically relevant Black Swans may be a sudden breakthrough in quantum computing or technology providing clean, safe *and* cheap energy. For most actors, the 2007/2010 financial crisis has been such an event.

Examples from direct operations are the problems in the ramp-ups of the Airbus A380 production in 2006-2008 and of the B787 Dreamliner in 2003-2011. In the case of the A380, the highly customized harnessing became a problem because many of the cables were too short. It was beyond the capability of imagination and, thus, of awareness of all actors that different departments worked with different versions of design software.

The B787 case is marked by a long list of various disruptive events across major systems of the whole aircraft that, at least in this accumulation for management, were hardly imaginable. The disasters turned into additional costs of 6.1 Billion $ for Airbus and into an estimated 12 – 18 Billions $ loss for Boeing. For more details, the reader should refer to the "Catalogue of Catastrophe" [8] [9].

Neither a SWOT analysis nor, e.g., a simulation-based analysis will find a Black Swan. A possible solution would involve a systematic effort of *explorative learning*, of encouraging and facilitating creative lateral thinking or of taking advantage from diversity in the teams rather than fostering uniformity. For organisations and for actors, it is an act of balancing the discipline to act in conformity to standards or agreed proceedings on one side and on the other the intelligence of questioning standards and proceedings as potential habitats of Black Swans.

III. PREVIOUS WORK

The concept of risk landscapes, as it is discussed in this paper, goes back to a number of intra-industrial as well as collaborative research projects with industry. They delivered the empiric base of the concept. The major work includes air cargo logistics, selected airport ground operations, inflight catering systems, complex technological ventures, and, currently, small series production in aviation industry and work with a group of small enterprises on innovation strategies.

TABLE II. OVERVIEW OF STUDIES

Interactive Tracking	1995/96 Volkswagen, LH-Cargo
CL Knowledge Integrator	1998/2001 CargoLifter Project
RFID-based intelligent inflight catering	2007/10 Airbus (main partner)
Production Management	2012/15 Airbus, Iacobucci
For Fife SME – innovation strategies	2014/15 Group of SME

Adaptiveness to unplanned events, including the management of related risk and impact, has been the recurrent theme. The very first project became the primer. The idea to integrate virtually any aspect of acting under uncertainty into the concept of risk landscapes and realising the impact of accelerated disruptive innovation emerged from the projects described in this section (see Table II). The use of semantic technology and multi-agent systems was an early choice.

A. *Interactive Tracking*

In 1995, the Strategic Research Team of Lufthansa Cargo and Volkswagen Transport (VWT, the transportation unit of Volkswagen) agreed in a project that analysed methods to improve the response of the factory in Germany to urgent orders for spare parts by satellite factories. The work was done in collaboration with a team from the Technical University of Braunschweig [12] [13].

Satellite factories (in the study located in Mexico and in Johannesburg) assembled Volkswagen cars. Most of the parts were produced in a German factory and, by standard, sent by ship to their destinations, being too slow in case of emergency. To both destinations, the factory-to-factory air-transport time was about one week with five flights offered per week. Thus, in case of required response times to an emergency order of two to three weeks, up to ten flights could be used. The question was *how to exploit this flexibility to improve the flexibility of the customer*.

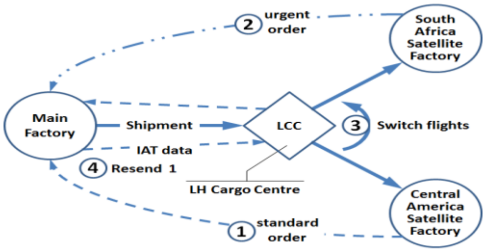

Figure 4. Interactive Tracking, Functional Scheme.

A typical case (Figure 4) is a request for parts ordered by the factory in Mexico. After packing in the main factory, the parts have entered the Lufthansa Cargo transport pipeline. From that moment on and until it arrived at its destination, the shipment became "invisible". Just after the shipment to Mexico "vanished", an order from Johannesburg arrived asking for a similar, but not identical, mix of parts. However, due to the fact that the production programs of the satellites were similar, the same applied to emergency demands. To avoid interruptions in the production line, the shipment should arrive the same week.

Since RFID was not yet available at that time, the idea was to mark relevant shipments using a barcode, in order to make shipments fully traceable for Volkswagen. Allowing for simplification, the process involved labelling shipments in the first Lufthansa station, and, then, scanning and storing there, as well as in any subsequent station, respective data, such as the airway bill, the place and time etc. The aforementioned data were maintained in a disk that accompanied the shipment until it was loaded into the actual flight. Provided that the staff was properly trained and disciplined, this process allowed to find and redirect shipments on order of VWT to the destination with the higher urgency; in Figure 4 to Johannesburg.

The management of operations' risk was not an explicit issue in the project. However, we were aware that logistic systems, able of handling a higher level of detail and a faster response to problems, can significantly increase operations' performance. It was a strategy to manage a "landscape of risk in the nutshell". Together with the Kenan Institute of Private Enterprise (University of North Carolina), the idea was picked up by an international researchers' network [12].

B. *CargoLifter Knowledge Integrator (1995/2002)*

The CargoLifter Project (CL, Germany 1995-2002) aimed at designing, producing and operating airships of a size of 260 m. length, 65 m. width and 82 m. height: a 'flying crane' for transporting goods of up to 160 metric tons up to 8 by 8 by 50 meters in size. The project failed because of its overcomplexity. It included the job of an airframer (design and production), of an airline (flight operations), of a ground infrastructure provider (each parking position of this airship would require about a square kilometre) and of a logistics company specialised on complex special transport projects [14].

To substantiate and justify investments into this project, a major challenge was to intelligently link the task to establish a valid business model and the task to support the acquisition of financial resources on one side with the technological progress of airship development on other side. In order for this to be accomplished, the strategic research team of CL specified the '*Knowledge Integrator*' (KI, realisation by Magenta, London and Prof. G. Rzevski, Open University, Milton Keynes).

A first purpose was to estimate the financial performance of airship operations' networks including airship, infrastructure, customer sites, orders, etc. as nodes. The specification of the market relied on different sources: (1) real data from members of a global group of industrial lead-user about their projects, (2) data from global market and benchmark studies, and (3) even potential competitors like shipping lines, and (4) by the airship engineering team providing estimates parameter values of the overall airship operations' performance as specified. All data were regularly updated.

Thus, the model *integrated* the knowledge available about performance parameters estimated on the base of the progress of technological and operations' design, various market studies and on knowledge inhered in questions of investors, lead-users, banks, authorities, press, etc. Based on that pre-knowledge a Bayesian process was started by specifying and simulating operations scenarios with the goal to deeper understand and improve knowledge a posteriori. In this way, technicians stepwise improved their estimates of the performance of the airship and its impact while share- and stakeholders reviewed their ideas about expected returns or the capital to be bound in development and operations.

Above all, the idea was to systematically grow the KI to the capabilities of a fully-fledged operations' management system dynamically acting on unplanned events. Although the term "risk landscape" was not used at that time, the model fully meets the definition. In the bankruptcy of the CL project, the developed software and the majority of the documentation unfortunately perished.

C. RFID-based Intelligent Catering Systems (2007-2010)

iC-RFID was a strategic research project, funded by the German Federal Ministry for Economic Affairs and supervised by the Program Management group of the German Aerospace Centre [13]. The purpose of this collaboration of five industrial partners and three research institutes was to design and demonstrate functionality and business cases of end-to-end integrated RFID-based inflight catering systems.

CESAR (Configuration and Evaluation of Service Systems in Air-Catering with RFID, [15]) was one of the subprojects, designed and implemented by the research team at Cologne University of Applied Sciences. It was a prototype of a multi-agent system integrating major RFID-enabled functionality of inflight catering, novel service models and further technological innovation. The model included the specification of catering services by airlines (food, beverages, sales items etc.), production and packing service content by air-caterers, selected airport ground operations, ground transport and exchange services (highloader trucks for unloading and loading aircraft galleys) and main aspects of the rotation of service equipment, such as trolleys. At the same time, the flights in the airline networks, standards of aircraft producers and of particularities of aircraft interiors, such as galley layouts and equipment, as well as legally enforced regulations were considered.

The MAS assisted the management to maintain service levels in case of unplanned events by analysing discretions to act of all active nodes in the scene, by proposing solutions, organising the implementation of decisions and tracking the effectiveness of action. However, *CESAR*, a prototype of a context-sensitive real-time operations system, was only able to respond to events that had already occurred and not to the risk of events. Methodology, such as tracking of criticality and evaluating event risks on the basis of behaviour, enabling the uptake of proactive action, was not implemented.

D. Adaptive Production Management (2012-2015)

ARUM is a collaborative project with 14 partners. The ARUM project concentrates on two use-cases from the aviation industry about the production of aircrafts and aircraft interiors. ARUM is co-funded by the European Union in the 7th Research Framework Program (GA 314056) [15]. One of the use cases focuses on the *ramp-up* of production, which is one of the most critical phases in the life-cycle of a complex product like aircrafts. Stories about the Airbus A380 or the Boeing B787 'Dreamliner' are known [8] [9]. These ramp-ups are marked by possibly fatal problems of technical maturity of components and processes as well as by poor learning curves as a consequence of the small series in aviation industry. ARUM provides MAS-based planning and scheduling systems that capture and process unplanned events in large scale and are prepared for risk-sensitive methods.

Airlines, as the final customers and operators of aircrafts, expect innovation that saves costs or improves services. The trend is to pack them into refurbishments of existing programs instead of ramping-up all new options with the start of a new aircraft program. Thus, ramp-up scenarios will also appear during the lifecycle. In this context, colleagues of the National Technical University of Athens (co-authors of this paper), as well as from Manchester University, Certicon, Prague and CUAS among others coordinate the development of semantic models (ontologies) [1] [16] [89].

In this work, it became clear to us that accelerating change is *the* pervasive force driving issues like repetitive ramp-ups, challenges to managerial workflows, further fragmentation of learning curves, and not at least inconsistencies and losses of effectiveness of semantic models. Moreover, it is the source of disruptive change and increasing unpredictability. In collaboration with Almende, Rotterdam, strategies that enable managerial workflows to keep pace with the quickening of technological change and that effectively support learning strategies will also be implemented.

E. What Matters: Lessons Learned

1) Agility matters in terms of adaptiveness to any change. The IAT project was motivated by ideas of an industrial Agile Management program [17] [18] that focused on impacts of fast action on unplanned events and of a clear dedication to the customer in terms of "We learn and solve your problem!". Regarding the reduction of uncertainty about the progress of a complex project, the Knowledge Integrator was inspired by this idea. It generated verifiable information that reduced risk and, thus, increased value of invested time, knowledge and capital. The iCRFID project targeted the exploitation of RFID and, in direction, of the Internet of Things, for planning and ad-hoc exception management. ARUM is about increasing adaptiveness and reducing uncertainty in the ramp-up of complex product-production systems.

2) Detail and local particularities matter. IAT tackled a very small fraction of VW shipments. However, the expected loss (impact) of any of these problems exceeded by far the efforts. *CESAR* again calculated potential benefits of corrective action. Both solutions intelligently handled details captured on the lowest level of sensing and acting in local operations. A similar approach is taken in the ARUM project.

3) Precision of mapping matters. It was unlikely that a shipment, deviated from Mexico, contained *exactly* the same mix of parts as ordered by the Johannesburg factory. However, it provided an intermediate solution. On the contrary, iCRFID and ARUM do not allow for variety. For instance, to ensure "delivery as promised" under all circumstances, item X, passenger Y or aircraft Z became active nodes represented by software agents that track and manage the way from a storage to the seat of the passenger. In case of failure, the affected agent issued a request for corrective action, e.g., by taking a spare item from the next source capable of solving the problem.

4) Multi-agent technology matters. Depending on the size of the scene, thousands of details may matter. At any time, each can become the critical resource and the change of any parameter may devalue an existing solution. Potentially fatal chains of events (propagation effects) or positive feedback circles that are caused by the behaviour of particular nodes or sub-nodes are additional reasons for stressing the need for capabilities to model and process objects and events on the lowest level of operations and detail. Therefore, in spite of limitations, particularly with regard to their potentially high load of communications, only MASs have proved to fully satisfy requirements to control ongoing operations under the condition that a particular detail here and now has impact on another particular detail and that solutions depend on the discretions to act that are available on that level.

5) Semantics matters, as operations systems can exhibit a high and dynamic heterogeneity of objects, processes, frameworks and terminology. Semantic modelling seems to be the only solution providing the flexibility and adaptability that effectively supports the R.E.A.L. processes.

6) The quality of managerial workflows matters. Effectiveness and efficiency of operations ultimately rely on efforts to systematically adapt and improve managerial workflows. In changing environments, this is a continuous task. This is the core of managerial excellence as defined in the Beste Fabrik program, active since 1995 and based on far more than 1000 industrial case studies in six European countries and run by seven major business schools [19]. In ARUM, we identified concrete needs for improvement as prerequisites of any effective use of intelligent tools. In the CargoLifter project, deficits of management workflows became the major reason of bankruptcy [20].

IV. A FRAMEWORK OF MANAGING RISK LANDSCAPES

A. R.E.A.L. – Realise, Evaluate, Act, Learn

Realise – evaluate – act – learn is a generic logical structure of proceeding that describes the behaviour of managed (active) nodes in RLs. These tasks are not trivial, either on the local level of individually responsible managers or on any level of integration, and particularly not in case of response to unplanned events. Even ideas may be lacking about what actually has happened and how to proceed further. Without question, managerial effectiveness relies on human factors like discipline and readiness to act, high attention with a sense for details *and* the big picture as well as on high communication skills. The poorer the management is, the less likely it is to keep pace with the propagation of events. As a result, the management fails to mitigate drawbacks or to take advantage from upside potentials.

1) Realise: Nobody can act on unperceived events. Perception may fail because of lacking training or lacking attention owed to human shortcomings. An event may not always be recognised and, thus, not communicated. Sensors may fail or be missing, signals may be filtered out or an event may be not properly read or vague in its meaning and need more knowledge to be understood [7].

