8th International Conference on Digital Enterprise Technology - DET 2014 – "Disruptive Innovation in Manufacturing Engineering towards the 4th Industrial Revolution

A solution for information management in logistics operations of modern manufacturing chains

Elisabeth Ilie-Zudor[a]*, Zsolt Kemény[a], Anikó Ekárt[b],
Christopher D. Buckingham[b], László Monostori[a]

[a]*Institute for Computer Science and Control, Hungarian Academy of Sciences, 1111 Budapest, Hungary*
[b]*Aston University, Aston Triangle, B4 7ET Birmingham, United Kingdom*

* Corresponding author. Tel.: +36 1 279 6195; fax: +36 1 466 7503. *E-mail address:* ilie@sztaki.mta.hu

Abstract

One dominant feature of the modern manufacturing chains is the movement of goods. Manufacturing companies would remain an unprofitable investment if the supplies/logistics of raw materials, semi-finished products or final goods are not handled in an effective way.

Both levels of a modern manufacturing chain –actual production and logistics- are characterized by continuous data creation at a much faster rate than they can be meaningfully analyzed and acted upon manually. Often, instant and reliable decisions need to be taken based on huge, previously inconceivable amounts of heterogeneous, contradictory or incomplete data.

The paper will highlight aspects of information flows related to business process data visibility and observability in modern manufacturing networks. An information management platform developed in the framework of the EU FP7 project ADVANCE will be presented.

© 2014 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY-NC-ND license
(http://creativecommons.org/licenses/by-nc-nd/3.0/).
Peer-review under responsibility of The International Scientific Committee of the 8th International Conference on Digital Enterprise Technology - DET 2014 – "Disruptive Innovation in Manufacturing Engineering towards the 4th Industrial Revolution"

Keywords: information management; big data; decision support

1. Introduction

In the present business landscape, manufacturing companies are not simple independent entities anymore, but parts of value-adding networks. When companies are adding value to physical goods, they become parts of multi-echelon networks (i.e. supply-chains, delivering goods and related services to the final customer [1, 2]. Operations management theory claims that controlling these multi-echelon networks of companies integrally can provide significant benefits [3, 4, 5]. Material flow transparency, specifically the visibility to inventories and deliveries in the whole supply network, is considered an imperative requirement for successful supply-chain management, and has been associated with significant efficiency and quality improvements [6, 7].

One dominant feature of the modern manufacturing chains is the movement of goods. Manufacturing companies would remain an unprofitable investment if the logistics of raw materials, semi-finished products or final goods are not handled in an effective way.

Both levels of a modern manufacturing chain –actual production and logistics- are characterized by continuous data creation at a much faster rate than they can be meaningfully analyzed and acted upon manually. Often, instant and reliable decisions need to be taken based on huge, previously inconceivable amounts of heterogeneous, contradictory or incomplete data. As highlighted by S. Bertolo [8], there is an emerging need to develop intelligent knowledge management systems that allow "to progressively integrate an organization's implicit knowledge into its formal business processes and to be able to expose both to third parties in the dynamic creation of virtual organizations as required by common business objectives."

The paper will highlight aspects of information flows related to business process data visibility and observability in modern manufacturing networks, and will present an information management platform developed in the
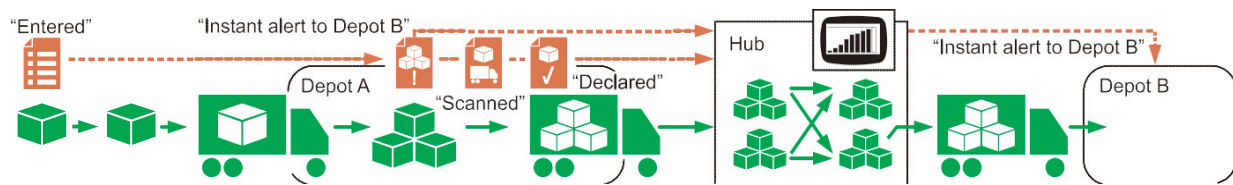
Fig. 1: Typical material flow and information flow in a hub and spoke logistics network [20]

framework of the EU FP7 project ADVANCE ("Advanced predictive-analysis-based decision support engine for logistics", http://advance-logistics.eu/).

2. Business process data visibility and observability

Manufacturing chains/networks suffer from sub-optimal dispatching, replenishment or production decisions due to three factors:

- the lack of information: in many cases information available at a given time and location, is not shared across the network or collected over time (e.g., inbound and outbound shipment units in a local warehouse, some of which—if shared across the network—would be already decisive for estimation or early detection of transportation demands at other locations);
- the massive amount of information available for gaining decision-critical insight is beyond the capabilities of human personnel to assimilate and understand;
- the large amount of information arising from constantly changing demands and associated profiling information (customer preferences, individual or batch-level product data, e.g., for perishable goods, service profiles, preferences and constraints) remains relatively "flat" from the point of view of a "bigger picture" (e.g., general trends) because its structure is obscured.

