5th International Conference on Experiments/Process/System Modeling/Simulation/Optimization
5th IC-EpsMsO
Athens, 3-6 July, 2013

INFORMATION MODELLING AND DECISION SUPPORT IN LOGISTICS NETWORKS

Elisabeth Ilie-Zudor[1], Anikó Ekárt[2], Zsolt Kemény[1], Dávid Karnok[1], Christopher Buckingham[2], Ricardo Jardim-Goncalves[3]

[1]Institute for Computer Science and Control, Hungarian Academy of Sciences
Kende u. 13–17, 1111 Budapest, Hungary
e-mail: {ilie, zsolt.kemeny, karnok.david}@sztaki.mta.hu

[2]Aston University
Aston Triangle, Birmingham, B4 7ET, United Kingdom
e-mail: {c.d.buckingham, a.ekart}@aston.ac.uk

[3]Universidade Nova de Lisboa, Centre of Technology and Systems
2829-516 Caparica, Portugal
e-mail: rg@uninova.pt

Keywords: logistics networks, hub-and-spoke, operational decisions

Abstract. *The paper presents challenges related to data handling, interoperability, transparency, as well as process modelling and prediction in logistics networks. While some of the interoperability and transparency problems are likely to surface in other business networks as well, the emphasis of certain specific problems (e.g., relation of material and information stream) are characteristic to the targeted spectrum of logistics operations. The paper summarises the solution concept and presents first practical insights gained with a recently introduced industrial pilot.*

1 INTRODUCTION

Many of today's challenges in business, transportation and manufacturing are answered by the emergence of networked organisational structures that have the potential of combining local flexibility and agility with the "mass advantage" of large capacities present either network-wide or at some key high-capacity components of the network. Logistics has been a notable proving ground of such structures and processes, resolving in the last couple of decades preference conflicts that seemed insurmountable before. Certain types of hub-and-spoke logistics networks, for example, utilise re-bundling of shipments to offer flexible, affordable and fast forwarding of less-than-truckload (LTL) consignments[1,2].

However, the emergence of networked structures itself is not without new challenges either. The majority of these problems can be traced back to three key aspects:

- *Transparency—seeing each other's relevant processes*. Today's networked organisational structures are made possible by information being more reliably available than before, allowing better—and timely—decisions to be taken regarding physical or business processes. This, however, requires a certain degree of *transparency*, i.e., sharing of information between selected partners or across the entire network. Since the latter has its costs (equipment and labour expenses) and risks (disclosing sensitive data), network members shall see information sharing as a form of *investment*, the return showing in extended opportunities and better support for local decisions. Once a member takes the step of sharing information, this must be carried out with a certain degree of reliability, i.e., with sufficient adherence to standards agreed upon regarding accuracy, timing, etc. From the opposite point of view, the participation of a member can be fruitful if the network is able to exploit shared data of de-facto reliability.
- *Interoperability—understanding each other*. Typically, networks do not homogeneously develop in all locations concerned, but are built of members who have already existed long before and established their own processes, data models and interpretations of their own. These formal models and interpretations need to be properly matched to allow connections across organisational borders. Most chores of establishing common ground are related to new members appearing in the network, however, the efforts of *maintaining* interoperability in everyday operation (i.e., adherence to common standards, network-wide introduction of data model or process changes) should not be underestimated either[3].

- *Handling large amounts of data—comprehension beyond a member's horizon*. Process transparency—especially in a network where loads of potentially relevant business processes can run simultaneously—introduces a new burden simply because data become available in amounts not encountered before, while their content of useful information remains limited. Recent years witnessed the emergence of data mining techniques and parallel processing to handle massive low-level data[4]. In some cases, data are *dispersed over the network*, and forwarding the right piece of information to the right location becomes another challenge—this is a typical phenomenon in many logistics networks. Meeting these challenges, on the other hand, greatly extends the horizon of understanding and predicting processes on a larger scale, and opens up new opportunities for operational and strategic decisions to the benefit of members as well as the entire network.

Focusing on an area in the logistics industry, the paper presents a project providing tools to address the above groups of challenges. Funded within the EU 7th Framework Programme, the ADVANCE project selected a nationwide logistics network as its object of research and proving ground, and the paper presents the development results and findings in this context, not omitting considerations of reusability.