2) Evaluate: The decision whether or not to act on *identified* unplanned events depends on thresholds of its relevance. In economic contexts, the relevance is expressed by their expectation value, which is the product of event risk and impact, with the event risk p of an unplanned event that has occurred being equal to 1). In regulated environments like aviation or health industry [21] [22], particular classes of events may be relevant by default in order to avoid quality hazards. If events and their contexts are clear, evaluations can be supported by ARUM technology or Big Data applications, and, in standard scenarios, possibly be automated.

3) Act: Acting on unplanned events (re-)establishes planned states by implementing a suitable policy. If there is time and planning capacity the plan can be updated. In some cases, rules or a proven best practice may be applicable. Elsewise it needs "educated improvisation", e.g., of an experienced dispatcher and the hope that it works.

4) Learn: Unplanned events are the reason and the resource for learning. Deep knowledge about a system derives from enduring observation of its behaviour in a large variety of operations' scenes and the review of many failures along R.E.A.L. processes. In case of disruptive change or innovation, contexts of learning may be lost, "old" technological and managerial knowledge may be devalued and new learning curves might start. It is very likely that experimental learning will have to support or even to replace practical experience. Nevertheless, while experimental learning analyses the behaviour of complex but widely known systems like a factory, a supply chain or a service system, explorative learning tries to avoid dependency on current knowledge [23]. It is far more permissive and allows for testing ideas that may be very strange in the eyes of domain experts. Ambiguity and complexity here are resources. De-learning is becoming a topic. The focus shifts to the management of transitions and the identification of re-useable knowledge.

Learning as a continuous effort is the backbone:

- *Operations' Intelligence* is the capability to effectively disambiguate complex scenes in all phases of R.E.A.L.
- *Real-time operations' control*: Many events need immediate action to answer to downsides or upsides.
- *Tracking of effectiveness*: The effectiveness of implemented policies has to be measured and analysed. And new events require further action.
- *Awareness of assumptions* is a core aspect of learning, like in Bayesian experiments that explicitly capture prior and posterior knowledge (see Section V-B).
- *Encouraged and augmented learning*: As failures and "strange ideas" become sources of learning and innovation, a culture needs to be developed, in order to elicit rules and to provide resources particularly for learning from failure and exploration. Organisational structures form the base of effective augmented learning, including the effectiveness of computer-based support, such as simulation programs. Carefully explored and deployed *data-driven business and operations intelligence* are about to become a further opportunity of learning [94] [95] [96].

B. *On Interactions of Forces Driving the Big V and on Related Control Problems*

From containerisation to servitisation and 3DP, the origins of the Big V are interplays of technological, economic and social developments that also drive the phenomenon of acceleration. Modelling interdependent risk needs to conceive and decode the driving forces and their impacts. From this point of view, velocity, variety or variability does not form the problem; their increase is. More importantly, the situation deteriorates with *acceleration* that hardly leaves a chance to accustom to a plateau or a rate of change. These accelerators are inherent to relations between nodes. Namely positive feedbacks, are relevant and, in consequence, resources that enable management to strategically and operationally control accelerating processes. Multiple facets of the accelerators need to be considered in a model.

1) The globalisation of the reach of almost any activity, the increasing informational connectivity of everything via the internet, the abstraction of businesses and operations as well as the competition by digitalisation are the driving forces. They are inseparably intermingled and locked into multiple amplifying feedbacks (Figure 5). Historically, the development is a stepwise facilitation of the exchange of everything by converging technologies; containers for the physical part, the internet for information and data as well as the virtualisation of services and, finally, the digitalisation that adds computability.

In terms of risk landscapes each of the phases boosts:

- the volume, variety and variability of actors (managed nodes in Figure 1),
- the resolution of these objects (sub-nodes, i.e., further detail, things or services), producing more volume, variety and variability that are relevant in terms of goals and risk of acting,
- the resolution of time, i.e., a higher frequency of unplanned events, particularly of two kinds: (a) creative destruction and (b) operational risk.

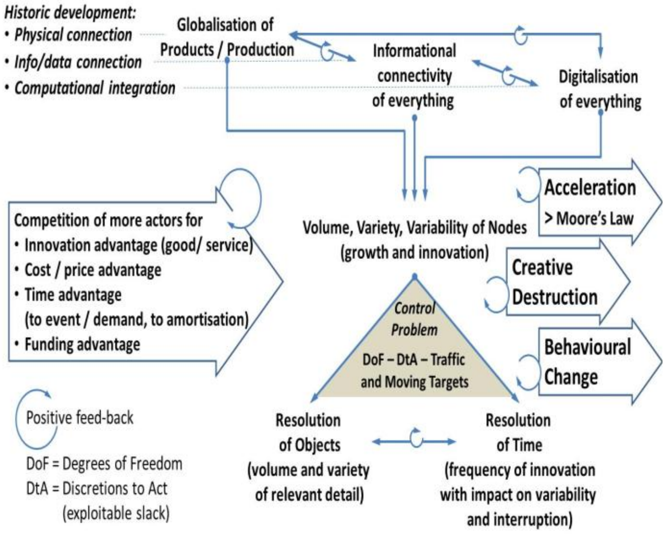

Figure 5. Driving Forces of the Big "V" and of Accelerating Change.

2) For given feeding grounds, competition is driven by the number of competitors, the information available to compete and the capabilities of mining, selecting and processing relevant data. Thus, competition shapes these developments by positive feedbacks. Basically, there are four ways to compete: (1) for the better product, (2) for lower costs and lower prices, (3) for speed of acting, and (4) for the access to capital that is required to pre-finance innovation or to cover risk. The choice of these strategies highly depends on the current conditions and the involved actors are seldom able of playing all combinations.

3) The resolution of detail may increase because technology consists of more parts, but also decrease with new technologies like 3DP. The problem, however, is not the volume but the potential criticality of detail and their managerial impacts [83]. In industry, more detail implies more types of stock keeping units (SKU), more supply chain complexity, and larger stocks, i.e., more working capital.

Thus, financial departments want less, while operations departments like to be on the safe side. Competing for access to capital, the winners are clear. In the IAT, project a VW manager answered a question about benefits: "*We expect a reduction of shipments and volumes of materials in transport. In the long run, we may also be able to reduce inventories*". That is to say, reduced inventory consumes improvements of control and, in consequence, more details become more and more critical.

4) Ultimately, *competition accelerates* because time always matters; the time to amortisation of an investment is crucial to capital cost. Innovation needs timing to the market and, at the same time, it acts as a force of creative destruction and disruption. Response time to demand is a major factor of service quality. Shorter lead and cycle times require to invest into process innovation and form programs eliminating cost drivers, working capital or reducing the time to amortisation, i.e., the time capital is committed to a particular business and therefore under risk [24] (see Figure 9).

Time competition is a primary and acceleration a secondary effect of the forces and the underlying economics of capital turnover. Moore's Law is not the issue; its impact is; hardware capacity doubles every two years, but the richness of exploiting this capacity grows faster. The same applies to behavioural change, e.g., the demand for service.

In consumer markets the demand for technology has been educated by hardware providers, like Apple, grounding their business model on short product life-cycles. Therefore, not only competition and market but also the life of people and turnover time of fashions or trends accelerate and, in the eyes of customers, change the focus of utility.

5) *The Ashby problem*: Asby's Law of Requisite Variety [25] requires controllers to dispose of at least as many degrees of freedom (DoF) as the system that must be controlled can exploit to produce uncontrolled (unplanned) events. It is about the *variety of behaviour* and includes DoF available from different constraints in *different dimensions of acting*.

In this respect, a solution may not become effective in production because of organisational failure, like a non-effective allocation or handling of responsibility. Control relies on effective constraints to behaviour, as well as on effective policies to respond to the upsides or downsides of unconstrained behaviour.

Ashby's subjects of control were technical systems: "*if variation is required* [to control behaviour], *there must be a source of noise* [yet unknown DoF] *or information* [about uncontrolled DoF]" [26]. A strategic scenario could be a search for drivers of unexpected fluctuations of sales. In operations, these drives could involve variations in the quality of supply. Nonetheless, real customer relationships, competitive games, operations' systems and markets are not complicated but complex. They exhibit emergent behaviour and are populated by positive feedbacks that mutate butterflies into gorillas, ideas into creative destruction or responsibility into a "not-invented-here" syndrome. Nothing of this can be reduced to the behaviour of single nodes and strongly dependent on contexts and history.

6) In the "slow motion" environments of Ashby, there was not enough energy in the system differentiating complexity from complicatedness. Nevertheless, Ashby's Law still holds. In order to lock into positive feedbacks and to grow exponentially, there must be nodes in the RL that dispose of appropriate DoF. For their detection to be feasible in time, business intelligence needs to get onto the track of indicators and of potentials to change. It is an issue of operations' intelligence to get onto the track of early butterflies, e.g., by analysing noise.

7) Discretions to Act: Almost any schedule allows to accommodate another appointment. If need be, contact persons of booked appointments may be asked to shift or to wait a little.

8) Multi-agent systems in the CargoLifter or the iCRFID project worked in a similar manner; if a delayed flight blocked two catering trucks in peak time, no truck had a slot available to take over. However, it is often possible to shovel capacity free by managing small shifts across the fleet, and, if necessary, involve further stakeholders. The use of Discretions to Act (DtA) implies using slack in a system as a resource. This is based on two capabilities: The first involves understanding slack as a resource of flexibility. The second one is to identify and effectively exploit the DtA.

However, they cannot be planned, but they can be constrained or used up by scraping the bottom of operations' resources. Optimal slack can only be tuned based on experience and simulation analysing patterns of noise and needs for flexibility. DtA are also hard to track, since they are a volatile resource that disappears (the truck is stuck in a traffic jam) and re-appears because another problem, like a cancelled flight, shovels time free.

Centralised control will barely handle interactions of volatile DtA and dozens further context parameters. Solutions emerge from trading DtA, but they are local and bound to the slices of time available to the individual actors. It needs *peer-to-peer systems* (P2P), i.e., nodes that are aware of their current states and, on that base, coordinate their local, individual decisions.

Collectively they are aware of the integrity of the overall process [27]. Further the nodes can realise and coordinate their exploitation of DoF and their exposure to risk across the network, e.g., in the same way that a car-to-car communication system uses DoF of individual cars for accident prevention and traffic control [28]. So a minor reduction of speed may avoid an hour wasted in a traffic jam.

Vast volumes and velocity of communication is the price to be paid for the advantage of widely autonomous P2P systems to control operations and to deliver indication for strategic decisions. For reasons of comparison, the *CESAR* prototype included about $10^2 – 10^4$ nodes, whereas a realistic full scale model would need about 10^6 nodes.

For a network of factories, the ARUM model may reach the same dimension. The current European air traffic involves about 25 thousand flights per day and it is expected to reach about 43 thousands flights per day by 2030.

C. The Challenge of Accelerated Creative-Destruction

"*The paradigm shift rate (i.e., the overall rate of technical progress) is currently doubling (approximately) every decade; that is, paradigm shift times are halving every decade (and the rate of acceleration is itself growing exponentially). So, the technological progress in the twenty-first century will be equivalent to what would require (in the linear view) on the order of 200 centuries. In contrast, the twentieth century saw about 25 years of progress, since we have been speeding up to current rates. So, the twenty-first century will see almost a thousand times greater technological change than its predecessor*" [29].

J.A. Schumpeter defined innovation as a force "*that incessantly revolutionizes the economic structure from within, incessantly destroying the old one, incessantly creating a new one*" [30]. Investors see here a portfolio of options to commit capital to novel technologies and business models. In IT-driven markets, a system or a model may be characterised as "old" within a year. But what is "old" in an accelerated competition for innovation that vice versa is an accelerating force? What is the risk that options to innovate or to establish a business vanish the next day in global, internet driven competition with brain-to-market times being cut by 3DP? In the digital business of speed, trading the shelf life of information counts in milliseconds and technology to improve response times is implemented as it appears.

In terms of RL, it is crucial that any of the developments discussed in the following paragraphs had or has unescapable disruptive strategic and operations impact. Nonetheless, all of them have been knowables, and, thus, left time to adapt. With Big Data and 3DP, this fact also changes the border between strategic and operations' aspects of RL erodes with further acceleration.

The list of firms that failed to adapt to change is long. As an example, on July 30th 2014 the CEO of Osram, the second largest producer or electric lights in the world, said: "*The whole industry is taken aback about the fast decline of demand for traditional products*" [31]. The whole industry? The source of failure is shortcomings of industrial and operations' intelligence. In the words of Clayton Christensen, "*Outmoded thinking and the tyranny of key performance indicators impede innovation that creates new markets and new jobs*" [32]. A brief overview:

1) In less than 20 years *the Internet* became the global exchange platform for information, data and data-based services allowing to directly or indirectly connect and control virtually any 'thing' by sensors and actors. Aware of planned or actual states, any object can also be directly (via embedded IT) or indirectly (via software agents) become an active node in an RL (Figure 1). Equally, any digitizable service, either complex infrastructure or application, is available as a scalable cloud service. Based on semantic annotation or ontologies, the Semantic Web [33] gives bytes and data a meaning and facilitates internet-based knowledge processing and intelligence. Web technology became the mainstream of semantic modelling, of data filtering and integration as well as of data-based services. The Internet became the largest imaginable agglomeration of data and the delivery room of the "Big V".

2) This data cloud is a bonanza, available for mining and exploitation in science, business, and the public sector or state agencies. As the Big V, Big Data (BD) are children of the Internet. BD capabilities and capacities are unprecedented accelerators *by cutting the time from data to business* by means of scanning, organising and analysing massive volumes and flows of data. BD are breeding grounds of new scientific practices [97], of the scientification of businesses and of new occupational profiles, of new business models – and the doom of others.