Transparency of processes—i.e., the ability to gain exact information about processes without notable time lag—has been highlighted in the literature as one of the key prerequisites for improved control or coordination of various processes in production and delivery [9, 10, 11]. In most cases, the following issues receive attention: accuracy of information, timing of information, and granularity of information. The former two properties primarily depend on the technologies employed–revealing that accurate and free-flowing information needs a maximal removal of human intervention at the critical points of data acquisition. While this is self-evident in other problem domains (e.g., control engineering), only the past few decades have seen a comparable spreading of automatic data acquisition on our scale of production and delivery, namely, the introduction of automatic identification (AutoID) techniques.

The granularity of information covers two aspects: the question of distinguishing individual instances vs. observing mere quantities, and the depth of observation (items, pallets, batches etc.). Currently, it is still widespread industrial practice to merely observe stock levels at a given location (referred to as account-based approach) [12], as in a number of applications this proves to be sufficient. The prevalence of

this approach is also shown, for example, by the still widespread use of the so-called EAN13 (European Article Number, ISO/IEC 15420) code for merchandise which does not distinguish individual instances by, e.g., a serial number. However, as product customisation is spreading, or more intricate customer or government requirements for product traceability are faced, more and more applications require the unique distinction of instances of the same class of articles (i.e., a so-called item-centric approach).

An extensive study dealt with the introduction of AutoID technologies and functionalities based thereon, and can be cited here as a comprehensive work explaining principles and current findings of the field [13]. The study identified four functionality levels based on identifiers:

- An identification infrastructure alone means the mere presence of an identification technology, and the possibility of meeting decisions on the spot, based on local knowledge about the identified object.
- Identifier-based operations already presume the existence of a central information repository which assists in meeting local decisions (e.g., if an entity is authorized to pass through a facility gate).
- Tracking-based operations give individual reading acts a meaning, since detecting an ID at a given time and place (plus, optionally, other conditions) implies that the item in question was physically present at the point of reading. This event is then stored in a database, so that it becomes possible to keep track of what occurred to it during its life cycle. The same applies to interactions with other instances.
- Advanced item-centric services become necessary if relevant parts of the life cycle take place under the authority of other parties, or the complexity of the processes require to maintain transparency across organizational borders. In such cases, item-related data and services (e.g., notification, subscription) are shared among process participants with proper access restrictions.

The highest functionality level massively exploited in the industry is the layer of tracking-based operations [13], however, some early proprietary solutions as well as commercially available business integration frameworks already reach into the layer of advanced item-centric services.

The spreading of AutoID based solutions, also created issues in need of further resolution, as this technology enables–and competition even forces–logistics and trading companies to take into account an order of magnitude of data they never experienced before. To give an approximate picture: RFID explodes the information related to traditional order processing by a factor of 10,000 or 100,000, and sales

slip line processing explodes the usual order processing data by similar factors.

Even if some re-thinking in information granularity (border of item-centric vs. account-based solutions, designating item groups as unique instances etc.) may tame the data amounts in some selected cases, most industrial users will inevitably face the burden of an information flood. Concerns of being unable to handle such amounts of data are realistic–mostly, the inability to cope with the volume of such data even hampers their exploitation for operational decisions.

It can be concluded that the technological background of low-level process transparency is far better developed and more widespread in industrial applications than the efficient harvesting of the resulting masses of information for gaining insight on a higher level of abstraction. This is especially so for logistics operations, given that logistics networks create billions of items of information every month, which all need analysing if instant and strategic decision making are to be properly integrated.

The last few decades have seen the emergence of logistics networks that successfully bridge the apparent gap of fast and low-cost shipping of small - typically less-than-truckload - consignments. One proven way of such solutions is a multi-level structure that implements, from an operational point of view, a class of hub and spoke networks (Figure 1), while from an organisational perspective it is an enterprise network of a "major player" with a large-throughput core structure of hubs, contracting smaller local logistics providers also called operating depots, with collection and delivery services [14]. The key to the success of such networks is the bundling and re-bundling of shipments [15, 16] that make transporting less-than-truckload consignments economically feasible. Enterprise interoperability has been recognized as a paradigm vital for improving processes of operations spanning enterprise borders [17, 18, 19] and recent years have witnessed much progress also related to hub operations.

The complexity of dealing with data present in hub and spoke networks is increased by the fact that the data needs to be captured and utilized from a multitude of different data sources. Examples of the kind of data are:

1.   Historical data (pallet flows, orders, scans in and out, normal volume per day-of-week, per week-in-month, per month etc...)

2.   Semi-static hub data - e.g. vehicle time-slots, bay allocations, hub-depot travel time matrix, pricing.

3.   Dynamic hub data - e.g. pallets currently held at the hub.

4.   Dynamic depot data - e.g. new pallet orders, new pallet labelling events, vehicle manifests etc.

5.   Adjacency data - extrapolating trends from one region and applying them to a nearby region, or using trends between distant regions which are linked by pallet flow patterns.