The remainder of the paper is organised as follows. First, the logistics scenario specific to the project is presented, and the addressed challenges are highlighted. Second, a conceptual framework of the solution is drawn. Next, an outline of the implemented solution platform is given, and conditions of an industrial pilot are summarised. Finally, first experience with the adaptation to the pilot case is summarised. While ADVANCE is already nearing its completion and the paper is more than a work-in-progress report, it should be noted that the pilot application is currently being introduced, and much practical experience is yet to follow in the next months.

2 PROBLEM STATEMENT—THE HUB-AND-SPOKE CASE

In the past, shipping of small—less-than-truckload (LTL)—consignments was perceived to reflect conflicting requirements as they had to be fast *and* affordable. In the last couple of decades, a new form of logistics networks has emerged that achieve economical vehicle utilisation levels and range of operation by re-bundling shipments at well-defined points of the network[1,2]. Material processing is built around a *hub-and-spoke* structure with the spokes performing local pick-up and delivery, while high-throughput hubs support the re-bundling of consignments. In practice, the following three stages are characteristic for these networks (see also Figure 1 for an example):

- An *inbound phase* consisting in collection of consignments by the spokes and dispatching them to the hub (i.e., bundling consignments by area of origin);
- *Re-bundling* by destination in the hub facility;
- An *outbound phase* where spokes clear the re-bundled shipments from the hub and deliver them to their destination.

Hub-and-spoke logistics networks can benefit from a *division of activities along skills and capacities*. Inbound and outbound steps are best carried out by subcontractors that can more flexibly adapt to local conditions of an area, and mostly utilise the same fleet for their section of inbound and outbound movement. Re-bundling in high-throughput hubs is then performed by a large central player. This composition of the network enables participants to *combine complementary strengths*, most notably, local adaptivity and high performance. Nevertheless, such multi-level hub-and-spoke structures do present efficiency challenges, many of which can be traced back to the three main problem areas of present-day networked enterprises:

- *Transparency* may be less than perfect for the given requirements. While it is usually left to the central player to elicit the subcontractors' willingness to share the data needed—and the central player does shape its business policy accordingly—it depends very much on the subcontractors how much discipline is exercised in de-facto data sharing practice. Despite the participants' willingness, operations may suffer from incomplete or erroneous information if data are committed with imperfect timing (information stream may lag behind the material stream) or accuracy (incorrect data may lead to wrong decisions).
- *Interoperability* may need to be established at additional cost. While a central player is usually powerful and organised enough to provide subcontractors with a standard and mandatory IT solution (thus leaving the major part of IT adaptation to the subcontractor), this may not be the case in the fusion of several networks of comparable size where each of them has already established its own data models and procedures.
- The *large amounts of low-level data* present as much a challenge as they mean an opportunity. In multilevel logistics networks, low-level data are often dispersed over the entire network, and processing them for long-term forecasts or support of strategic decisions does not consist in tackling their outstanding amounts alone. In the latter case, pre-filtering raw data and forwarding the adequate selection to the correct point of the network at the right time is perceived as a new challenge.

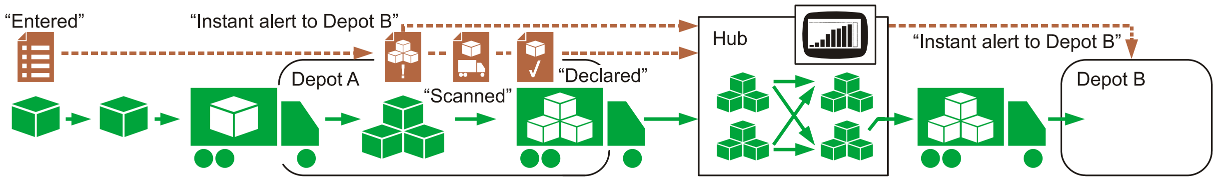

Figure 1. Typical material flow and information flow in a hub-and-spoke logistics network

3 SOLUTION CONCEPT OF THE ADVANCE PROJECT

In order to help logistics networks overcome the aforementioned difficulties, the ADVANCE project aims to provide a solution platform that a given network can tailor to its own needs in bridging information gaps related to organisational connections.