3) The actual value of goods is equal to the services they deliver: a car that fails to start is no asset but a problem. The result is a growing shareconomy. Thus, *servitisation* is a strategy to enrich goods by services (e.g., cars by assistant systems) or, above, to sell not the goods but the service they enable, a view that implies that customers learned to evaluate products in a different way: A taxi is a *car as a service*, a capacity shared by taxi passengers, virtual and scalable in the sense of virtual computer capacity.

The number of people sharing instead of buying cars doubled since 2004. The ambitions may be too low, but by 2020 Mercedes plans one Billion Euros of revenue by car sharing (less than 1% of sales in 2012) and by 2030 BMW wants to make more money with data than with cars. Recent studies indicate that in agglomeration areas a shared car replaces up to 32 bought ones. The change of social values changes a global industry.

4) 3-D Printing (additive manufacturing) revolutionizes the production of most material goods including transplantable tissues and organs [34], food [35], fabrics, clothes, toys, dinnerware, buildings [36], and parts for aircrafts [36] [37] [38] [39].

"*The most exciting thing about 3DP is that complexity is free. The printer doesn't care whether it makes the most rudimentary or most complex shape*" [40]. Moreover, lot size equal to one is almost for free, large parts of tooling and a large variety of supply become needless, working capital is remarkably reduced and economies of scale are reduced to learning how to improve and operate printers [41].

Printing as a service will become a mainstream of production, provided close to place and time of demand. In fact, Amazon just announced the launch of a 3DP-Store for customizable goods [42]. More important, *3DP digitalises the way from brain to market*: the essentials are the file coding the design, materials that match functional requirements and a printing device capable of processing the materials in the right time and with the right quality.

Anybody able to create a new design or functionality the proceeding will be to experiment, learn, optimize, and sell. Additionally, times from science to business become shorter: printable materials with new properties or printers with new capability will immediately change options of design and production.

V. SEMANTIC MODELLING OF RISK LANDSCAPES

The scale of problems of "integrating everything" into a processible model requires an ambitious semantic effort, the more as the concept of risk landscapes does not reduce the heterogeneity or variability of operations and the need for managerial knowledge. It actually reduces the semantics to describe and analyse risk and its propagation to a few concepts: event risk, impact and the expectation value of events that controls the relevance of acting.

Risk-effectiveness relies on the performance to coordinate decisions across the network, and not on the effectiveness of nodes to meet local goals. Event risk and relevance are landmarks that help to navigate and to coordinate acting in the landscape. This section describes basic concepts of semantic technology, the potential parallels to observations and discussions in geosciences, and questions deriving from the unprecedented increase of the Big V and of the accelerating speed of change.

A. On the Subject of Semantic Models

Does *this* sentence have a meaning? Not on its own resources but if this paragraph provides it with a context. What is the meaning of water dropping on your head? Outdoors, it may be '*it rains*', whereas, indoors, it may imply that '*the pipe leaks*'. In the same respect, what is a lightning? Members of a primitive tribe may take it as an omen. In an invitation, the term '*ball'* may get its meaning by the request to wear a ball gown or a soccer jersey. Obviously, meaning lies neither in the terms nor in observations (perceived signals). It is in the context, that is, the network of associations on our disposal to accommodate the sentence, the drop or the lightning, in a non-ambiguous way that "makes sense" rather than irritates. In this respect, different observers of a lightning may use the same term but may not share the same interpretation, because they do not share the same knowledge model.

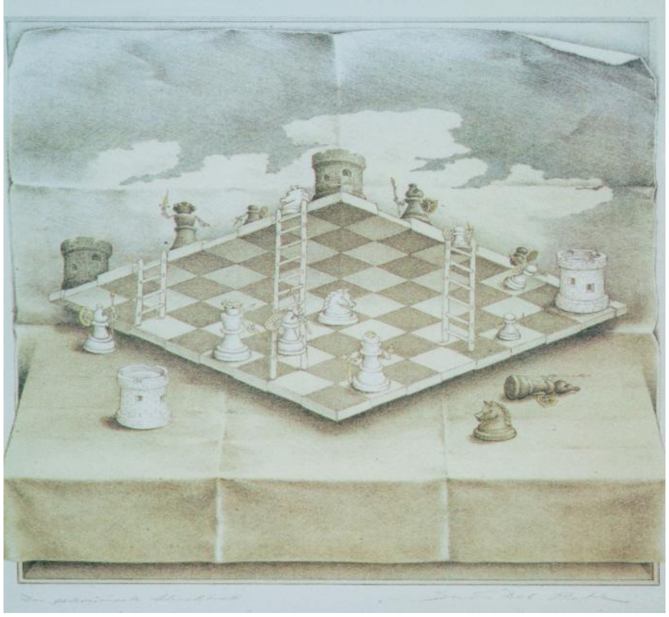

Figure 6. Sandro Del Prete: The curved chessboard [43].

A short insertion clarifying the use of terminology in this paper may be useful:

- *Perceived signs or signals* of any type indicate the arrival of an event and need interpretation.
- *Semantics* is the meaning of signs that emerges from associations to other meanings (context).
- *Concepts* are disambiguated meanings, i.e., have clear associations to other concepts. Conceptualisation is the process to disambiguate a network of associations.
- *Terms* are names for concepts. A definition is the disambiguation on the level of names.
- *Knowledge* is the network of concepts that enables an actor to act in a real or in an imagined environment. Sharing of knowledge requires a sharing of meanings and of terminology.
- *Intelligence* (also business or operations' intelligence) is the performance of adapting knowledge and new events to each other and to maintain or improve the consistency and to reduce ambiguity of knowledge.
- *Ontology* is a conceptualisation of being, pragmatically a model of knowledge about a domain that can be shared, based on an agreed terminology with an adequate precision and consistency.

Figure 6 illustrates, however, that the proceeding of perception, disambiguation, and acting may fail. Signs may have different meanings, qualitative data may tenaciously resist structuring and workable data may not be available.

This is at the heart of the management of RL that is not made from figures, but from the meaning of figures, events, or stories (narratives) that provide context [93]: why bought hedgefonds bonds of a bankrupt Argentina? How meaningful is the definition of a car as "immovable property"? Which knowledge is shared by the Osram CEO with '*the whole industry'*? How can we identify, and what is the meaning of a Butterfly or a Black Swan? All this is about contexts.

In terms of semantics: hedgefonds have a different ontology and value risk differently because they have a different business model, young urban citizens and traditional customers see different values in the same issues. The reasoning of Osram management was mislead by knowledge that did not match reality.

There is no value-free knowledge because any step in the R.E.A.L. process implies decisions and each decision relies on *currently available pre*-knowledge, that is preliminary knowledge built from assumptions and values motivating initial ideas that prime the proceeding and questions, e.g., to be answered by a Big Data application.

Butterfly effects or Black Swans frequently are not realised because they are beyond perceptions and beyond imagination or because they are in conflict to our fundamental assumptions and to values that organise our world. Therefore, the educated gut feeling of an experienced operator frequently performs better than any system.

In comparison with the CargoLifter KI program, the idea is to systematically mature the model by unifying the formulation of questions, growing the awareness of underlying assumptions, clarifying the terminology, understanding interdependencies of the domains involved and agreeing in a common ontology that enables effective collaboration. This is not far from the challenge to model and capture information in an RL. In contrast, environments marked by accelerated change and an increasing resolution of object and time imply repeating efforts of semantic re-engineering with phases of high or of greatly reduced matureness.

B. Potential Reference Strategies from Geosciences

In the first version of this paper [1], we agreed to use a model from geosciences as a reference. Lacking experience with modelling and processing complex risk landscapes it looks plausible to compare them with weather models of a continental scale. Interestingly, looking for examples for modelling and processing of very large volumes of heterogeneous data, we found relevant work in geosciences.

Managers must decide with the knowledge they have, which is an aspect that justifies a short review of Bayesian models below. Not making a decision is the worst decision not at least because it divests of learning and in practice not "the" optima, but because best possible solutions are achieved. Kristine Asch describes this problem on the example of creating "*a harmonised geological dataset for the whole of Europe and adjacent areas*". To get an idea of the multiplicity of constraints of managing RLs, it would be useful to refer to [44].

Though they differ in terms of resolution or dynamics (probably except for the case of high-energy atmospheric processes), there is indication that geosciences and risk landscapes face similar requirements regarding semantic methodology. Both include large sets of numeric data and extensive non-numeric content including vague concepts or narratives to be disambiguated and formalised, because indication and evaluation of change often does not lie in the figures but in their interpretation.

The distribution of work of geo-scientists and managers in risk landscapes need interdisciplinary collaboration and depend on individual perception. Both need modelling to maintain the connectivity to lowest levels of possible data sources as well as to allow for the largest possible variety and the lowest practicable degree of reduction. This process needs terminologies that not unify but align variety.

In slow-motion environments, standards make work easier. But, this can hide indicators of change. Under conditions of accelerated change, the tolerance to some ambiguity may be useful. Nonetheless, interoperability is paramount in both areas. Upper ontologies serve as cross-mapping hubs for satellite ontologies that hold more specific local concepts. Reducing the ontological problem to a terminological one, a systematic restriction of interpretation is a way to reduce confusion [45].

Based on an example of the ARUM project, the specification of an element in a risk landscape could be 'bracket-is_a-resource' completed by 'resource-has-dimensions' and 'owner-is_a-dimension', 'bracket-has-owner', etc. Each step narrows the space of possible interpretations. The proceeding is handy and compatible with modelling RLs. It also allows for playing with constraints, e.g., for explorative learning or testing of alternative modelling strategies.

C. Allocating Semantic Modelling in the Organisation

For the semantic engineering of large risk landscapes, this suggests to establish networks of actors (nodes), gathering and modelling local data according to principles discussed above, including a sense for deviances from expected behaviour (contradictions to a hypothesis), a training issue. In practice, the task should be allocated to knowledge management departments.

Modern knowledge management departments are service centres offering products customized to corporate or operations' strategies and tasks. They capture and disseminate tacit and explicit knowledge, as well as encourage and facilitate the direct exchange in the organisation, e.g., via social media, linked data infrastructures, or personnel rotation programs and allocations to projects requiring collaboration of departments. They are engaged in formalising knowledge, in related internet projects, in content management and quality management programs for KM-services, content or effectiveness of knowledge. KM is involved into the support of Big Data applications, e.g., by developing semantic filters for heterogeneous mass data. Thus, KM is the most important candidate to implement n semantic infrastructure for operations' intelligence and managing RLs. Thus, the goal is that well educated knowledge managers assist operations' line managers and feed the model.

D. Open Questions – some Examples

In the following, we aim at exemplifying a few topics that may need further discussion. One of these issues is that risk landscapes, to a large degree, are concerned with human behaviour of perceiving, evaluating and modelling human behaviour. Problems of biases or path dependencies ("history matters") of thinking, modelling or acting may serve as examples.

1) "Typical terms" are explained by narratives rather than defined formally. They typically populate social, economic or political domains that widely escape from pure formal descriptions [46]. Lieto and Frixione identify the problem in the formal constraints of description logic underlying the formalism of ontologies. *"As far as typical information is concerned, such formalisms offer only two possibilities: Representing it by resorting to tricks or ad hoc solutions, or, alternatively, ignoring it."* Both, obviously, cannot provide a satisfactory solution.

What is more, the concept of constraints does not work if the problem is not in lacking attributes but in the issue. For example, try to disambiguate the attributes '*motivation'*, '*curiosity*', '*sharp-eyed*', '*commitment*', '*decidedness*' in a formal way (proposed by a senior advisor of a recruitment agency to select chief executives). The separation relies on the narratives, also called "case studies".

The content and particularly the interpretation of narratives, case-studies or attributes are case- and context-sensitive. For example, try to imagine the meaning of the CEO attributes before and after the financial crisis. In the classes, we experience that stories are often differently interpreted. The example about the hedgefonds and Argentina suggests that yet the term 'crisis' can have a negative and a positive connotation.

2) Acceleration, a main topic of this paper, also challenges knowledge management; spill-over learning, is about identifying that knowledge has the potential to "survive" disruptive change. This has already been an issue in high-tech domains. What will become an issue of de-learning when combustion engines of cars are replaced by electrical ones? The about 140 components and related manufacturing techniques, tools, machines, etc. and knowledge may outlast this in the museums or in shops for classic cars. If events of that format happen more frequently, the semantic aspect connected to this is not only to identify perishing and new knowledge. Above all, it is about the paradigm of modelling.

General concepts like 'endurant' (a time-independent observable) and 'perdurant' (only observable over a span of time) or 'universal' (generic term, a pure concept) and 'particular' (subtopic or individual in a generic class) are terms on the highest level of abstraction in ontologies. In case of fast change, the meaning of such concepts may become vague.

3) Is there a semantic support of creativity? In the ARUM project, we made a first approach to construct a simple ontology for the engineering of paper cups based on a matrix of problems to be solved (e.g., subject of innovation), previous solutions (probably outdated) and the attributes, the only category that survived change. As constituents of objects, attributes establish relationships between old and new types of cups but also to further objects, e.g., to the type of coffee shop or drink.

4) This implies that featureless objects (a) don't support relationships, (b) cannot be real in terms of '*perceivable*' and (c) that the perceptiveness and relevance of real objects is in their attributes and the relations they support. This interpretation fits to the definition of knowledge as a network of associations and concepts as the population of these networks.

Formally, "property-is_a-endurant" and "thing-is_a-perdurant" are two statements that hold. Properties also are 'universals' that need instances ('Green' is_constituent_of 'green leaf'). The main point of taking attributes as primary and objects as secondary elements in a conceptual hierarchy is the idea that there are relationships that are not between objects.

This is a radical constructivist approach: objects are constructs of perceived qualities, such as '*green'*. Since attributes are also source of relationships, this approach may pave a path to find new options; there is a potential to augment creativity.