6.   Travel time data - e.g. digital mapping with road speeds, live traffic feeds, GPS tracking etc.

7.   Economic data - e.g. state of the economy, consumer spending, consumer buying habits and whether the current month contains four or five weeks.

8.   Other external data - e.g. weather, public holidays, knowledge of major events (for example, the 2012 Olympics in London), potentially even using commercial geographic census-type business data.

A very conservative estimate is that in such networks 40,000 data items are created every minute, which can be handled neither by human operators nor by conventional data-analysis approaches. Instead, new and sophisticated data mining techniques are required to pluck out the important patterns in the shimmering data display and present these patterns in a form that can be understood and acted upon by human decision makers in real time. The trillions of potential relationships and dependencies in the data items are subject to combinatorial explosion and are absolutely not amenable to systematic number crunching: new, rapid, and focussed intelligence is required, guided top-down by the human decision-making experts and bottom up by the data mining approaches. The top-down element is partly informed by longer-term analyses of the billions of data accumulating over days and weeks but the bottom-up analysis depends on working as fast as possible on thousands of new pieces of information arriving across the network every minute.

The deficiencies resulting from inefficient utilisation of resources due to limitations in processing   localised information in larger amounts and over larger ranges can be mitigated if companies

- invest in sharing and collating information over temporal and hierarchical ranges, and
- introduce methods of analyzing the collected data comprehensively enough to cover the sources of network operation deficiencies.

In view of the amount and span of information (both in terms of location as well as dynamically over time) modern artificial intelligence methods can provide a substantial advantage by their abilities to collate and filter the available information, identify phenomena of importance, and provide decision support, forecasts or early warnings for human personnel. This, again, requires:

- data to be obtained and processed in machine form so that data mining can be applied;
- methods for automatically   identifying   relevant information;
- and the intelligence for relating this information to selected decision priorities for the network operation.

In many cases, the acceptance of decision support systems (and therefore, the "return on investment" in data sharing and processing) is leveraged beyond a critical level if the output of automatic processing is in a human-interpretable form (i.e., it "makes sense" with the experienced operator), allowing human assessment of the decision quality.

3. Information management through a predictive-analysis-based decision support platform

In order to support networked companies addressing these challenges and increase responsiveness of logistics business processes to internal and external dynamics, a predictive-analysis-based decision support platform was developed as

part of the FP7 project ADVANCE (http://advance-logistics.eu/).

The ADVANCE platform provides a dual perspective on transport requirements and decision making dependent on the latest snapshot information and the best higher-level intelligence, featuring the following key functionalities:

- allows companies to extend their already existing infrastructure towards better information sharing;
- provides means for exploiting this information for better operational decisions;
- presents automatically generated results in a human-interpretable way;
- facilitates the alignment of artificial and human expertise so that they can cross-validate and collaboratively adapt the system as the knowledge domains evolve.

Research in ADVANCE focused on exploiting locally existing information (potentially including data that can be made available with a minimal addition of tracking infrastructure) by performing adequate pre-processing steps, making it available to a widened circle of the targeted network (or other corporate structure), identifying models for decision support, and making human-interpretable predictions.

Quantitative improvement in local decisions. Locally available information was gathered, pre-processed and exploited with a higher efficiency than before, allowing for the modelling of processes, the extraction of interpretable process features (in other words, formalized expressions of process status), and the prediction of adequate measures for improving or maintaining given process states. The ADVANCE framework pursues real-time responses to the decision problems.

Quantitative improvement of collaboration processes network-wide. Here, the task consisted in gathering and using Quantitative information regarding selected aspects of collaboration (e.g., transportation capacities and decision preferences of logistics service providers or suppliers within a supply chain) were gathered and used for, optimisation of the network as a whole. Once the information of interest is available for sharing within the network (especially among partners appearing "nearby" in a functional chain), decisions can be met to the benefit of improved local management of actions without the burden of lagging and coarse central coordination. This is of special importance if certain actions are typically carried out in chains spanning several local network members (e.g., handing over a shipment to the next forwarder).

3.1. Perspectives of decision support in ADVANCE

The ADVANCE solution offers decision support to both type of actors in a hub and spoke network, i.e. both hub and depots.

The decision support from the Hub perspective

Improved responsiveness of business processes is expected to lead to (1) a reduction in pallet handling time at the hub, which may lead to an improvement on delivery times further downstream in the transport network and (2) better use of warehouse capacity.

