3.1 Publishable summary

ADVANCE

Summary of Project Context and Objectives

Logistics networks typically accumulate OVER 1 BILLION new items of information each month (customer orders, pallets, trailer images, postcodes, depot data, GPS tracking of vehicles, etc.), generated every minute of every day by thousands of pallets travelling on hundreds of trailers for more than one million customers scattered across hundreds of thousands of postcodes, each with multiple different service requirements. Every second, thousands of data items come on stream at any point of the network and need analysis to guide short-term decisions about lorry deployment (within minutes) as well as longer term plans for carrying capacity.

Experience, however, shows that the mere presence of large amounts of low-level data can hardly be of help in network operation where data has to be available at the right place, at the right time, and in the right form to be of any use (i.e., rarely in the particular way a given piece of information is already present in the system). This holds even more for networks where a certain degree of organisational and operational heterogeneity exists (e.g., due to the network being composed of independent actors, or the contractual framework allows for more diversity in business and data handling practices). The need for thoroughly (re)structured and transformed information poses a challenge in view of the amounts of data to be processed: The patterns and dependencies that exist in the 50 million or more data elements created daily can only be meaningfully processed by intelligent data-mining approaches linked to strategic decision making based on longer term analyses of billions of pieces of information—well noted, this goes far beyond handling current or recent information only.

The ADVANCE project targeted these information-related challenges in a specific class of logistics applications (with the possibility of adaptation to other fields of usage): in hub-and-spoke logistics networks made up of one large central player and several local subcontractors. Main targeted outcomes of the ADVANCE project were:

- An open-source dataflow design and execution framework to enable users to efficiently develop prototypes tailored to their specific needs—including various forms of interfaces for actors at different points of the network;
- Advances in the underlying theoretical knowledge, with emphasis on model building, prediction and optimisation (in order to properly prepare data for decision support), as well as cognitive modelling of human decisions (in order to bring the data and associated alternative choices to the user in the form that actually assists them in making operational decisions);
- Practical implementation of the theoretical advances that can be deployed at suitable points of the dataflows developed in the design and execution framework.

With these targeted outcomes, the ADVANCE project was expected to improve network operation efficiency (reduce excess costs, deadheading, congestions or foreseeable risks

of these) where process transparency or observability would otherwise be inadequate, partly inherently to the nature of the processes themselves.
Though not directly specified, it was also expected that outcomes of the ADVANCE project would shed light on existing transparency limits within a logistics network, and provide valuable insight for long-term, strategic decisions on infrastructure dimensioning and IT investments.
Implemented results of the ADVANCE project were to undergo field tests in an application pilot, conducted by an industrial consortium member operating a hub-and-spoke logistics network.

Summary of S&T Results/Foregrounds

Having started in October 2010, the project has now completed all work planned, and achieved additional results which are likely to lead up to long-term improvements within the logistics network and work processes of the industrial consortium member. While many additional results were based on experiments with the industrial partner's confidential data (including data models), the findings of these have found their way back into the open-source solution framework and will be available to new users outside the current sphere of usage.

In-depth insight in the application field

The first months of the project, as well as the time of preparation before official kick-off, were spent conducting studies and interviews in close collaboration with the main industrial partner of the project. Findings were substantial for the success of the project, and refined initial assumptions to be more in keeping with the reality of industrial practice:

- The analysis of interviews and properties of business models helped us identify the exact class of application scenarios and the challenges that can be addressed. As solution algorithms often have to be exactly tailored to the given problem, such deep knowledge is essential for meeting a proper choice.
- On-site interviews have revealed the specifics of daily operation—these were especially valuable in selecting the problems that can be efficiently tackled by introducing advanced IT. We have examined and identified points where fast decisions are met by long-term routine, and the overhead of "smart" support would obstruct the process. Keeping clear of such interference will not merely bring an efficiency advantage—it will forego the serious pitfalls some other experiments did not avoid.
- The detailed analysis of business models and corporate behaviour has already yielded notable results, and more lies still ahead—sorting out limitations and typical business behaviour will enable us to set correct strategic goals not even the industry could identify so far.

Strengthening the fundaments—test bed and solution framework in one

During the project, the ADVANCE framework provided a test bed for experiments and field tests of putting theoretical results at practical work—now, by the end of the project, it was refined to a solution framework other users can build their own solutions on. Much of the year 3 efforts regarding the solution framework were exerted to incorporate findings of field tests into the fundaments laid down in previous project years—the domain-specific practical knowledge gained in a sensitive industrial environment was now safely brought back to improve the solution base which will be accessible to the public.
The fundamental framework (the ADVANCE engine) is in essence an environment where the solutions of our targeted application domain can be built and run. Assembling a specific solution would mean a large assortment of technical details to care about: we would have to sort out and properly match data models of the information handled, we would have to ensure proper timing of operations or sequences, etc. Many of these chores pose no technical challenge in the individual case, but require lots of almost automatic development work. The purpose of the ADVANCE engine and the associated development environment is to take this unnecessary burden off the shoulders of developers and researchers, shorten development time, reduce the risk of errors where this can be done automatically, and let system architects and process analysts concentrate on tailoring the solution to their specific case. The ADVANCE framework allows the easy design and testing of data flow networks from the perspective of network structure, business logic, as well as data models. Already the preliminary domain examinations revealed the importance of data models in the targeted application domain. To this end, the ADVANCE framework is augmented by a set of ready-made data models that adapt to the logistics domain in general, and can serve as a starting point of easier customisation.
With the design, development and testing of the ADVANCE engine already completed in previous project years, the final stage of the project was a time of “distilling” safely re-usable additions from the domain-specific and case-specific findings of the industrial pilots tested in the last project year. In a number of cases, this was far from trivial—field tests did require proprietary solutions at times, to a major part to speed up prototype deployment and testing. Now, all these unique solutions had to be harmonised with the off-the-shelf platform by stripping them of everything confidential and case-specific, and making them general and versatile enough to be a standard building block in a customisable construction kit.