Admittedly far from the topic of this paper, but probably close in terms of semantic modelling, a kind of this problem also appears in physics. The deeper physicists drill into the particularities of the micro-cosmos, the more the properties used to describe an object of interest (electron) dissolve the meaningfulness of these objects and, even beyond, the concept of "object" in general. What is more, an electron is a bundle of properties and the object character of an electron under some conditions renders useless.

"*Today more and more people think that not things are relevant categories but the structures between them. This so-called 'structural realism' is a far more radical break with any conventional atomistic conceptualisation of the material world than any variant of ontologies built on particles and fields*" [47].

This is not the place to discuss structural realism theory. It shall just shed light on the point that turning the ontological pyramid upside down, although in a different context, may be worth a deeper analysis.

VI. KNOWLEDGE MODELS OF RISK LANDSCAPES

A. *The Knowledge Base of Managing Unplanned Events*

"*An ontology is a specification of a conceptualization.*" with the further explanation: "*A conceptualization is an abstract, simplified view of the world that we wish to represent for some purpose. Every knowledge base, knowledge-based system, or knowledge-level agent is committed to some conceptualization, explicitly or implicitly*" [48]. This section elaborates on the knowledge base about risk related to unplanned events and reduces the scope to *Risk Management* as described in ISO 31000 (2009), to the task of managing negative or positive impacts of events under uncertainty [48].

Time passes in Figure 7 from left to right. The cones represent the universe of past and future events: Given a maximum speed of propagation, events outside delimitations can neither be causes nor effects of the event in the middle.

Those to the left are causes and those to the right effects of the one in the centre. The right cone is the one of management answering to unplanned events. The left cone refers to the responsibility of analysts who, based on history data or simulation, focus on causes of events or paths of propagation and deliver pre-knowledge for the estimation of risk.

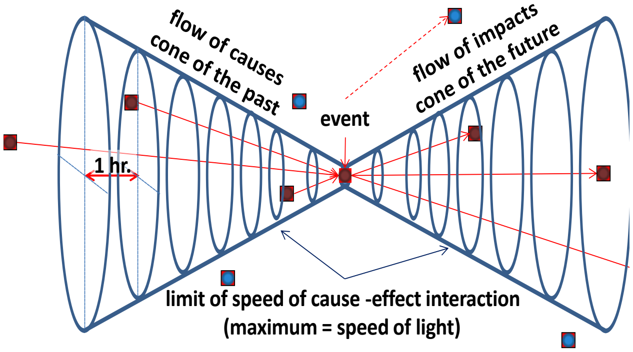

Figure 7. Cone of Cause-Effect Relations and Propagation.

In complex dynamic systems the problem is that it can be very hard, if not impossible, to trace causes or to repeat scenes in simulations. Here we therefore leave analysis to the back-office and focus on the job of managers. Certainly causes matter. But when an unplanned event has arrived, managers are required to appropriately handle a few basic parameters to understand its relevance as well as to act, that is to plan events that recover the situation:

- The Event Risk *(ER)* is equal to a stochastically or statistically defined probability *P*, with $0 \leq P \leq 1$, where 1 and 0 represent certainty of occurrence or non-occurrence.
- The positive or negative (monetary) impact (*I*) of an event is experienced by at least one victim or beneficiary.
- We added the parameter of *awareness* (*A*) to the model as a prerequisite of managerial acting. For example, a competitor's attack may become aware too late. The factor of awareness depends on factors that may antagonize managerial effort, like implicitness, ambiguity, ignorance, taboos, hubris, "*unknowables*" or "*unknown knowables*" [49].
- The *relevance* of events is equal to the *Economic Expectation Value* $eEV = P{*}I{*}A$.

Risk landscapes are "*the set of all (possible) events in the managerial universe*" [50] developing from interacting risks. So the value of *P* may be a function of other incidents: the risk of a denial of service attack depends on the probability to hijack a sufficient number of computers. Forward chaining of events is represented as a risk of transit and impact may be mitigated or eliminated by other events (consider noise cancellation) or meet a well prepared victim. In turn, a beneficiary of a lottery may not be impressed by the prize.

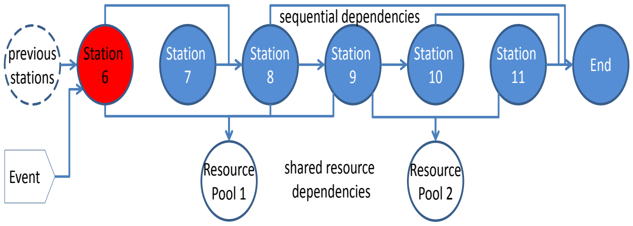

Figure 8. A Simple Pert-diagram of Dependencies.

Unplanned events with a serious impact switch the mode of acting of management from “handle planned operations” into “recover planned status of operations”. The background lies in the arguments of the formula $eEV = P{*}I{*}A$:

If a substantial economic expectation value would have been identified in previous planning, a proficient management would have planned for that contingency. If P or I are undervalued, it arrives as unplanned event, and if A is close to 0, the event may be an *unknown knowable*.

In the simplest case, RLs connect work-stations along technical dependencies as indicated in Figure 3. But not all stations may be directly connected. A little more complicated model is shown in the Pert diagram [51] in Figure 8: A failure in station 6 may affect station 7 by stopping work in station 8 and shared resources may open another path of propagation. A policy is a plan with the intent to reduce downsides or catch an upside of the related unplanned event, both calculated from the eEV of the triggering event with P=1 and A=1 but still with an impact to be validated.

The evaluation follows the ways of propagation and, therefore, is calculated by the target stations until propagation is stopped. Considering first-level effects only, the impact of an unplanned event in station 6 is equal to the sum of impacts in stations 7 to 11 and consequential idleness of resources. As a process, a policy employs resources and has costs. These resources may be (a) implementational ones like cable-ties that temporarily replace proper brackets for cables, or (b) contingency buffers that may even have been allocated to another purpose. Lifetimes of policies are either limited to the time it needs to find and implement a new or recovering the initial solution, or until a new event asks to change the policy.

B. A Bayesian Model for Estimating Risk

1) The Knowledge Integrator in the Cargo-Lifter project was an experiment in Bayesian inferencing with the ultimate goal to provide investors with a distribution of probabilities about the flow of returns on invest (in terms of capital commitment: the time to amortization) in dependency on technological and organisational progress of work in the project. Simplified, the task is to estimate the likelihood that a dice used in a game is ideal, (analoguously: the project is promising for investors) if it shows in 60 rounds 30 times a 6, i.e., a strongly left-skewed (Figure 9).

The KI, a MAS-based simulator, was used to evaluate alternative strategies for operations’ scenarios: What is the economic value of the strategy OS for market OM and an operations’ performance OP? Strategies here stand for the managerial options to exploit market potentials, e.g., in terms of the profitability of operations. OM, the market environment, was a distribution of market models that included information about volumes or price elasticities, and, particularly, studies with real industrial data. With this high quality of inputs, the distribution of OM was considered to be given.

OP, the performance of the airship in this market, was estimated on the base of technological progress regarding functionality offered by the airship (AF) and related ground (GI) and air infrastructures (AI) of airship operations. An example for GI is the efforts required to exchange load with industrial locations or with sea and river vessels, and for AI, the constraints of an airship operations certificate (issued by air authorities).

If the quality level of the specification of AF, GI and AI is equal to the quality of the market scenarios (OM), Bayesian inference would be able to provide a model that answers the questions about returns on investment depending on the progress of the project. But, although the teams improved, there was finally not enough time left to generate a sufficient quality of estimates, because underlying technological knowledge could not be built fast enough to maintain trust.

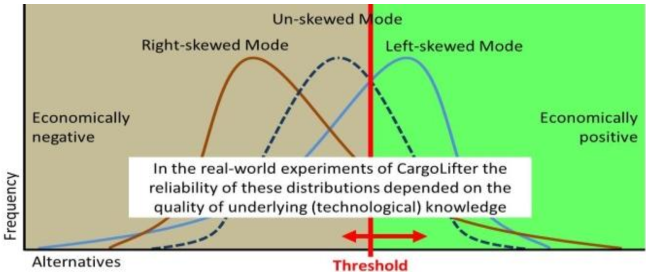

In the real-world experiments of CargoLifter the reliability of these distributions depended on the quality of underlying (technological) knowledge

Figure 9. Economic Evaluation of Statistical Modes.

In general, this may be typical for the very early stages of a complex venture and it is not questioned that such a project needs to go through an even painful period of learning. But exactly this needs to build confidence of investors, lead users and the public and not at least self-confidence of the developers about the quality of work to overcome uncertainty: It is to be expected that a traceable path of milestones and lessons learned is a valid strategy to achieve a positive economic expectation value (Figure 9).

This proceeding delivers concrete managerial options [90][91]: If the project performs as expected or even better, a path can be continued and new promising results may justify increasing investment. Reversely, poor or negative progress in overcoming uncertainty, justifies abandoning at least a branch of the path. Finally, if there is a reason to wait for further information, decisions may be postponed for a certain time.

In this early phase of development, the managerial experiments are identical with the number and the distribution of executed managerial options. Thus, although the example of the dice is undercomplex, the analogy holds: A distribution with positive results should have a negative skewness, i.e., the decisions should cluster around a value ≥ 4, a strong 6 may justify increasing investment, and a positive (economically negative) skewness of the distribution, i.e., clustering continuously around a value ≤ 3, is a reason to abandon.

The intelligence to turn this into a self-stabilizing process is a core element of the so-called managerial excellence [19]. Executable real (managerial) options are the major tool. In the ADVENTURES project, the respective theoretical model was elaborated by scientists of the Otto Beisheim School of Management (Germany), INSEAD (France), The Wharton School (USA), and in collaboration with the strategic research team of the CargoLifter project [92]. In a sub-project, a first landscape of risk of the venture had been also elaborated.

Real-options models allow translating the problem of a complex venture into the concept of a landscape of risk and the logic of Bayesian inference. It does not mean that any sufficient quality of "priors" is given in the very beginning but that there is a structured proceeding to control the development of that quality and to make reasonable decisions, i.e., to specify thresholds (e.g., minimum progress or maximum failure, Figure 9) and, accordingly, to execute options.

2) Modelling a RL implies learning and a basic quality of data: The examples of the CL-case and of ramp-ups in aviation industry as addressed in the ARUM project illustrate the challenges of building models that are consistent in terms of their data and the semantic model.

In the CL case, it took about two years to get a first, consistent database and a first ontology spanning across the needs of technical engineering as well as of operations and market engineering. In this process we also learned how the development of consistent data depends on a consistent ontology.

The fact that it needed two to achieve this state became an indicator of the weaknesses of the project. Actually it failed because of the unability to structure – within thresholds of time set by the share- and stakeholders – a basic RL for expected outcomes: sufficient returns on invest. Clearly, the capability to learn became the ultimate limitation to the project.

In contrast, the Airbus A350 program does not start from the scratch at all but builds on a large experience of aircraft construction. Nevertheless, as the examples of the introduction of carbon-fibre technology or of lithium-ion batteries show, the inclination of the learning curve is the paramount parameter of success.

Comparing, a simple internet research delivers many examples indicating that 3D-printing will accelerate the speed of innovation in aviation industry (See I-C-7 and [8][9]). Thus, the ontology has to support a frequent and potentially disruptive change.

3) *A main aspect of the proceeding* is the complexity of the semantic model. This does not imply that all nodes share all their semantics. But collaboration needs a shared core of semantics (Figure 10). It includes the option that a core-ontology of landscapes of risk fits into a larger context. For instance, a core-ontology for production und ramp-up conditions has been elaborated in the ARUM project [80]. Considering that unplanned events are the main issue of risk management, an ontology for events is elaborated in the following section. To maintain maximum compatibility, this ontology is designed according to the standards of the semantic web.

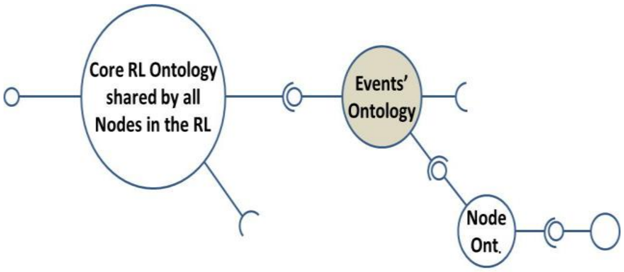

Figure 10. Network of Ontologies.

4) *No decision is a decision*: Managers, judges and doctors must make decisions, thus, almost inevitably, start from incomplete data and prior knowledge and employ methods that can ground acting on a minimum of plausibility (a state of a best practice, etc.) including complex estimates about coupled and conditioned event risk. They learn from experiment and from (even fatal) failures.

Therefore the ontology needs to support the Bayesian logic of "*inverse probability*" (inferential statistics, a term in early references to Bayes): "*... in practice one can check the dependence on prior distributions by a sensitivity analysis: comparing posterior inferences under different reasonable choices of prior distribution (and, for that matter, different reasonable choices of probability models for data)*" [52].

Section C provides a first model of an ontology, designed according to the W3C standards and aiming to match the requirements of modelling a landscape of risk.

C. A Formal Ontology for Events' Management in an RL

In order to capture the aforementioned concepts, the ontological model shown in Figure 11 was created. The actual event is represented by an individual of the *Event* class and it is linked to the appropriate *EventType* individual through the *hasEventType* property.

The purpose of the *EventType* class is to semantically describe an event. The *triggers* object property enables the expression of the propagation of an *Event*, that is, the occasion when an event causes the triggering of further events. An *Event* is associated with multiple datatype properties.

Namely, the *hasRelevanceValue* property denotes the relevance value of the event, which is compared to the Relevance Threshold (*hasRelevanceThreshold*) in order for responsible roles to decide whether this event has to be handled.

The *hasRisk* datatype property reflects the probability of the event being triggered, whereas the *hasImpactValue* denotes the monetary value of the inflicted impact. Finally, the *raisedAtTime* and *includesComment* properties represent the time the event was raised and any additional comments on the event, respectively. Further datatype properties can be defined and associated with an event, in order to capture all the required information for an event of a specific event type.