Hub objective 1: Reduction of pallet handling time at the hub. It is expected that the sooner information on incoming and outgoing pallets is available, the more efficiently the hub can be operated. Efficient operations lead to less handling time. The sequence of tipping pallets from inbound trucks should be assessed carefully with respect to the sequence of loading pallets on outbound trucks. The tip and load sequence should be optimised jointly. The tip and load sequence mainly depends on the number and characteristics of the consignments being announced and manifested by member depots. Any delay in processing this data may lead to a suboptimal sequencing process. The sequence is also dependent on factors such as last minute changes in depot decisions, delays of incoming trucks as a result of traffic conditions. Direct availability of all data regarding pallet arrivals is essential for efficient hub operation.

Hub objective 2: Better use of warehouse capacity. When a truck is tipped at the hub, each pallet from it is temporarily placed in the bay assigned to the depot that has to take the pallet from the hub. The capacity of the bays is limited. More complete, accurate and timely information on incoming and outgoing pallets made available to the hub is expected to improve the use of this limited bay capacity. Examples include levelling the number of pallets in the bay, better anticipating last minute changes in the arrival of pallets to the hub site. Better use of bay capacity will result in trucks spending less time at the hub and better use of hub site capacity in general. In this way, a growth in number of pallets that need to be processed at a hub can, for example, be absorbed without an expansion of the number of warehouses. As a result of pallet handling optimisation and improved bay capacity usage at the hub, an increase in responsiveness of business processes should result in less unpredicted delays for member depot trucks.

The information on incoming and outgoing pallets can be used to better align the tip and load sequence to the pallets that actually come in at the hub, also making better use of the bay capacity. The decision for a specific sequence is expected to reduce the number of delays for outgoing pallets. When a delay of an inbound truck still results in a delay of some outbound truck(s), the delay information should be available in advance. The member depot can be notified about the delay and can therefore better prepare for it.

The decision support from the Depot perspective

From a member depot perspective, improved responsiveness of business processes should lead to increased truck capacity utilisation. Member depots are responsible for picking up at the hub all pallets going to their own delivery territory. For efficient operations, member depots seek to utilise truck capacity well both on the way to the hub and back to the depot.

Depot objective: increased truck capacity utilization

The number of trucks to be sent to the hub is determined by the expected number of pallets to be collected from the hub and delivered by that member depot. Truck capacity on the way to the hub can be used by sending pallets into the network. It is expected that earlier information on the number of trucks that have to be sent to the hub increases the chances for member depots to fill their truck capacity on the way to the

hub. Additionally, the pricing policies within the network should be analyzed with respect to detailed data on pallet flows to the hub. Dynamically changing pricing policies, based on truck capacity usage and aimed at providing incentives to bring in pallets to the network, should increase truck capacity utilization for member depot trucks.

3.2. TheADVANCE decision support software framework

The ADVANCE decision support software framework includes support for both hub and depot operations via the ADVANCE Live Reporter (ALR) and the Depots Collaboration Tool (DCT).

3.2.1. The ADVANCE Live Reporter

The ADVANCE Live Reporter architecture, as depicted in Figure 2, comprises six element types:

1.   At the top of the architecture, end-users are provided with information through a (generic or dedicated) user interface.

2.   The information that is presented through the user interface is assembled by the Analytical Process Engine (APE). The APE is the heart of the ALR, and performs data analysis by using and combining several software modules (later referred to as "blocks"). The analytical process engine may get part of its input from APEs of other organisations, whereby users allow or disallow the sharing of selected information with selected partners.

3.   A business analyst may use the flow editor to deploy the blocks, which are stored in the repository. In order to do so, multiple blocks can be "combined".

4.   A schema editor that is used by a business analyst to define and enhance the information needed by users (and in intermediate process steps).

5.   Operational data that are collected accumulate in the data storage. A data store interface is used to provide the analytical engine with the data required for analysis and to store intermediate results.

6.   At the bottom of the architecture, application interfaces are designed to convert data from existing systems into data the ADVANCE system can use.

When large amounts of data have to be channelled and processed—much of it has to succeed without significant time lag, congestion or data loss. This calls for a solution which is resource-lean, capable of high throughput and allows immediate response.

The reactive paradigm, where dataflows are processed in a push-based manner, can very well suit these requirements [21]. Here, processing blocks react in the presence of incoming data only (i.e., they do not "pull" data from sources but respond immediately if input appears), while output is emitted in a "fire-and-forget" manner. Special mechanisms, such as buffers, can be implemented to prevent data loss or allow synchronous operation—however, this is only done where specifically needed, relieving most of the network of unnecessary computational demands.