- The ADVANCE platform offers a *dataflow framework* providing resource-efficient means for reproducing existing data flow structures and prototyping new structures or procedures.
- The dataflow framework features a *reactive engine* with runtime type handling support for operations on dataflows with application-specific data types.
- A *graphical flow design interface* supports the design of custom dataflows, provides an extensive library of functional blocks for data switching and processing, and supports dataflow design with advanced *type inference tools* that help flow designers locate data model problems prior to running the flows. Type handling tools also make it easier to negotiate different data models when individual participants or networks merge their dataflows and need to map each other's data models without compromising resultant cross-network transparency.

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 examines 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).

While ADVANCE has aimed to provide system architects and solution developers with tools for network-wide solutions, the same holds just as much for the *operational level* of everyday business in a logistics network:

- ADVANCE has implemented front-end *user interface elements*, now undergoing field tests in an industrial pilot, which present data already found in the system, as well as added by the project's novel modelling and prediction components, to enhance an operator's overview and awareness of the current process situation and supply data for decisions in manageable form.
- ADVANCE also provides a *decision modelling and support* tool which maps possible choices and their degrees of support close to the way human decisions are made. This allows the efficient elicitation of knowledge already present as decision routine, and suitable evaluation methods allow feedback for subsequent improvement of the knowledge support structures. This makes the decision support tool *evolvable*, contrary to many earlier solutions.

4 IMPLEMENTED COMPONENTS

4.1 Runtime environment: reactive framework

In the targeted logistics application, 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[5]. 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 the reactive paradigm or equivalent concepts are already available[6,7,8], 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 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.

While the paper would not allow a more detailed description of the framework, the reader is kindly referred to an earlier publication associated with the ADVANCE project where related literature and available technologies are covered in a broader overview[9].

4.2 Type system and type handling

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[10] are an absolute 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 be able 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 that is assumed 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*[11,12,13] is employed to perform set operations (e.g., subset, superset, intersection, union) on structured type definitions. While the volume of this paper does not have the space to accommodate more details, the reader is referred to an earlier publication related to the project which presents the algorithms and related work in more detail[14].

4.3 Model building and prediction

If a given case of transparency limitation is inherent to the way the network operates, one will encounter the need for estimating/forecasting 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 type 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 (AOI)*[15] 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 short equidistant time points (i.e., below the magnitude of an hour), each of these making a prediction based on data available at the given time point. Models were trimmed down with attribute selection and kept mathematically simple to reduce computational demand and keep the principles comprehensible for personnel with limited theoretical background. While one given 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. Here the reader is referred to a further publication which presents project results on prediction in detail[16].

4.4 Decision support and knowledge elicitation

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 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.

Present-day 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[17]. 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, nevertheless, 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 expert systems.

In order to overcome this shortcoming, the decision support system implemented in the ADVANCE project follows the *Galatean model* that is in close keeping with the way human decisions are met, and enables 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[19]. 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, it does not have to remain at that—if required, 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 further detail, the reader is referred again to another publication focusing specifically on the application of the Galatean model in the context of the ADVANCE project[19].

4.5 Graphical developing and runtime control environment

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 (see also Figure 2).

The flow editor provides the user with an assortment of processing *blocks* that have inputs (except for blocks providing constants, or channelling in data from outside the runtime environment) and outputs. While the number of inputs and outputs 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 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.

4.6 Possibilities of front-end user interfaces

Front-end user interfaces for operating personnel—who use runtime functionalities built on the flows—are largely independent of the integrated design and execution control environment (i.e., flow editor and flow engine). As a consequence, the particular design and structure of runtime user interfaces are largely subject to criteria of the targeted application scenario regarding many design and implementation choices. Web-based interfaces, however, represent the most natural choice, as the flow execution environment runs in a web container already (specifically, Tomcat was used during development of the current implementation). To date, only a limited set of front-end elements was produced in a generic test implementation, nevertheless, it is planned that parts of the (proprietary) industrial pilot will serve as a prototype for front-end GUI elements offered with the open-source dataflow framework.