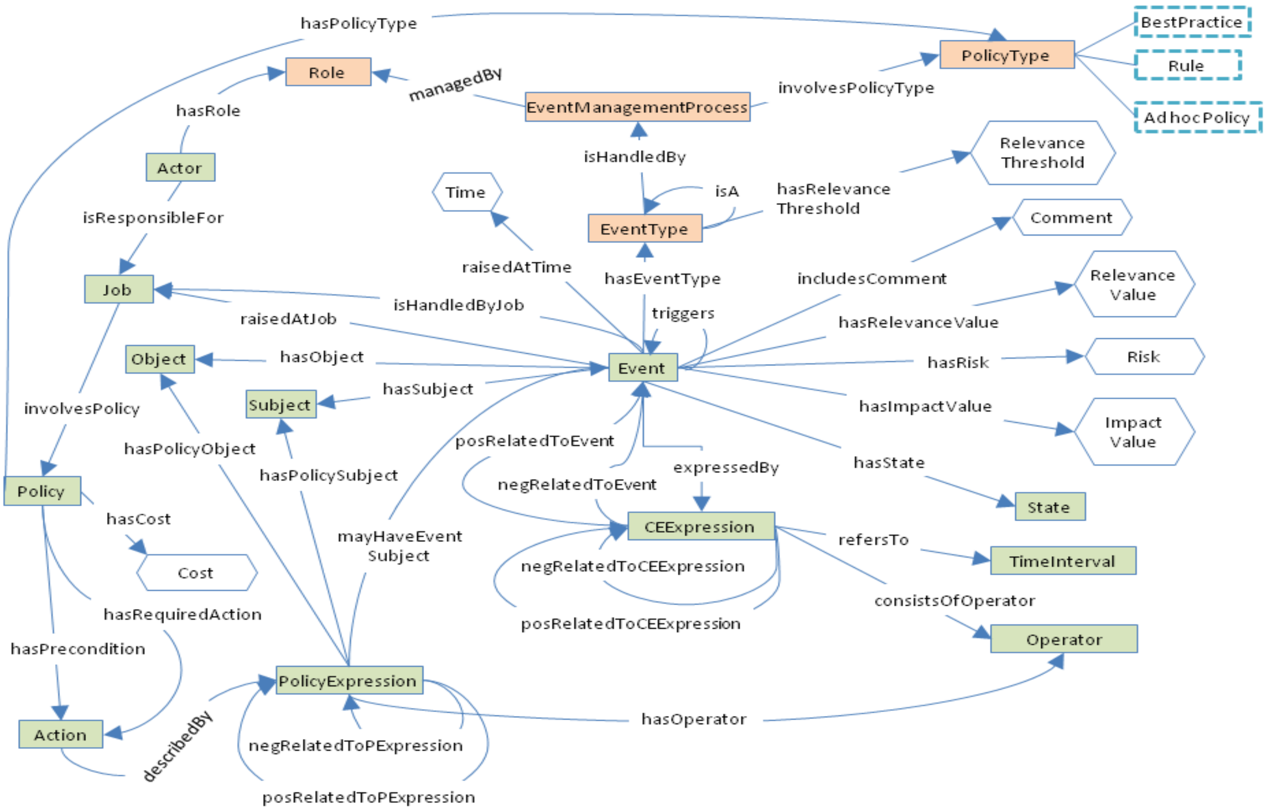

Figure 11. Events Ontology.

An *Event* is associated with the *Subject* class through the *hasSubject* object property, this way expressing the individual that caused the triggering of the event. Apart from a subject, an event may also involve an *Object*, that is, an instance affected by the event. The individuals of the *Object* class are associated to an *Event* via the *hasObject* object property.

The *Job* refers to the atomic task of a scheduled process, during which the event was raised. This association is reflected by the *raisedAtJob* object property.

Through the Job instance, the context of the event can be obtained, following the relations between instances of the class of the Core Ontology that has been developed in the ARUM project [16].

An individual of the *EventManagementProcess* class denotes the process that has to take place in order for the event to be managed effectively. An individual of this class is associated with an individual of the *EventType* class via the *isHandledBy* object property.

The *EventManagementProcess* may serve as a specification of the aforementioned *Job* class. Additionally, the properties *managedBy* and *involvesPolicyType* are defined, linking an Event Management Process to the responsible organisational *Role* and to the appropriate individual(s) of the *PolicyType* class.

The *Role class* is associated with the class EventManagementProcess, depicting the organisational roles that are responsible for the management of an event of a certain type. The specific actor bearing a role is modelled by the *Actor* class, which is associated to the *Role* class through the *hasRole* property. Additionally, the object property *isHandledByJob* is defined, in order to link an *Event* with the scheduled event handling *Job*.

Additionally, an event is associated with an instance of the *State* class through the *hasState* property, denoting if the event is being handled at the moment (Active State) or if it has already been handled (Inactive State).

An Event Management Process involves policies that either mitigate the impact of the event or take advantage of its positive outcome.

An individual of the *Job* class is associated with the *Policy* class through the *involvesPolicy* property, in order to depict the policies that are applied for handling an event. In this case, though, the aforementioned individual of the *Job* class needs to be specified as an Event Handling Job.

A *Policy* can be associated with its *cost* through the datatype property *hasCost* and with a *PolicyType* via the *hasPolicyType* object property.

An individual of the PolicyType class serves as the semantic description of an individual of the Policy class. A PolicyType may fall into three categories:

a. a Best Practice, referring to a policy that has already been documented but not in an identical context, rather than a similar one,
b. a Rule, which reflects a policy that has been applied in the past in an identical situation,
c. an Ad hoc Policy, which is applied in case of no prior knowledge.

The *Action* class is used to express the precondition(s) and the required action(s) that form a *Policy*, in an “if-then” structure. An individual of the *Action* class is linked to an individual of the *Policy* class through the *hasPrecondition* and *hasRequiredAction* properties.

An individual of the *PolicyExpression* class is used to describe an *Action*. Specifically, it is formed by a *Subject* (*hasPolicySubject*) and an *Object* (*hasPolicyObject*), which point to any ARUM Core Ontology class. In the case of the *Subject*, it may also refer to an individual of the *Event* class via the *mayHaveEventSubject* property.

Finally, a *PolicyExpression* is associated with an individual of the *Operation* class through the *hasOperator* property, in order to form a complex Policy Expression. In this case, multiple individuals of the *PolicyExpression* class are linked together via the *posRelatedToPExpression* and *negRelatedToPExpression* object properties.

The purpose of the *CEExpression* class is to enable the modelling of complex events in the form of expressions. The individuals of the *CEExpression* class are associated with an instance of the *Operator* class through the *consistsOfOperator* property, one or more instances of the *Event* class through the *posRelatedToEvent* and *negRelatedToEvent* properties and zero or more individuals of the same class through the *posRelatedToCEExpression* and *negRelatedToCEExpression* properties. For example, considering the atomic events A, B and C, complex events (A AND B) and ((A AND B) OR C) can be modelled.

An instance of the *TimeInterval* class is associated with a *CEExpression* instance via the *refersTo* property, in order to depict the time difference between the triggering of two events, atomic or complex. Individuals of the *Operator* class serve as parts of a *CEExpression* or a *PolicyExpression*. Examples of individuals are the AND, OR, XOR Operators.

D. *Ontology Service*

Once an event is triggered, decision makers needs to handle it, by exploiting every available piece of information. The role of the Ontology Service is to provide access to the Events Ontology presented in the previous section.

The Ontology service works on two different levels: (a) providing a set of Java libraries to be used as an API to access the attributes of the created objects, as well as their associations with other objects, and (b) as a service, receiving requests from other services in the form of messages. The functionality offered by the Ontology Service is presented in the rest of the section.

One of the responsibilities of the Ontology Service is to provide access to the Events Ontology described in the previous section. The access to the ontological data is required for multiple purposes. Namely, upon the triggering of an event, the actual event, along with information relevant to the context of the event, needs to be stored.

The storage of ontological data is achieved by an internal triple store, a special type of a database, maintaining information in the form of subject-predicate-object triples. It has to be noted that events maintained in the Ontology Service triple store may be delivered by sources maintained by legacy systems, by sensors installed within the factory, introduced by actors in the shop-floor of the factory, etc.

Apart from storing events, the Ontology Service offers the capability of performing queries to the semantic data. The queries can be either applied by invoking the appropriate method offered by the Ontology Service API or it can be expressed in SPARQL [98], which is a query language for RDF. Queries may be performed in the context of retrieving past events with specific characteristics or, in general, events that were raised during a predefined time period, within a specific context.

Based on the semantic type of the event, the Ontology Service is responsible for providing the decision maker with the information regarding whether the event has to be handled or not, by accessing the appropriate value of the relevance threshold and comparing it with the corresponding relevance value, which are reflected in the Events Ontology by means of data type properties.

Then, in the case where the unplanned event needs to be handled, the involved parties have to examine similar events that are logged in the internal triple store, as well as events that were triggered due to the initial event and, finally, infer the probability of them being triggered again. The afore described procedure is accomplished by invoking the appropriate methods of the Ontology Service.

Finally, the Ontology Service offers the functionality of accessing policies, by applying criteria, such as the policy type to be applied on an event of a given type. Furthermore, the most effective policy can be retrieved, by performing comparisons between the attributes of the event under consideration and identical events that were triggered in the past or similar ones that were raised in an identical context. This enables decision makers to select the appropriate policy, based on existing knowledge. If such an event has not been raised again in the past, a new policy can be designed, by following the structure defined by the Events Ontology.

The Ontology Service API makes use of the Apache Jena Framework [53], in order to provide the required functionality of accessing the Events Ontology. Jena is a free, open source Java Framework for building Semantic Web applications and is composed of a number of APIs interacting together to process RDF Data. It fully takes advantage of the RDF data model, by representing semantic information in the form of a graph. This graph is formed by nodes, representing the subject or the objects of a statement, and by edges that are defined by the predicate of a statement.

VII. COMPUTING LANDSCAPES OF OPERATIONS’ RISK

A. *The Scale of the Computing Problem*

Agent-based modelling and simulation (ABMS) provides means to handle RLs with thousands of distributed nodes as well as means to connect, e.g., a large variety of legacy systems. The choice of algorithms that code a realistic behaviour of agents is a core aspect of modelling. Examples are algorithms to check eligibility of resources to serve in a particular process (workers, tools or components need specified skills) or economic algorithms to minimize idle resources.

Depending on constraints, methods control whether, e.g., agents of components of the product may be active in factory operations or passive in a phase of transfer as one of many shipments in a ship or as one of many stock-keeping units in a storage. An agent may also represent an item that has become critical due to an event or has been taken under control because of an estimated risk. In the example, each node that represents a workstation, i.e., its share in the breakdown of work (Figure 3) or resources assigned to it, pools of workers, inventories of resources etc. can be represented by one agent, or, if a deeper resolution is required, further agents may represent elements of their substructures (sub-nodes).

In an agent-based model of a risk landscape, the number of nodes is equal to the number of active agents and this number may further change due to the dynamics of the system as events may activate nodes that have been passive before and vice versa. Considering similarities to existing HEC applications, the computational scale of an RL may compare to a weather model with a number of geo-cells equal to the maximum number of nodes of an RL.

While the number of nodes in RLs may be smaller (but dynamically change), the number and variety of interactions is noticeably higher. And while weather models have clear inputs like temperature or humidity, managerial models may have to deal with the question “Is it a problem about humidity?” or with the fact that human behaviour (awareness) may have immediate impact on events’ risk and impacts.

However, the computational scale is also driven by response times and needs to specify and compare options to disambiguate interacting events or to identify and implement optimal policies.

Further problems emerge from interactions of operations’ domains (Figure 2) if, e.g., one unplanned event drives a lattice tree of potential propagations in production (accruing backlogs) and in parallel in logistic (withdrawals of inventories to avoid quality hazards), or in engineering. Notably, the propagation of an event in a domain can take a “deviation” across another one.

In any case, the processing of complex RL calls for capacity on a level that matches Big V problems, thus hardware and software systems of the classes of High-End and, potentially, High Performance Computing. In this section we describe “candidate” technologies: Cyber-physical systems and applications of Big Data technologies that capitalise on such capacities. Their fit to the management of risk landscapes is obvious, while MAS (as used in the projects described), allow handling of structures and dynamics in any relevant resolution, may lack the required scalability [54].

B. *Cyber-physical Systems (CPS)*

The acronym CPS stands for a widespread integration program embracing various domains from the automation of buildings, car management and communication-based traffic management to factory and supply-chain automation by integrating technological cyber- and real-world (physical) systems into an adaptive operations’ automation system.

“*Besides further research ..., a fresh look at CPS also requires a new transdisciplinary engineering approach. As we speak about hybrid systems including electronics, mechanics, software and other technical components, new approaches towards integrated systems modelling approach, a coherent design theory and related design, analysis and simulation tools become indispensable. But, cyber-physical systems are not just a self-contained and isolated ensemble of technical components but are often embedded in a social context to form a socio-technical system. In such systems people are embedded in complex organizational structures and interact with complex infrastructures to perform their work processes. A holistic approach towards human factors, including usability of interfaces and functionality, intuitive machine operating, and seamless coordination of human and machine behaviour are of outmost importance to avoid erroneous system behaviour*” [58].