Given these expected advantages, the reactive paradigm was chosen to be deployed in the fundamental dataflow framework of the project. Several frameworks implementing
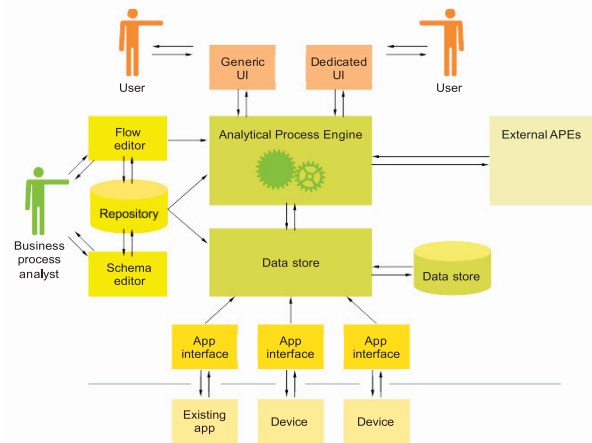

Fig. 2. ADVANCE Live Reporter architecture.

the reactive paradigm or equivalent concepts are already available [22, 23, 24], however, none of these was found to efficiently support the strongly typed dataflows required in our setting. In order to fill this gap, reactive4java was implemented which is now also available as an open source framework (http://sourceforge.net/projects/advance-project/) in itself. The runtime environment built upon the reactive framework offers fundamental functional blocks for processing, channelling and buffering data. More complex processing blocks can be composed of basic blocks, or, if better efficiency is required, application-specific custom blocks can be designed and implemented. In line with similar environments, practice has shown that generic blocks can be used for developing simple conceptual prototypes and switching around dataflows, however, solutions that do eventually find practical use have to rely on the higher efficiency of purpose-made custom blocks. (A more detailed description of the framework is offered in [25].

Dataflows in logistics applications are typically strongly typed, and reliance on structured data models is common in both communication and storage of logistics-related data (although actual structures may be lightweight compared to, e.g., manufacturing). Therefore, type definition and data modelling tools are a must for solution frameworks targeting the logistics sector, both for handling of types within one consolidated network and making data models negotiable when information channels of different participants or networks need to be interfaced. To this end, an XML-based type system was developed in the ADVANCE project that is meant to support all typical logistics-related operations with customisable data models.

When dataflows and their processing methods are concerned, specialisation is advisable. As opposed to inheritance, this means that the type definition starts out with the most complex set (structure) of attributes and removes the ones not needed for the given type. In order to allow such an approach, a data structure was designed in the ADVANCE project to cover all needs within the targeted application range. While this makes the solution kit, as-is, suitable for a

specific application range only, the XML-based type definition makes it easy to redesign the initial type for other application ranges.

In the application context of ADVANCE, semi-automated type negotiation and type handling methods became necessary, both at runtime and during development and debugging of dataflows:

- During development time, users are provided with a verification tool to examine if the constructed dataflow network is consistent with regard to possible input/output types.
- Still during development time, a type probe is at the users' disposal to check the typing of selected wires in the dataflow diagram.
- During build and runtime, types are determined dynamically. This is necessary because typing of data sent within the same dataflow might vary as well.
- In order to meet these requirements, structure-based type inference [26, 27, 28] is employed to perform set operations (e.g., subset, superset, intersection, union) on structured type definitions. More details on the algorithms and related work can be found in [29].
- The reactive runtime environment and the concept of blocks connected by dataflow channels make a graphical flow design environment a natural choice. Therefore, flow design, compilation and execution control were unified behind a graphical IDE front-end in the ADVANCE project.
- The flow editor provides the user with an assortment of processing blocks that have input (except for blocks providing constants, or channelling in data from outside the runtime environment) and output. While the number of input and output is fixed for a given type of block, suitable data switching blocks are provided for merging or replicating data. Construction of composite blocks is also supported, as is implementation of custom blocks where complex functionalities (as specific data processing tasks) would call for the more efficient form of coding block functions right away. The blocks can be connected via bindings that will work as typed data channels at runtime. For this reason, extensive support is given to detect possible type conflicts already during design time: detected type conflicts can be noticed immediately, and a type probe functionality can display the current type of a given binding.
- Once a flow design is complete, it is verified, and possible errors are reported and highlighted in the flow diagram. Upon successful verification, the graphical diagram is transformed into an XML flow description that the runtime environment will use. The execution control environment contains access control with various levels of access rights, and also allows the execution of multiple dataflows by maintaining different execution realms, i.e., separate runtime containers.

In addition to a "bare" framework, the ADVANCE project has also conducted research on techniques that allow more insight into a network's processes than routine surveillance and common aggregation techniques would do. To this end,

ADVANCE examined modelling and prediction algorithms that can:

- detect trends and patterns over a longer timeframe or over a wider area in the network than the attention of human decision makers could capture, and
- make predictions or interpolations where process observability may otherwise be insufficient or impossible (e.g., due to "legacy" business processes, lack of reporting discipline, or the nature of process scheduling).

If a given case of transparency limitation is inherent to the way the network operates, one will encounter the need for estimating data that would otherwise play an important role in decision processes, e.g., the pre-allocation of transportation assets and scheduling of tasks for more efficient or more balanced utilisation of network resources. For this reason, model building and prediction were another focal area in the ADVANCE project. While research was primarily working towards satisfying the requirements of the project's main industrial pilot, genericity and adaptability of the results to other similar scenarios was just as much observed.