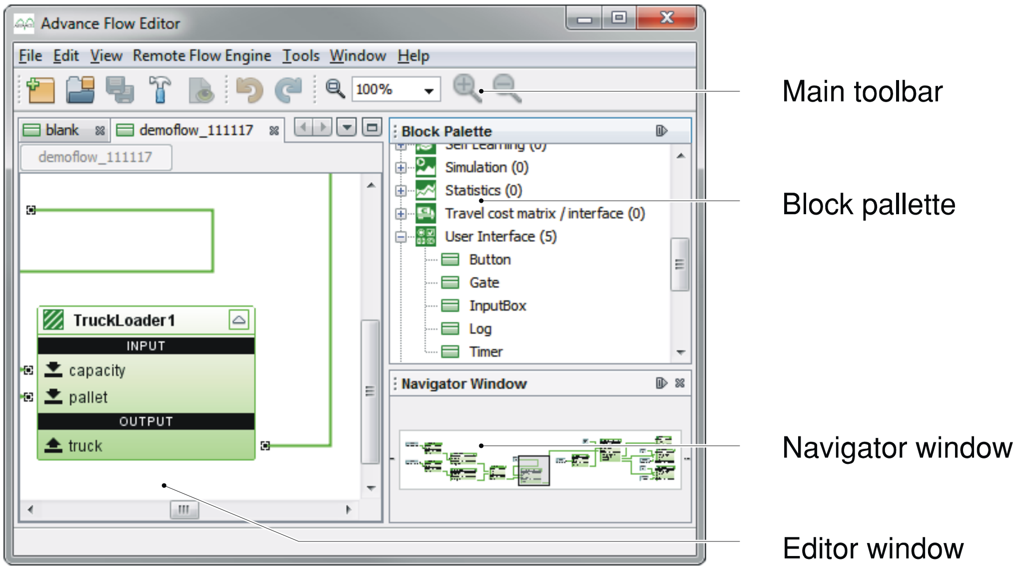

Figure 2. Screenshot of the ADVANCE flow editor showing functional blocks in a dataflow

5 PILOT APPLICATION—FIRST PRACTICAL FINDINGS

5.1 Targeted user groups

In the final year of the ADVANCE project, an application pilot was implemented and is currently undergoing field tests in a real-life logistics environment, with most components being deployed at a major hub of a nationwide multilevel logistics network. The following targets are pursued with the pilot:

- Improving data transparency by better utilising data that is already there;
- Bridging transparency gaps and causality-related lack of information by predictions (day-by-day for the next few days, and a regularly updated forecast during the current day);
- Providing selected groups of runtime users (mainly hub personnel) with new views at current and predicted status of shipments and resources that better facilitate their routine decisions and actions;
- Providing means for issuing instant alarms and display highlights in critical situations;
- Suggesting actions to be taken in critical situations (i.e., an implementation of the decision modelling and support presented earlier).

Additionally, the pilot application is also subject to constraints inherent to the use case:

- Processes in logistics networks with tight deadlines are elaborate results of a longer evolution. Therefore, the new solution should be introduced with as little intervention with existing processes as possible, and allow manual override or reverting to previously established routine.
- Each particular user group should be assumed to have its own specific requirements regarding response time (both for data displayed and for back-end calculations), useful appearance, granularity and selected amount of displayed content. It is also important to observe how much time various types of users have available to comprehend the displayed content and interact if necessary.
- Once the implemented solution meets approval in practice, it is expected to evolve along with the underlying processes. Therefore, a certain amount of flexibility needs to be reserved for future development (space for additional controls, views, etc., as well as some flexibility with data models). While it is difficult to give a formal guarantee for flexibility, whenever design and implementation decisions are taken, possibilities for future enhancement should be taken in consideration.

Prior to design and implementation of the pilot solution, a thorough survey was conducted with personnel operating the network at different points and at different levels in the decision hierarchy. As a result of the interviews, points of relevance were identified where the ADVANCE solution could bring the most useful added value. From these, three user groups were identified which will be covered by the initial pilot:

- *Hub personnel at the top level* of operational decisions observe inbound and outbound processes for the entire hub. While they decide on instant action most of the time during a shift (this kind of activity being best supported by minimal-interaction desktop screens), they may also examine forecasts and progress of shipments on the order of magnitude of days (best served by desktop screens where details can be unfolded on demand) to prepare for major actions in the coming days, if any are necessary.
- *Hub personnel at warehouse level*, guiding unloading/loading for one warehouse (or a pair of them), is exposed to extreme time pressure at peak throughput and is best served by a web interface on a mobile

device (typically a tablet computer). The interface should efficiently present all data needed for decisions at a single glance, therefore, *colours and bar lengths prevail* (as opposed to numbers and text).