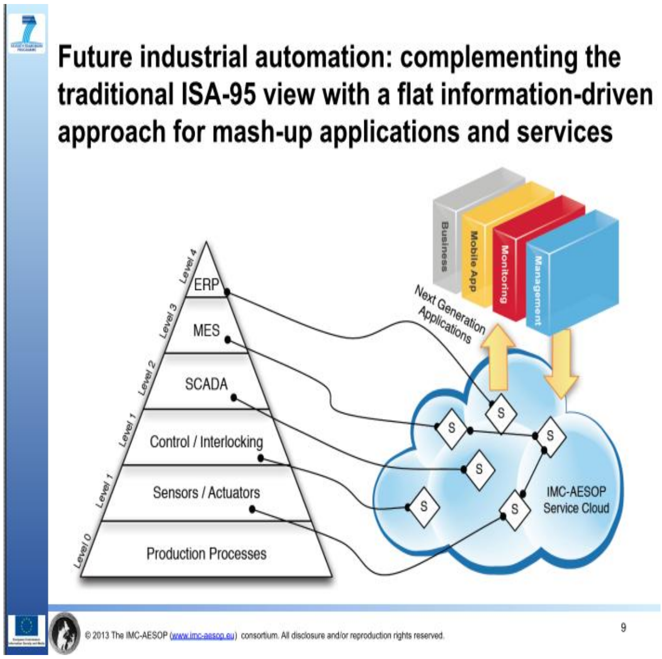

Future industrial automation: complementing the traditional ISA-95 view with a flat information-driven approach for mash-up applications and services

© 2013 The IMC-AESOP (www.imc-aesop.eu) consortium. All disclosure and/or reproduction rights reserved.

9

Figure 12. Enterprise Control Hierarchy (left) into a Service Cloud (right) (Factory-case, accordingly to [55] [56] [57])

Figure 12 is adopted from a presentation about the IMC-AESOP project about industrial automation. [55] It depicts a model of a factory application. On the left, it shows an enterprise control hierarchy from low level atomic activity in the workflows up to the overall planning of processes and resources. The hierarchy implies that any instance on a higher level captures and manages many instances of the level below.

The figure shows that CPS integrate the whole control hierarchy of an enterprise from the lowest level of sensors, actors or controllers embedded into robots, machines or attached to materials (RFID), and middleware like System Control and Data Acquisition or Management Execution Systems (SCADA, MES), on the highest level Enterprise Resource Planning Systems (ERP) as services into a “*cyber-system in the cloud*” [60].

An example proves that this infrastructure is compatible to RL-models and the concepts of risk, impact, or unplanned events. On the lowest level of the control hierarchy, an RFID reader may transmit data that a particular “thing” has been read and “now” is “here”. On the highest level, an ERP system may send data about "things" that "now" are expected to be "here", and, if prepared, can provide information about the impact in the case that this plan fails (unplanned event). The compliance to semantic web standards as well as the option to employ semantic technologies allow CPS to cope with heterogeneity of objects and context-dependency of events. Thus, an RL management system can be one of the "next generation applications" mentioned in Figure 12.

But as CPS are restricted to well defined technological objects and relations, the impact of semantic technologies is limited [88], Data generated, stored or processed are available for further processing by Big Data or All-in-Memory applications (e.g., Apache Hadoop [61]), thus not at least for the management of an RL.

C. Big Data Applications

The jungle of data is growing: "..., *there are approximately 4.4 zettabytes (4.4 trillion gigabytes) of information in the digital universe and this number is expected to reach 44 zettabytes by 2020 ... In 2013, by the IDC's* [International Data Corporation] *count, there were 187 billion "connectable things" on the planet, of which 7% were connected to the digital world. By 2020, the number is expected to rise to 212 billion, with 15% ... generating new data*" [60].

Gartner defines Big Data not as a specific architecture or as a particular service but by their scalability to the Big V, namely volume, velocity and variety [61]. Big Data software is a diverse set of technologies and applications, a fast growing crowd of children of the avalanche of data produced by all ways the internet and related technologies are used.

In the Hype Cycle 2013, Big Data is put on the "Peak of Inflated Expectations" (followed by the "Trough of Disillusionment", next to Consumer 3D Printing and ahead of the Internet of Things). The "Plateau of Productivity" is estimated to be reached in 5 – 10 years (while the IoT may take more time to mature). Developments may vary in different markets or due to the size of companies. But the inflation of the data universe is inescapable:

Any organisation that not effectively adopts respective technologies will get a hard time: "*The machine is the problem: A solution is in the machine*" [63].

The ubiquitous production and availability of data in science, in social networks, in business operations, or in commercial or public surveillance systems, etc., actually from any source that could be imagined, fundamentally the digitalisation of almost everything, generates data clouds of zettascale *volume* and high-*velocity* data flows of an unprecedented *variety* and increasing *variability* but also a questionable *veracity*.

Very likely, Big Data applications will develop capabilities that enable the management of very large risk landscapes: The convergence of Big Data and High Performance Computing is standing to reason:

"*The intersection of these two domains is mainly driven by the use of machine learning methodologies to extract knowledge from big data, and we see an increasing number of platforms that are combining these capabilities to provide hybrid environments that can take advantage of data locality to keep the data exchanges over the network at a manageable level while they offer high performance distributed linear algebra libraries*" [64]. A comparison of HPC and Apache Hadoop Big Data architectures is available in [65].

Business intelligence applications are examples that are relevant for RL management. Supply chains or production lines are sources of massive volumes of data. In the Internet of Things, a typical high-resolution and high speed environment also the velocity of data flows is significant. And to a large degree these data are even well structured. Intelligent algorithms can identify patterns in these flows, support the reading of changing patters (e.g., as indicator of impending criticality) or analyse weaknesses in the system or the effectiveness and efficiency of managerial policies and best-practices.

On life-cycle level, patterns can help to isolate sources of ramp-up problems like weaknesses in the design of the product-production systems, ineffective practices, etc. The analysis of communication flows in social media can provide early indication of a changing customer behaviour, like variations in the trend towards a shareconomy.

Global business platforms are a new development, e.g., collecting massive social data consumer and businesses in the first instance for improving and diversifying their own business. Finally they may become also providers of business intelligence services.

The inflation of data also stands for growing resolution and variety, both implying more sources or targets of risk as well as for acceleration (see Figure 5). The closer to real-time, the less volume and the more velocity and variety will be the core of the task to specify, find, prepare and process the *right* data to support the management of risk landscapes timely (with respect to velocity) and properly (regarding to variety, variability and veracity).

Data-intelligence will have a significant potential to improve business and operations' intelligence in the management of risk landscapes. They identify early problem indicators, analyse the vulnerability of business- and of operations' systems, track interactions between degrees of freedom in dimensions of operations, scanning and analysing drivers of change or positive feedbacks.

The job of data scientists is to develop time-effective semantic models and algorithms for mining meaningful data and asking meaningful questions for further processing and technical support of decision making in the light of valid strategies, finally for the purpose to support R.E.A.L. processes in strategic and operations' management. Beyond, substantial capabilities of "creative criticism" will become paramount: In digitalised, fast changing, heterogeneous, and potentially compromised environments the value is rather in the questions than in the answers.

For illustration we borrow a case from the financial market: Sketchily the financial crisis may have been caused by under-complex models or algorithms respectively by highly leveraged lending that had not been questioned. On a deeper level the reasons are in unquestioned institutional failures that blocked the emergency exits: "*As long as the music is playing, you've got to get up and dance ...*" (Charles Prince, Chairman and CEO of Citigroup Inc. after an about 80 billion USD loss in assets) [67].

Another possible trap may lie in community effects that trigger positive feedbacks and risk. For example, unquestioned predictions in stock markets can become self-fulfilling prophecies, like outcomes of search engines depends on hidden rankings by the search engine operator.

D. Multi-agent Systems (MAS)

MAS are no to CPS or BD-applications but may rather be enablers for intelligent simulation, planning and scheduling as well as to cope with the dynamics and the resolution of detail of RL. An example may be a local unplanned event driving the global criticality of a particular resource that, to mitigate the problem, now need particular care by activating a respective software agent. Given the computational scale of RL the problems may be the scalability of MAS and the need to control the typically high loads of communication MAS impose to the infrastructure.

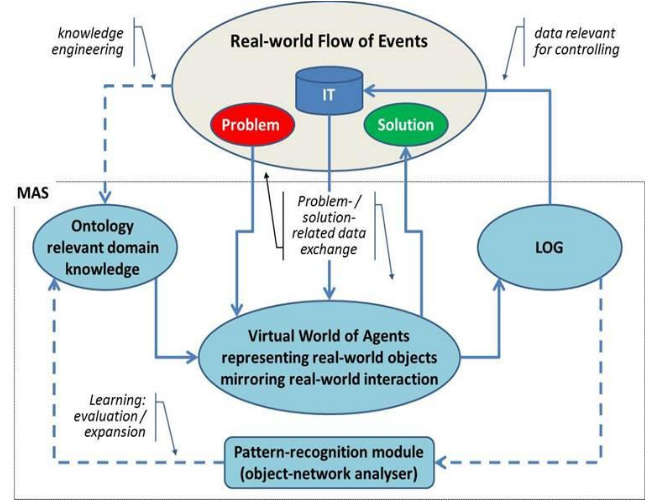

Figure 13. Architecture of an MAS mimicking a real in a virtual world [68].

1) The CargoLifter Knowledge Integrator and CESSAR were realised as multi-agent systems that mimic a relevant part of the world on the base of an ontology. Software agents were applied to objects representing a relevant risk (e.g., a failure would interrupt operations). Relations of agents were designed accordingly to a service-driven logic, analogously but not in full conformity to standards of a service-oriented software architecture [86].

Multi-agent systems (MAS) [86][87] consist of software agents that collaborate in swarms to pursue common, while each pursuing its individual goals. As semantic agents they share a set of concepts (the ontology) that enables to reason about scenes in a compatible way and to coordinate action so that local activity of agents becomes effective in terms of global goal. Semantic agents may be designed for learning from outcomes of decisions, e.g., by comparing current and previous scenes. In the CL or the CESSAR model, agents pursued individual goals on a virtual market by providing services and procuring services needed for their business. Both were stand-alone systems that may exchange data via Internet but not based on Web technology standards. Figure 13 depicts the architecture of these models.

2) There are advantages of this approach: Service-oriented modelling in general is intuitive and the architecture of agent systems allows to design highly complex and multi-layered service supply chains or to integrate any relevant resolution of detail are further advantages of these systems:

- The set of services can be increased by adding agents with respective demands and offering capabilities (truck transport + cleaning + maintenance ...)
- The volume of services is controlled by constraints of capacity only (availability of trucks for service).
- The resolution of object can be increased by adding agents as providers of sub-services via is_part_of relations and service dependencies (in the iCRFID project, e.g., truck → tank → sensor → gas station → pump_#)
- Organising of managerial tasks by multiple holons (e.g., linking all fuel-sensors into a system wide fuel-management *and* into a truck maintenance system).
- True 1:1 interactions between classes (ARUM project: sensor ↔ station) or individual agents (iCRFID project: good_X ↔ passenger_Y) and between events can be implemented as 'objects'.
- The resolution of time is equal to the number of events / unit of time (frequency) and depends on response-times of the MAS to unplanned events.
- Rationales of acting and best practices are implemented in the economics of agents (e.g., transaction costs depending on degree of propagation).
- Systematic exploitation of discretions to act (DTA / slack) by iterative negotiations [84].

3) MAS have not joined the mainstream: *"Despite considerable progress, it seems that the challenges ... encountered at early days still hold. In particular, the adoption of AOSE* [Agent-Oriented Software Engineering] *principles in the academia, and even more so in the industry, is limited*" [69].

Besides the lack of scalability, this is likely also owed to the fact that MAS are able to solve complex problems because they are complex systems: *Solutions are sensitive to start conditions and non reducible to the behaviour of individual agents, thus emergent and hardly reproducible*. From practical experience as managers or engineers, who typically claim to *control* systems, we know that MAS quite plainly prove that this claim mostly is an illusion: "*It works, but I don't know why*." (an experienced dispatcher at Cologne Airport after experiments with *CESSAR* (iCRFID project)).

The scale of RLs require MAS to be redesigned for running on larger and likely also parallel HEC computing infrastructures [70]. Also scalability is still an issue, although holonic MAS architectures enable larger system [71][72]. But also this strategy will hardly solve the problem for very large loads of communication. For instance, a full scale industrial version of *CESSAR* (iCRFID project) where any unplanned event can drive impact across 18,500 flights, operated by 4,460 aircrafts that connect 1,860 destinations for almost 17 Million passengers per day [73].

The development of hybrid MAS is another strategy to reduce the load of communication by employing highly performant mathematical solvers for the calculatory jobs and semantic agents to structure scenes and problems as well as to evaluate results or to support related learning processes.

4) Architectures of CPS and MAS have common features: MAS consist of service-oriented autonomous agents interacting in a cyberspace accordingly to FIPA standards (Foundation for Intelligent Physical Agents) [74][88]. CPS consist of service oriented *applications* (embedded software included) that interact accordingly to web-service standards in the Internet of Services. But while an MAS is a self organising swarm of autonomous, goal-driven agents, a CPS rather compares to an orchestra expected to deliver a particular performance that has been composed by management: the plan of operations, e.g., of a factory or of a complex network of supply chains. Nevertheless the "orchestra" should be able of changing music on demand, e.g., of an unplanned event.

In terms of the R.E.A.L. framework, the equivalent to an alternative piece of music is an alternative plan, either in the form of defined rules or, more complex, as elaborated best practice. MAS provide more flexibility to extend and adapt models or create ad-hoc peer-to-peer relationships between agents respectively actors. And finally MAS are able to simulate and analyse complex systems while a CPS is an architecture that as such needs to employ integrated services, e.g., a MAS for simulation.

5) Big Data may include MAS-based services, e.g., for dynamic planning and scheduling, for simulation and analysis or as semantic reasoners. MAS on the level of high-end or high-performance computing are in reach of development [70], particularly for offline computing task (i.e., not as real-time controllers in a CPS). Promising approaches may be to use BD applications to feed into MAS or, vice versa, to enable Big Data by MAS, e.g., for experiments in risk management for strategic scenarios: "*Who may consider this information to be valuable? ... "What would happen if we provide our product or service free of charge? What if a competitor did so? The responses should provide indications of the opportunities for disruption, as well as of vulnerabilities*" [75]. In operations, 1:1-designed scenarios can be simulated like *"What is the impact of an organisational change to vulnerability?" "What are the limits of current best practices to mitigate impacts of a particular class of unplanned events?"*

E. Limitations to Effectively Parallelizing MAS

MAS are generic distributed systems. This may suggest that agents in a MAS act in parallel. But this is not true. Most MAS are deployed on Microsoft standard software and agents' decision making and communications are sequentially scheduled by allocating capacity slots to tasks or threads. If parallel processors are available and supported by the operations system typically tasks can be distributed. Also holonic architectures can be processed in parallel [76].