Initial surveys regarding transparency gaps in hub-and-spoke logistics networks revealed that the majority of these can be mapped reasonably well onto event chains whose observability may suffer some imperfection at present time, but reconstruction (e.g., for cross-validation in model learning) is possible from historical data. Typically, the events are notifications of shipments entering a given status, while data of interest to be predicted are transportation demands arising from material accumulating at a given point of the network. The term advance order information [30], is commonly used for the former type of events, and in the project, a combined additive and multiplicative model, often used in the AOI context, was assumed for the examined case. Given the quickly changing processes of the targeted scenario, separate models were created for relatively frequent equidistant time points, each of these making a prediction based on the most up to date data available at the given time point. Models comprehensible for personnel were built using machine learning with feature selection.

Short-term predictions and greater visibility of various quantities of interest to the day-to-day running of the freight transportation network are provided. The principal predictions include expected numbers and sizes of pallets for the end of the day, next day or two days in advance at different aggregation levels, for different service levels, for both collection and delivery depots and the hub.

While a specific selection of models and methods was found best fitting the scenario examined in the project, the nature of the approach allows a wider variety of methods to be deployed and tested interchangeably, enabling future users to elaborate choices in algorithms best serving their scenario (details on project results on prediction can be found in [14]).

3.2.2. Depots Collaboration Tool

A logistics network typically has key points where decisions (e.g., resource assignment, scheduling of processes) must be met that have strong impact on the performance of the overall system. In hub-and-spoke networks, these are usually related to the dispatching and timing of inbound and outbound
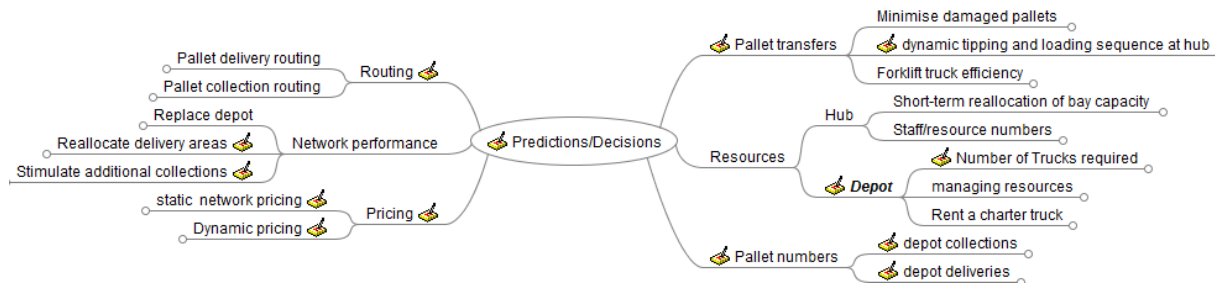
Fig. 3. Initial mind map template

shipments in a way that constraints (e.g., delivery deadlines) are met while resources are used as efficiently as possible (e.g., less deadheading vehicles). In present day networks, these decisions are met by humans, i.e., highly skilled personnel with sufficient routine who can make sound choices even when information needed for the decisions is incomplete, lagging behind relevant events, or simply unreliable.

Current network transparency and complex—often probabilistic and difficult-to-formalise—behaviour of operations would make it risky, if possible at all, to fully automate decision making [31]. However, a human decision maker can still be aided by a decision support system that presents data and offers alternative choices in a way that improves the operator's insight and awareness of the given situation. Logistics networks are subject to frequent changes regarding resource capacities and operation requirements, making it necessary to re-tune the decision support system and allow its constant evolution through user feedback—a challenge that is not met by many decision support systems.

The ADVANCE Depots Decision Tool serves two purposes:
- Mapping of existing decision mechanisms and staff preferences using a purpose-built cognitive modeller;
- Decision support tool putting the modelled knowledge to work with live data received from the network infrastructure.

The underlying principles of modeller and decision tool make it easy to gradually refine the model during use—also in collaboration over several decision points—and keep up with the evolution of the processes as well.

The decision support system implemented in the ADVANCE project follows the Galatean model supporting the way human decisions are met, allowing human interpretation (and evaluation) of the entire decision structure. The approach builds upon the finding that expert decisions are not met by rules but by weighing up degrees of support for premises contributing to a decision [32]. In the Galatean model, simply connected decision trees are set up with membership functions tuned to the perfect conditions for a given decision. Such decision trees can be initially set up by structured interviews with decision makers, and can be subsequently refined. During use, degrees of support are percolated through the tree, giving an overall degree of support for the final decision at its root. As opposed to many other decision support systems, the user can browse through the decision tree and see a comprehensible explanation for the

top-level result. This enables the user to match the suggested choice with his/her own picture of the world and fine-tune the system as needed.