Exploring the scientific backgrounds

Several areas of optimisation and artificial intelligence had to be relied on to solve the information handling problems targeted in the project, and considerable efforts were spent on selecting and harmonising solution components from all these areas into a practicable solution that holds its ground in a live logistics environment. In the first months of the project, the modelling and global optimisation of logistics networks—primarily the so-called hub-and-spoke type characterising the project’s key application scenario—were surveyed. It was revealed that the required methods are highly specialised for each case, conveying two implications:

- the solution framework should be able to host a wide spectrum of algorithms in order to guarantee the targeted versatility, and

- researchers, developers and users should be made aware of the high degree of required specialisation to prevent pitfalls at less experienced companies.

The perceived degree of specialisation proved to be a valuable guide for our focal research efforts which were, in the first year, directed towards the following areas:

- pre-processing of large data sets for analysis,
- use of incomplete data appearing in collaborative networks,
- machine learning and generic optimisation algorithms.

Research in these areas went on throughout most of the project life span.

Getting involved and being aware—interaction with industry, science and the public

The outstanding commitment of the main industrial consortium member, Palletways, gave the consortium the opportunity of gaining a realistic picture of the application area, revealing aspects of importance and keeping research and development in the project always close to the reality of everyday industrial operation. The constant collaboration involved users from all targeted points of deployment—this ensured a co-evolution of views and revealed domain-specific insight that was not explicitly recognised even by the experts and users involved. This additional experience, mostly gained within the last project year of pilot deployment, will be of advantage while approaching the targeted users, once the solution framework finds interest with new user groups in the logistics industry.

During its three-year life span, the project was promoted to the technical and scientific community, as well as the general public, in various ways and on numerous occasions. Scientific publications, among them refereed journal articles, have been presented, published, or are in preparation to appear beyond project closure. Coverage of the project by news reports informed the general public as well as the industrial community about the project's goals and advances. In numbers, 8 press releases, 16 conference papers, 5 journal publications, one book chapter, over a dozen major user workshops and many other forms of public appearance, including presence at EU events and fairs, have been elaborated/held.

The road ahead—what's beyond the project

At the beginning of the project, we already had specific plans of ensuring that the project output will live on in the years to come, available to any prospective user, and kept flexible to incorporate refinement that would not yet be in sight during the project years. While we still are committed to this plan, additional opportunities have also presented themselves in collaboration with our industrial partner. The last year of pilot deployment shed light on possible improvements in aspects and processes not targeted by the original project scope. At the time of writing this report, we have already gathered the infrastructural and organisational requirements of these improvements that can serve as a guide for meeting network-wide strategic decisions. It is reassuring to know that we are ready to be there once these long-term measures are implemented and set the stage for even more advanced and comprehensive support for logistics operations.

Outcomes at a glance

Open-source environment

The core output of the project is the dataflow design and execution framework with the following main components:

- Flow modeller: an environment to develop dataflow definitions that will be run using the execution environment;
- ADVANCE engine: the runtime environment for working the dataflows;
- Specialised modules for model building and machine learning: these are available as special data processing modules to work with the data conveyed in the flows;
- The ADVANCE Live Reporter: a consistent set of user interfaces augmenting the dataflow components, assembled to serve typical user needs expected in the targeted logistics domain.

Decision support tools

Research and development of decision tools comprised a separate track that went parallel to dataflow-related activities. Aimed, in our specific case, at decision tasks of subcontracted depots, the decision tools serve 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.

Additional outcomes

In addition to the original goals, the scope of the ADVANCE project was extended after the first project year by the elaboration of simulation, planning and scheduling tools to improve efficiency—and, ultimately, throughput capacity—of hub and depot-to-hub operations. To this end, the following software components were developed:

- A simulator for hub (especially shed-level) operations that can be used for preliminary verification, as well as quick estimation of effects of a certain decision;
- A high-level planner for depot-to-hub operations (in particular, determination of optimal truck load composition and truck departure);
- A low-level scheduler for hub operations that can also be used as a fast estimation tool for best and worst cases expected in a given situation.

In addition to the actual software components, a important part of the outcomes was a map of possible or essential transparency improvements. These are likely to play a key role in roadmaps for future IT investments and process layout changes of the coming years.

Further Information

Official project web site: **http://www.advance-logistics.eu/**

Contact: project coordinator Dr. Elisabeth Ilie-Zudor, ilie@sztaki.hu

Consortium:

Participant organisation name	Part. short name	Country	Web-site
Computer and Automation Research Institute, Hungarian Academy of Sciences	SZTAKI	Hungary	www.sztaki.hu
Aston University	Aston	United Kingdom	www.aston.ac.uk
University of Groningen	RUG	The Netherlands	www.rug.nl
Technology Transfer System srl	TTS	Italy	www.ttsnetwork.com
Palletways UK	Palletways	United Kingdom	www.palletways.com