The variety of agents' operations systems that allow for real parallel acting of MAS is very limited. Problems lie in the internal communication of agents. Besides the volume of data traffic the messaging protocols of agents are hardly compatible with operations' systems like MPI that are used in parallel computing.

In the context of our work the Repast HPC platform has been analysed [56]. This technology supports parallel agents' activity in an HPC environment, supports large models and enables the communication between agents. However Repast is based on an internal time model the platform is unable to continuously exchange information and synchronize with external systems in real-time / real world. It does not support scheduling or dynamic planning of ongoing operations, that is "online" with actual processes. In consequence Repast is no tool that can be integrated into a service cloud and its use is restricted to simulation.

F. HEC computing architectures and scenarios

With many modern and often dynamical and interactive application scenarios, the term "high performance" is covering demanding applications that are on the one hand compute- and on the other hand data-centric. It is a common understanding that parts of the respective scenarios will support the exploitation of parallelism for their implementation.

With all available high end and high performance systems and architectures the hardware and software issues cannot be separated. The requirements from algorithms and application scenarios lead to solutions favouring the different architectures. In the case of increasingly big data scenarios the attributes of the data and usage are a most important factor.

With the workflows and algorithms the most major attributes of the data, namely volume, velocity, variability, and vitality, mark the physical requirements of needs for communication and data locality. The respective software components have to be adapted in order to fit these requirements, which have to span from distributed to centralised resources, creating robust, reliable, and intelligent software components and workflows.

High End Computing (HEC) systems range from a desktop computer, through clusters of servers and data centres up to high-end custom supercomputers. Resources can be physically close to each other, e.g., in a highly performant compute systems, or the compute power can be distributed on a large number of computers as with most Grid and Cloud computing concepts. Mostly, these architectures are used for task-parallel and data-parallel problems in classical capacity computing.

High Performance Computing (HPC) systems are based on architectures with a large number of processors, for exploiting massive parallelism. Commonly used models are Massively Parallel Processing and Symmetric Multi-Processing, used with the concept of local islands. Due to physically shared memory usage and compute communication, the physical architectures with these HPC systems are different.

Handling of RM processes will therefore focus on distributed components. Due to the physically different structure of highly distributed and massively parallel resources, the following aspects can be considered.

In the case of HEC, e.g., Cloud Computing, these components can be system resources acting autonomously like servers, being connected by external network means, being the ideal resources for events processing at capacity level. HEC resources can provide efficient means for massively distributed tasks. The non-availability of resources can be handled on a job or task base.

In the case of HPC, e.g., common with Scientific Computing on Supercomputing resources, the components can be internal network resources only, compute nodes on the one hand, being controlled by a management network and software, and management nodes on the other hand.

The communication intensive modelling especially for the overall results and visualisation as well as the pre- and post-processing for the models will be suitable for use of HPC resources. In order to optimise the efficiency and economic use of the HPC resources and minimising the effects of job size fragmentation these resources should be used for a defined class of suitable large tasks within the workflow. Available resources can be configured as distributed HPC resources within the network provided for the described systems. Regarding the demanding network requirements Software Defined Networks (SDN) [77] can provide modular and efficient solutions for these purposes.

VIII. CONCLUSIONS AND FURTHER WORK

In most countries, listed firms are required to include a formal analysis of corporate risk in annual reports. Theoretical and descriptive parts are delivered in narrative form. The standard of underlying risk models is based on actuarial methodology that also may deal with relevant operations' risk.

They provide an integrating system of strategies, similar to those applied in insurance business and with similar problems as discussed in this paper: "*A major challenge here is a more substantial and realistic description and modelling of the various complex dependence structures between risks that show up on all scales*" [78]. But integrated risk modelling and processing, as addressed in this paper, is far too detailed and complex to be by this rather formal approach.

Although our work is in an early stadium, the industrial use-cases provide confidence that the particular computational approach discussed above will add a new strategy to risk management under exceptional circumstances in real economy. For operation and management, it is appropriate to focus on risk landscapes as networks of nodes and of related service levels [79][80]. Therefore, events described and related processes can be handled with less interference if services are defined and interfaces for the processes are created.

This is important for the HEC, HPC, and communication resources required. For HEC processes, this can be done on a service level cloud base, whereas for the HPC resources available in research environments, this mostly will have to be assisted by service level agreement policies.

This is important for the HEC, HPC, and communication resources required. For HEC processes, this can be done on a service level cloud base, whereas for the HPC resources available in research environments, this mostly will have to be assisted by service level agreement policies.

In both fields of semantic modelling and computation of industrial landscapes of risk, further work is to be done. The most crucial issues are

- To elaborate a formalised architecture of RL, based on the network of nodes, but consistently including the large variety of structural and dynamic aspects on the required level of detail.
- To develop an effective Bayesian strategy of capturing and improving estimates of event risk and related impact from responsible managers. The issue is that hybrid models require to link semantic conceptualisation with Bayesian methodology [38] that significantly goes beyond the eEV-model used in this paper. Another aspect is that relations between ontological and process-based reasoning (things and flows) may have to be revised [41].
- To deliver a first concrete industrial model of a risk landscape.

ACKNOWLEDGMENTS

This work is partially supported by the EU FP7 Programme, ARUM project, GA- No. 31405.

REFERENCES

[1] U. Inden, D. T. Meridou, M.-E. Ch. Papadopoulou, A.-C. G. Anadiotis, and C.-P. Rückemann, "Complex Landscapes of Risk in Operations Systems Aspects of Processing and Modelling," The Third International Conference on Advanced Communications and Computation (INFOCOMP 2013) IARIA, Nov. 2013, pp. 99-104, ISSN: 2308-3484, ISBN: 978-1-61208-310-0.

[2] A. Dixit and R. Pindyck, "Investment Under Uncertainty", Princeton University Press, 1994.

[3] H. Bruch and S. Goshal, "Management is the Art of Doing and Getting Done," Business Strategy Review, Sept. 2004, vol. 15, pp.4-13.

[4] K. Ishikawa, Introduction to Quality Control, Productivity Press, 1990.

[5] E.-K. Boukas and R. P. Malhame, "Analysis, Control and Optimization of Complex Dynamic Systems," Springer, 2005.

[6] A. M. Brandenburger and B. J. Nalebuff, "Co-Opetition," Crown Business, 1996.

[7] N. N. Talib, "The Black Swan: The Impact of the Highly Improbable (2nd Edition)," Random House, May 2010, ISBN: 978-1400063512.

[8] Why Project Fails. *Airbus – A380*. [Online]. Available from: http://calleam.com/WTPF/?p=4700/ 2014.11.14

[9] Why Project Fails. *Boeing Commercial Aeroplanes*. [Online]. Available from: http://calleam.com/WTPF/?p=4617 2014.11.14

[10] F. Krings and O. Krone, "Konzept zur interactiven Transportsteuerung im Rahmen globaler Verkehrssysteme," Technische Universität Carolo Wilhelmina Braunschweig, 1996.

[11] R. Kohlenberg and R. Schulze, "Interoperabilität von verschiedenen Barcodesystemen zur Sendungsverfolgung in der Logistikkette VW – VWdM," Technische Universität Carolo Wilhelmina Braunschweig, 1997.

[12] GLORI. *Global Logistics Research Initiative*. [Online]. Available from: www.glori.com/research.htm/ 2014.11.14

[13] Bundesministerium für Wirtschaft und Energie. *iC-RFID: RFID-gestützte Servicesysteme*. [Online]. Available from: http://bmwi.de/DE/Themen/Digitale-Welt/Internet-der-Zukunft/internet-der-dinge,did=360478.html/ 2014.11.14

[14] aerospace-technology.com. *CargoLifter CL160*. [Online]. Available from: www.aerospace-technology.com/projects/cargolifter/ 2014.11.14

[15] R. Franken and U. Inden, "CESSAR – Configuration and Evaluation of Service Systems in Air-Catering with RFID," Cologne Business Working Papers, Dec. 2011. ISSN: 2192-936.

[16] ARUM. *Adaptive Production Management*. [Online]. Available: http://www.arum-project.eu/ 2014.11.14

[17] S. L. Goldman, R. N. Nagel, and K. Preiss, "Agile Competitors and Virtual Organisations: Strategies for Enriching the Customer," Wiley, 1995.

[18] K. Preiss, S. L. Goldman, and R. N. Nagel, "Cooperate to Compete: Building Agile Business Relationships," Wiley, 1996

[19] Industrial Excellence Award. *Benchmarking Management Quality for European Competitiveness*. [Online]. Available from: http://de.industrial-excellence-award.eu/home/ 2014.11.14

[20] P. Hermanns, "Organizational Hubris – Aufstieg und Fall einer Celebrity Firm am Beispiel der CargoLifter AG", Kölner Wissenschaftsverlag, 2012. ISBN 978-3-942720-33-5.

[21] R. Bea. *Approaches to achieve adequate quality and reliability*. [Online]. Available from: http://ccrm.berkeley.edu/pdfs_papers/bea_pdfs/quality-and-reliability.pdf 2014.11.14

[22] S. A. Srinivasa Moorthy, "Lifecycle Challenges in Long Life and Regulated Industry Products," ICoRD'13, Lecture Notes in Mechanical Engineering, Springer India, pp. 833-844, 2013, doi:10.1007/978-81-322-1050-4_66.

[23] A. McAfee and E. Brynjolfsson, "Big Data: The Management Revolution," Harvard Business Review, pp. 61-68, Oct. 2012.

[24] Investopedia. *Working Capital*. [Online]. Available from: http://www.investopedia.com/terms/w/workingcapital.asp 2014.11.14

[25] W. Ashby, "The Law of Requisite Variety," An Introduction to Cybernetics, Chapman & Hall, 1956, pp. 206–212, ISBN:0-416-68300-2.

[26] W. Ashby. *The W. Ross Ashby Digital Archive*. [Online]. Available from: http://www.rossashby.info/journal/page/4158.html 2014.11.14

[27] U. Inden, G. Lioudakis, and C.-P. Rückemann, "Awareness-Based Security Management for Complex and Internet-Based Operations Management Systems," Integrated Information and Computing Systems for Natural, Spatial, and Social Sciences, IGI Global, 2013, pp. 43-73, ISBN 978-1-4666-2190-9.

[28] SKYbrary. *Airborne Separation Assurance Systems (ASAS)*. [Online]. Available from: http://www.skybrary.aero/index.php/Airborne_Separation_Assurance_Systems_%28ASAS%29 2014.11.14

[29] R. Kurzweil. *The Law of Accelerating Returns*. [Online]. Available from: http://www.kurzweilai.net/the-law-of-accelerating-returns 2014.11.14

[30] J. A. Schumpeter, "Capitalism, Socialism and Democracy". Taylor & Francis e-Library, 2003, ISBN 0-415-10762-8.

[31] Reuters Deutschland. *Osram-Mitarbeitern stehen Kündigungen ins Haus*. [Online]. Available from: http://de.reuters.com/article/companiesNews/idDEKBN0FZ1H120140730 2014.11.14

[32] C. Christensen, "Das Dilemma der Kapitalisten," Harvard Business Manager, 2014.

[33] T. Berners Lee, J. Hendler, and O. Lassila. *The Semantic Web*. [Online]. Available from: http://www.scientificamerican.com/article/the-semantic-web/ 2014.11.14

[34] The University of Sydney. *A step closer to bio-printing transplantable tissues and organs*. [Online]. Available from: http://sydney.edu.au/news/84.html?newsstoryid=13715 2014.11.14

[35] U.S. Army. *Chow from a 3-D printer? Natick researchers are working on it*. [Online]. Available from: http://www.army.mil/article/130154/Chow_from_a_3_D_printer__Natick_researchers_are_working_on_it/ 2014.11.14

[36] 3D Print Canal House. *What is a 3D Print Canal House?*. [Online]. Available from: http://3dprintcanalhouse.com/what-is-the-3d-print-canal-house-2 2014.11.14

[37] Engineering & Technology magazine (E&T). *Metal 3D printing promises revolution in aerospace*. [Online]. Available from: http://eandt.theiet.org/news/2013/oct/metal-3d-printing.cfm 2014.11.14

[38] Travel Daily Asia. *Airbus plans to make aircraft using 3D printers*. [Online]. Available from: http://www.traveldailymedia.com/205079/airbus-plans-to-make-aircraft-using-3d-printers/ 2014.11.14

[39] 3D Printer. *GE Announces Production 3D Printing and Stock Goes Up*. [Online]. Available from: www.3dprinter.net/ge-announces-production-3d-printing-stock-goes 2014.11.14

[40] 3D Printer. *British RAF Fighter Jets Fly with 3D Printed Parts for the First Time*. [Online]. Available from: http://www.3dprinter.net/british-raf-fighter-jets-fly-3d-printed-parts-first-time 2014.11.16

[41] A. Reichental. *How 3D printing will turn us all (back) into makers: Avi Reichental at TED2014*. [Online]. Available from: http://blog.ted.com/2014/03/19/how-3d-printing-will-turn-us-all-back-into-makers-avi-reichental-at-ted2014/ 2014.11.16

[42] D. Etherington. *Amazon Launches A 3D Printing Store With Customizable Goods*. [Online]. Available from: http://techcrunch.com/2014/07/28/amazon-launches-a-3d-printing-store-with-customizable-goods/ 2014.11.16

[43] EXPLORA Frankfurt Science Centre. *Museum + Wissen + Tech + Arts*. [Online]. Available from: http://www.explora.info/pressepix/pressepix9.php 2014.11.16

[44] K. E. Ch. Asch, "A Very Practical Geoinformatics Project: The Reality of Delivering a Harmonized Pan-European Spatial Geoscience Database," Geoinformatics 2007 - Data to Knowledge, GSA, May 2007, pp. 4-5.