For configuring the ADVANCE system so-called mind maps to represent the evolving knowledge about logistics operations held by expert interviewees were used. Mindmaps constitute one of the most intuitive aids for note-taking, brainstorming, and generally organising ideas, where the central idea (a decision class, for example) is situated in the middle and the subconcepts radiate outwards in ever more detailed subdivisions until the edges are reached with no further child nodes.

A semi-structured interview method was used for gathering requirements based on a schedule derived from an initial mind-map template [32] (see Figure 3). Analyses of the current and desired operational decisions were recorded in an emerging ADVANCE mind map template, which was then used to inform subsequent interviews. Interview data were analysed, and added to the mind map, with the mind map being periodically validated by in-depth interviews scrutinising the content and structure. This iterative process led to increasingly stable decision hierarchies showing how input data relates to output decisions. The complete mind map resulting from the interviews was used to specify the detailed end user and industrial functional requirements, covering many operations for improving efficiency in hub-and-spoke networks.

Building decision support systems around psychological models ensures that the most relevant data is used, including empirical sources directly accessible by the machine and those coming from human experience, and furthermore ensures the bridging of gaps in human interpretability and feedback to the decision support system.

4. Conclusion

Often, information in manufacturing chains is used locally in standardized work processes (e.g., shipment tracking data for a given region) and is, for the span of its existence, available in some structured form. In a number of cases (e.g., retailing), work processes are well-defined, and even though little information is recorded, it would be fairly straightforward to introduce structured data for further processing (e.g., storing of check-out steps distinguishing unique items or individual batches a given piece of merchandise belongs to). The amount of data, as well as their

relatively flat nature, requires them to be pre-processed, if not for storage, then, at least, for efficient sharing. Depending on the form of the raw data, the required depth of processing may range from simple aggregation to the extraction of patterns or data mining.

Even if relevant information is highlighted, this is rarely enough to directly support human decisions, since operators can hardly overview the data sets and extract relevant information to the degree the decisions would require. Therefore, computational intelligence is needed to analyze the data, detect patterns and build models (in essence, application-related assumptions), and eventually meet predictions regarding tendencies or effects of certain decisions. This can be then delivered to human operators in the form of suggestions or indicators which can be assessed by the operating personnel. In most cases, the models need to be refined over several iterations (or have to undergo continual changes for certain dynamic environments), relying on feedback from both data and humans (e.g., plausibility measures).

The ADVANCE decision support software framework relies on machine learning and cognitive modelling to deliver a practical solution that is both specialised to the industrial case study and is general enough to be adapted in other logistics scenarios. In fact components of the framework are independent and could be used in entirely different domains, where the problems to be solved have similar characteristics.

References

[1] Christopher, M., (1992), Logistics and Supply-chain Management, Pitman publishing, London.
[2] Lambert, D. and Cooper, M., (2000), Issues in Supply-chain Management, Ind. Marketing Management, Vol.29, pp. 65-83.
[3] Burgess, R., (1998), Avoiding Supply-chain Management Failure: Lessons from Business Process Re-engineering, Int. Journal of Logistics Management, 9/1, pp. 15-23.
[4] Mentzer, J., DeWitt, W., Keebler, J., Min, S., Nix, N., Smith, C., and Zacharia, Z., (2001), Defining supply-chain management, Journal of Business Logistics, 22/2, pp. 1-25.
[5] Norek, C.D.,Pohlen, T.L., (2001), Cost Knowledge: A Foundation for Improving Supply-chain Relationships, Int. Journal of Logistics Management, 12/1, pp. 37-51.
[6] Gunasekaran, A., Ngai, E.W.T., (2004), Information systems in supply-chain integration and management, European Journal of Operational Research, 159/2, pp. 269-295.
[7] White, R. and Pearson, J., (2001), JIT, system integration and customer service, International Journal of Physical Distribution & Logistics Management, 31/ 5, pp. 313-333.
[8] Bertolo, S., (2006), From Intelligent Content to Actionable Knowledge: Research Directions and Opportunities Under the EU's FP7, 2007-2013, in R. Meersman, Z. Tari (Eds.), OTM 2006, LNCS 4276, pp. 1125-1131.
[9] Michel, R: RFID, (2005), Where's the Beef?, Modern Materials Handling, 60(2), pp. 29–31.
[10] Dejonckheere, J.; Disney, S. M.; Lambrecht, M. R.; Towill, D. R., (2003), Measuring and Avoiding the Bullwhip Effect: A Control Theoretic Approach, European Journal of Operational Research, 147(3), pp. 567–590.
[11] Jansen-Vullers, M.H.; van Dorp, C.A.; Beulens, A.J.M., (2003), Managing Traceability Information in Manufacture. Int. Journal of Information Management 23(5):395–413.
[12] Monostori, L.; Ilie-Zudor, E.; Kemény, Zs.; Szathmári, M.; Karnok, D.; (2009), Increased transparency within and beyond organizational borders by novel identifier-based services for enterprises of different size. CIRP Annals – Manufacturing Technology, Vol. 58, Nr. 1, pp. 417–420.