- *Depot operators* at subcontracted collection and delivery partners are, in many regards, in a similar situation to top-level hub personnel. The desktop screens for this user group will share many characteristics with top-level hub interfaces, however, perspective and displayed data will differ from those meant for the entire hub.

Existing IT infrastructure made web interfaces a feasible choice, including the screens for mobile devices. The back-end is built to exploit known data more efficiently, as well as learn process models and make predictions of selected values with frequent updates for the current shift and once per day for the next few days. A considerable part of the implemented solutions is specific to the application case and the associated business processes and, therefore, has to be considered proprietary development. Nevertheless, findings of the application pilot will contribute to the improvement and further development of the generic, open-source components.

5.2 First findings

While a full evaluation of the application pilot is still in progress, first findings could already be gained—to a considerable part related to the application scenario itself.

During research of the specific application case, the evolution of business processes and models were examined, also in view of historical data of the past years. These revealed that a certain "core" of data models and business processes tends to remain constant over the years, and there is little variation expected from network to network. This implies that mapping of different data models may be less demanding for a considerable part of all structures used, even for a fusion of entire networks, than initially assumed, and many of the elaborate data model matching methodologies may only be required in specific problem areas where the common aspects of business processes are too few to establish common ground. Examples for these potential problem areas are service types and classes of shipment events in a complex network. Some variation regarding these may occur even within the same network if adaptation to local needs is more worthwhile than following a network-wide rigid standard.

Another, more difficult, challenge is the varying adherence to formal process specifications. In retrospect, it was found that IT developments—which do require a formal specification—prefer to rely on business process descriptions that are already formalised to some extent (which may decrease the efforts of transforming an application scenario into software specification). Logistics practice is, however, full of deviations from formal procedures, for a number of reasons:

- Fulfilling of binding delivery guarantees has priority over adherence to formal rules. Therefore, *ad-hoc measures* are taken in critical situations which can establish themselves as routine procedures later on. This was perceived as the main reason for established processes being much more complicated and full of—apparently random—exceptions than solutions started in theory, with a "clean sheet".
- Interpretation of terms and specifications may vary, and some of these local differences appear to be self-supporting for a longer time. Any solution designed for logistics networks should be able to tolerate such variations to a reasonable degree.
- As long as decisions are finally taken by humans, operations are likely to be plagued with errors and lapses of discipline (either in taking action or reporting it to the system). These can be mitigated by internal data correction and sanitisation (whenever correct data can be deduced with some workaround), and by automatic feedback that *motivates* participants to perform better (e.g., by presenting a "what if" picture highlighting the benefits lost due to incorrect action).

It should be clear that even if correction and motivation measures are implemented, the system needs to be fault-tolerant and able to adapt to unexpected intervention. Also, no logistics process can be perfectly modelled due to the considerable number of *ad-hoc* decisions taken. Therefore, instead of attempting the impossible in obtaining a fully-fledged model of the network, more effort must be spent on making the system capable of capturing unexpected changes and following the actual state of processes as closely as possible.

6 CONCLUSION AND FUTURE WORK

The paper reported on the EU FP7 project ADVANCE which targeted gaps in data interoperability and transparency in logistics networks. To this end, a dataflow design and runtime framework was implemented with extensive support for strongly typed dataflows and computationally lean push-based execution of functions. In addition, the scope of the project also included research of model building and prediction applicable to logistics networks on the operational level. Being in its last year, the ADVANCE project is now introducing an industrial pilot tailored to the needs of a specific hub-and-spoke network. With first practical findings regarding operation and weighting of challenges being already available, it was found that a consolidated network rarely poses major challenges related to data models. Instead, diversity of de-facto practices on the operational level (non-formalised on-the-spot decisions, variations in interpretation and strictness in data transparency) proved the main challenge, emphasising that this class of logistics operations often requires spontaneous adaptation and deviation from rigid rules. This implies practical limits of formalising business logic and sheds light on the importance of reliably capturing unexpected events in the system. At the current stage of the project, the major part of practical experience is yet to be gained, but savings are expected due to the solution allowing more efficient asset utilisation. Findings of the industrial pilot will certainly contribute to the selection of supported functionalities in the solution framework offered by the project.