[45] W. Kuhn, "Semantic engineering," Research Trends in Geographic Information Science, Springer, pp. 63–76, Jun. 2009, doi. 10.1007/978-3-540-88244-2_5, ISSN: 1863-2246, ISBN: 978-3-540-88244-2.

[46] M. Frixione and A. Lieto, "Towards an Extended Model of Conceptual Representations in Formal Ontologies: A Typicality-Based Proposal," Journal of Universal Computer Science, vol. 20(3), pp. 257-276, Mar. 2014.

[47] M. Kuhlmann. *Was ist real?*. [Online]. Available from: http://www.spektrum.de/alias/titelthema-quantenfeldtheorie/was-ist-real/1286309 2014.11.16

[48] ISO Store. *ISO 31000:2009, Risk management - Principles and guidelines*. [Online]. Available from: http://www.iso.org/iso/catalogue_detail?csnumber=43170 2014.11.16

[49] N. A. Doherty and A. Muermann, "On the Role of Insurance Brokers in Resolving the Known, the Unknown and the Unknowable," The Known, the Unknown, and the Unknowable in Financial Risk Management: Measurement and Theory

Advancing Practice, Princeton University Press, pp. 194-209, 2010.

[50] J. L. Synge. *Events and Spacetime* [Online]. Available from: http://www.phy.syr.edu/courses/modules/LIGHTCONE/events.html 2014.11.16

[51] J. A. Rooke, L. J. Koskela, and D. Seymour, "Producing things or production flows? Ontological assumptions on the thinking of managers and professionals in construction," Journal of Construction Management and Economics, Taylor & Francis, vol. 25 (10), pp. 1077-1085, Oct. 2007, ISSN: 0144-6193.

[52] A. Gelman, "Prior distribution," Encyclopedia of Environmetrics, Wiley, vol. 3, pp. 1634-1637, 2002, ISBN: 0471 899976.

[53] The Apache Software Foundation. *Apache Jena*. [Online]. Available from: http://jena.apache.org/index.html/ 2014.11.17

[54] P. Leitão, U. Inden, and C.-P. Rückemann, "Parallelizing Multi-gent Systems for High Performance Computing," The Third International Conference on Advanced Communications and Computation (INFOCOMP 2013) IARIA, Nov. 2013, pp. 1-6, ISSN: 2308-3484, ISBN: 978-1-61208-310-0.

[55] S. Karnouskos, "Cloud-based Cyber-Physical Systems in Industrial Automation," The 39th Conference of the IEEE Industrial Electronics Society (IECON 2013), IEEE, Nov. 2013.

[56] The IMC-AESOP Consortium. *IMC-AESOP Project: Architecture for Service Oriented Process.* [Online]. Available from: http://www.imc-aesop.eu/dl/13-4285%20AESOP%20broschyr_skiss4.pdf 2014.11.16

[57] A. W. Colombo, S. Karnouskos, and T. Bangemann, "A System of systems view on collaborative industrial automation," IEEE International Conference on Industrial Technology (ICIT 2013), IEEE, Feb. 2013, pp. 1968 – 1975, ISBN: 978-1-4673-4568-2, doi: 10.1109/ICIT.2013.6505980.

[58] B. J. Krämer, "Evolution of Cyber-Physical Systems: A Brief Review," Applied Cyber-Physical Systems, Springer New York, 2014, pp. 1-3, ISBN: 978-1-4614-7335-0, doi: 10.1007/978-1-4614-7336-7_1.

[59] Apache Software Foundation, "Apache Hadoop," [Online]. Available from: http://hadoop.apache.org / [Last accessed 16/11/2014].

[60] A. Gibson. *Growth of Big Data a big challenge for business*. [Online]. Available from: http://www.strategic-risk-global.com/growth-of-big-data-a-big-challenge-for-business/1408171.article 2014.11.16

[61] Gartner IT Glossary, "Big Data," [Online]. Available from: http://www.gartner.com/it-glossary/big-data/ [Last accessed 16/11/2014].

[62] N. Heudecker. *Hype Cycle for Big Data, 2013*. [Online]. Available from: https://www.gartner.com/doc/2574616 2014.11.16

[63] Y. Poullet, "EU data protection policy. The Directive 95/46/EC: Ten years after," Computer Law and Security Review, Elsevier, 2006, vol. 22, Issue 3, pp. 106-127, doi: 10.1016/j.clsr.2006.03.004.

[64] F. Villanustre. *Big Data Technologies Narrow the Gap between HPC and the Enterprise*. [Online]. Available from: http://www.isc-events.com/bigdata13/pressreleases-reader/items/big-data-technologies-narrow-the-gap-between-hpc-and-the-enterprise.html 2014.11.16

[65] S. Jha, J. Qiu, A. Luckow, P. Mantha, and G. C. Fox. *A Tale of Two Data-Intensive Paradigms: Applications, Abstractions, and Architectures*. [Online]. Available from: http://arxiv.org/abs/1403.1528 2014.11.17

[66] H. Adams and P. Varadan. *Fighting Financial Crime With Data*. [Online]. Available from: http://www.accenture.com/SiteCollectionDocuments/PDF/Accenture-Fighting-Financial-Crime-with-Data.pdf 2014.11.17

[67] Reuters. *Ex-Citi CEO defends dancing quote to U.S. panel*. [Online]. Available from: www.reuters.com/assets/print?aid=USN0819810820100408 2014.11.17

[68] G. Rzevski, "A practical Methodology for Managing Complexity," Emergence: Complexity and Organization–an International Transdisciplinary Journal of Complex Social Systems, vol.13, pp. 38-56, 2011.

[69] A. Sturm and O. Shehory, "Agent-Oriented Software Engineering: Revisiting the State of the Art," Agent-Oriented Software Engineering, Springer Berling Heidelberg, 2014, pp 13-26, ISBN: 978-3-642-54432-3, doi: 10.1007/978-3-642-54432-3_2

[70] P. Leitão, U. Inden, and C.-P. Rückemann, "Case Studies for Parallelising Multi-Agent Systems for High Performance Computing," International Journal On Advances in Software, IARIA, 2015 (to appear)

[71] H. Van Brussel, J. Wyns, P. Valckenaers, L. Bongaerts, and P. Peeters, "Reference Architecture for Holonic Manufacturing Systems: PROSA," Computers in Industry, Elsevier, Nov. 1998, vol. 37, Issue 3, pp. 225-276, doi:10.1016/S0166-3615(98)00102-X.

[72] A. Koestler, "The Ghost in the Machine," Hutchinson & Co, 1967.

[73] Star Alliance. *Travel the World with the Star Alliance Network.* [Online]. Available from: http://www.staralliance.com/de/about/member_airlines/ 2014.11.17

[74] IEEE computer society. *Foundation for Intelligent Physical Agents (FIPA)*. [Online]. Available from: http://www.computer.org/portal/web/sab/foundation-intelligent-physical-agents/ 2014.11.17

[75] J. Bughin, M. Chui, and J. Manyika. *Clouds, big data, and smart assets: Ten tech-enabled business trends to watch*. [Online]. Available from: http://www.mckinsey.com/insights/high_tech_telecoms_internet/clouds_big_data_and_smart_assets_ten_tech-enabled_business_trends_to_watch/ 2014.11.17

[76] P. Leitão and J. Barbosa, "Adaptive Scheduling based on Self-organized Holonic Swarm of Schedulers," IEEE 23rd International Symposium on Industrial Electronics (ISIE), IEEE, Jun. 2014, pp. 1706-1711, doi: 10.1109/ISIE.2014.6864872.

[77] A. Georgi, R. Budich, Y. Meeres, R. Sperber, and H. Hérenger, "An Integrated SDN Architecture for Applications Relying on Huge, Geographically Dispersed Datasets," The Third International Conference on Advanced Communications and Computation (INFOCOMP 2013) IARIA, Nov. 2013, pp. 129-134, ISSN: 2308-3484, ISBN: 978-1-61208-310-0.

[78] D. Pfeifer and J. Nešlehová, "Modelling and generating dependent risk processes for IRM and DFA," ASTIN Bulletin, Peeters, 2004, vol. 34, pp. 333-360, doi: 10.2143/AST.34.2.505147.

[79] C.-P. Rückemann, "Enabling Dynamical Use of Integrated Systems and Scientific Supercomputing. Resources for Archaeological Information Systems," The Second International Conference on Advanced Communications and Computation (INFOCOMP 2012) IARIA, Oct. 2012, pp. 36-41, ISBN: 978-953-307-737-6.

[80] C.-P. Rückemann, "Queueing Aspects of Integrated Information and Computing Systems in Geosciences and Natural Sciences," Advances in Data, Methods, Models and Their Applications in Geoscience, InTech, Dec. 2011, pp. 1-26, ISBN: 978-953-307-737-6, doi: 10.5772/29337.

[81] J. Ríos, F. Mas, and J. L. Menéndez, "Aircraft Final Assembly Line Balancing and Workload Smoothing: A Methodical Analysis," Kay Engineering Materials, Trans Tech Publications, Feb. 2012, vol. 502, pp. 19–24, ISSN: 1662-9795, doi: 10.4028/www.scientific.net/KEM.502.

[82] K. Sundaram and F. Tomesco. *Boeing Disputes India Report of $500 Million 787 Payment*. [Online]. Available from: http://www.businessweek.com/news/2012-03-14/boeing-to-pay-air-india-500-million-on-787-delays-india-says/ 2014.11.17

[83] E. Fleisch and G. Müller-Stewens, "High-Resolution-Management: Konsequenzen des Internet der Dinge auf die Unternehmensführung," Zeitschrift Führung + Organisation (ZfO), Schäffer-Poeschel, 2008, vol. 77, pp. 272 - 281, ISSN 0722-7485.

[84] U. Inden, S. Naimark, and C.-P. Rückemann, "Towards a Discretion-to-Act Control Architecture by Decoupling Modelling from Complexity," TMC Academic Journal, TMC Academy, Feb. 2013, vol. 7, ISSN: 1793-6020.

[85] L. Feigenbaum. *Semantic Web vs. Semantic Technologies.* [Online]. Available from: http://www.cambridgesemantics.com/de/semantic-university/semantic-web-vs-semantic-technologies/ 2014.11.17

[86] P. Skobelev, "Bio-Inspired Multi-Agent Technology for Industrial Applications," Multi-Agent Systems - Modeling, Control, Programming, Simulations and Applications, InTech, Apr. 2011, ISBN: 978-953-307-174-9, doi: 10.5772/14795.

[87] M. Wooldridge, "Introduction to Multi Agent Systems," John Wiley and Sons, ISBN: 978-0470519462.

[88] P. Leitão, J. Mendes, A. Bepperling, D. Cachapa, A.W. Colombo, and F. Restivo, "Integration of virtual and real environments for engineering service-oriented manufacturing systems," Journal of Intelligent Manufacturing, Springer US, Dec. 2012, vol. 23, pp. 2551-2563, doi: 10.1007/s10845-011-0591-8.

[89] U. Inden, N. Mehandjiev, L. Mönch, and P. Vrba, "Towards an Ontology for Small Series Production," Industrial Applications of Holonic and Multi-Agent Systems, Lecture Notes in Computer Science, pp. 128-139 , 2013.

[90] A. Huchzermeier and C. H. Loch, "Project Management Under Risk: Using the Real Options Approach to Evaluate Flexibility in R&D," Management Science, informs, Jan. 2001, vol. 47, Issue 1, pp 85-101, ISSN: 0025-1909, doi: 10.1287/mnsc.47.1.85.10661.

[91] S. Spinler and A. Huchzermeier, "Realoptionen: Eine marktbasierte Bewertungsmethodik für dynamische Investitionsentscheidungen unter Unsicherheit," Controlling und Management, Gabler Verlag, Mar. 2004, vol. 48, Issue 1, pp 66-71, ISSN: 1864-5410, ISBN: 978-3-663-01579-6, doi: 10.1007/BF03255757.

[92] MyBusinessCommunities. *ADVENTURES: Entwicklung von Bewertungsmethoden und Internet-basierten Handelssystemen für Optionen auf Start-Up Ventures, Dienstleistungen und Nicht-lagerfähige Produkte (Finanzdienste)*. [Online]. Available from: http://www.dl2100.de/projectdetail.php?PHPSESSID=e16628 ...ort1=4&projectid=9/ 2014.11.17

[93] R. Gulati, C. Casto, and C. Krontiris. *How the Other Fukushima Plant Survived*. [Online]. Available from: https://hbr.org/2014/07/how-the-other-fukushima-plant-survived/ 2014.11.17

[94] D. Weinberger, "Too Big To Know: Rethinking Knowledge Now That the Facts Aren't the Facts, Experts Are Everywhere and the Smartest Person in the Room," Basic Books, 2012, ISBN: 978-0-465-02142-0.

[95] E. Brynjolfsson and A. McAfee, "The Second Machine Age," W.W.Norton & Compnay, Jan. 2014, ISBN: 978-0-393-23935-5.

[96] J. Manyika et al.. *Big data: The next frontier for innovation, competition, and productivity*. [Online]. Available from: http://www.mckinsey.com/~/media/McKinsey/dotcom/Insights%20and%20pubs/MGI/Research/Technology%20and%20Innovation/Big%20Data/MGI_big_data_full_report.ashx/ 2014.11.17

[97] C. Anderson. *The End of Theory: The Data Deluge Makes the Scientific Method Obsolete*. [Online]. Available from: http://archive.wired.com/science/discoveries/magazine/16-07/pb_theory/ 2014.11.17

[98] W3C. *SPARQL Query Language for RDF*. [Online]. Available from: http://www.w3.org/TR/rdf-sparql-query/ 2014.11.17