[13] Kemény, Zs.; Ilie-Zudor, E.; van Bolmmestein, F.; Kajosaari, R.; Holmström, J. (2007), State of the art in tracking-based business, Public deliverable D3.1, version 1.3, EU FP6 STREP project TraSer (FP6-2005-IST-5).
[14] Welch, P. G.; Kemény, Zs.; Ekárt, A.; Ilie-Zudor, E., (2012), Application of model-based prediction to support operational decisions in logistics networks. Proc. of the 3rd Workshop on Artificial Intelligence and Logistics, Montpellier, France; SFB/TR 8 Report No. 031-08/2012, pp. 25–30.
[15] Zäpfel, G., Wasner, M. (2002), Planning and optimization of hub-and-spoke transportation networks of cooperative third-party logistics providers, Int. J. Production Economics Vol. 78, pp. 207–220.
[16] Wieberneit, N. (2008), Service network design for freight transportation: a review, OR Spectrum, 30/1, pp. 77–112.
[17] Jardim-Gonçalves, R., Popplewell, K. and Grilo, A. (2012). Sustainable interoperability: The future of internet based industrial enterprises, Computers in Industry 63(8): 731–738.
[18] Agostinho, C., Jardim-Gonçalves, R. (2009). Dynamic Business Networks: A Headache for Sustainable Systems Interoperability, On the Move to Meaningful Internet Systems: OTM 2009 Workshops, Lecture Notes in Computer Science Volume 5872, 2009, pp 194-204.
[19] Chen, D., Doumeingts, G., Vernadat, F. (2008). Architectures for enterprise integration and interoperability: Past, present and future, Computers in Industry 59(7): 647–659.
[20] Ilie-Zudor, E.; Ekárt, A.; Kemény, Zs.; Karnok, D.; Buckingham, C.; Jardim-Goncalves, R.: Information Modeling and Decision Support in Logistics Networks. Proc. of the 5th Int. Conf. on Experiments, Process, System Modeling, Simulation, Optimization, Athens, Greece, July 3-6, 2013, Vol. I, pp. 279-286, ISBN 978-618-80527-1-0.
[21] Elliott, C.M. (2009), Push-pull functional reactive programming, Proc. of 2nd ACM SIGPLAN symposium on Haskell, Haskell'09, ACM, New York, NY, USA, pp. 25–36, ISBN 978-1-60558-508-6.
[22] Liberty, J., Betts, P. (2011), Programming Reactive Extensions and LINQ. Apress,ISBN 978-1-4302-3747-1.
[23] National Instruments (2012), LabVIEW System Design Software. Official website, URL: http://www.ni.com/labview/.
[24] Chambers, C., Raniwala, A., Perry, F., Adams, S., Henry, R.R., Bradshaw, R., Weizenbaum, N. (2010), FlumeJava: easy, efficient data-parallel pipelines, Proc. 2010 ACM SIGPLAN conference on Programming language design and implementation, PLDI'10, New York, NY, USA, pp. 363–375.
[25] Karnok, D., Kemény, Zs. (2012), Framework for building and coordinating information flows in logistics networks, Proc. 14th int. conf. on the modern inf. tech. in the innov. proc. of the industrial enterprises, MITIP 2012, Budapest, Hungary, October 24–26, 2012, pp. 551–560.
[26] Kaes, S. (1992), Type inference in the presence of overloading, subtyping and recursive types, Proc. 1992 ACM conference on LISP and functional programming, LFP '92, ACM, New York, NY, USA, pp. 193–204.
[27] Odersky, M., Sulzmann, M., Wehr, M. (1997), Type inference with constrained types, Proc. of the Fourth International Workshop on Foundations of Object-Oriented Programming.
[28] Palsberg, J. (1995), Efficient inference of object types, Inf. Comput., Vol. 123, No. 2, pp. 198–209, DOI: 10.1006/inco.1995.1168.
[29] Karnok, D., Kemény, Zs. (2012), Definition and handling of data types in a dataflow-oriented modelling and processing environment, Proc. 14th int. conf. on the modern information technology in the innovation processes of the industrial enterprises, Budapest, Hungary, Oct. 24–26, pp. 561–574.
[30] Haberleitner, H., Meyr, H., Taudes, A. (2010), Implementation of a demand planning system using advance order information, Int. J. of Prod. Economics, 128/ 2, pp. 518–526.
[31] MacCarthy, B., Wilson, J., eds. (2001), Human Performance in Planning and Scheduling, Taylor & Francis, URL: http://books.google.com/books?id=0wBLHGX1_WIC.
[32] Buckingham, C.D., Buijs, P., Welch, P.G., Kumar, A., Ahmed, A., (2012), Developing a cognitive model of decision-making to support members of hub-and-spoke logistics networks, Proc. 14th int. conf. on the modern inf. tech. in the innov. proc. of the industrial enterprises, Bp, Hungary, pp.14–30.