ACKNOWLEDGMENTS

The research leading to these results has received funding from the European Union 7th Framework Programme (FP7/2007-2013) under grant agreement: ADVANCE nº257398 (http://advance-logistics.eu/) and IMAGINE nº 285132, (http://www.imagine-futurefactory.eu/).

REFERENCES

[1] 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.

[2] Wieberneit, N. (2008), "Service network design for freight transportation: a review", OR Spectrum, Vol. 30, No. 1, pp. 77–112.

[3] Ferreira, J., Agostinho, C., Sarraipa, J., Jardim-Goncalves, R. (2011), "Monitoring Morphisms to Support Sustainable Interoperability of Enterprise Systems", in: Meersman, R., Dillon, T., Herrero, P. (eds.): On the Move to Meaningful Internet Systems: OTM 2011 Workshops. LNCS, Springer Berlin, pp. 71–82.

[4] Choudhary, A.K., Harding, J.A., Tiwari, M.K. (2009), "Data mining in manufacturing: a review based on the kind of knowledge", Journal of Intelligent Manufacturing, Vol. 20, No. 5, pp. 501–521.

[5] 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.

[6] Liberty, J., Betts, P. (2011), Programming Reactive Extensions and LINQ. Apress, 1 ed., ISBN 978-1-4302-3747-1.

[7] National Instruments (2012), LabVIEW System Design Software. Official website, URL: http://www.ni.com/labview/, last accessed May 2013.

[8] 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.

[9] Karnok, D., Kemény, Zs. (2012), "Framework for building and coordinating information flows in logistics networks", *Proc. 14th int. con. on the modern inf. tech. in the innov. proc. of the industrial enterprises, MITIP 2012, Budapest, Hungary, October 24–26, 2012*, pp. 551–560.

[10] Hosoya, H., Pierce, B.C. (2003), "XDuce: A statically typed XML processing language", ACM Trans. Internet Technol., Vol. 3, No. 2, pp. 117–148, DOI: 10.1145/767193.767195.

[11] 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.

[12] Odersky, M., Sulzmann, M., Wehr, M. (1997), "Type inference with constrained types", *Proc. of the Fourth International Workshop on Foundations of Object-Oriented Programming (FOOL)*.

[13] Palsberg, J. (1995), "Efficient inference of object types", Inf. Comput., Vol. 123, No. 2, pp. 198–209, DOI: 10.1006/inco.1995.1168.

[14] 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, MITIP 2012, Budapest, Hungary, October 24–26, 2012*, pp. 561–574.

[15] Haberleitner, H., Meyr, H., Taudes, A. (2010), "Implementation of a demand planning system using advance order information", Int. J. of Production Economics, Vol. 128,No. 2, pp. 518–526.

[16] Welch, P. G., Kemény, Zs., Ekárt, A., Ilie-Zudor, E. (2012), "Application of model-based prediction to support operational decisions in logistics networks", *Proc. 3rd Workshop on AI and Logistics (AILOG-2012), 20th European Conf. on AI (ECAI), Montpellier, France, August 28, 2012*, pp. 25–30.

[17] MacCarthy, B., Wilson, J., eds. (2001), *Human Performance in Planning and Scheduling*, Taylor & Francis, URL: http://books.google.com/books?id=0wBLHGX1_WIC

[18] McDaniel, M., Busemeyer, J. (2005), "The conceptual basis of function learning and extrapolation: Comparison of rule-based and associative models", Psychonomic Bulletin & review, Vol. 12, No. 1, p. 24.

[19] 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. con. on the modern inf. tech. in the innov. proc. of the industrial enterprises, MITIP 2012, Budapest, Hungary, October 24–26, 2012*, pp. 14–30.