

MAGENTA Technology

Multi-Agent Logistics i-Scheduler for Road Transportation

Jonathan Himoff
MAGENTA Technology
130 Shaftesbury Avenue
London W1D 5EU, UK
Tel.: +44 (0) 20 7149 3750
jon.himoff@magenta-
technology.com

George Rzevski
MAGENTA Technology
130 Shaftesbury Avenue
London W1D 5EU, UK
Tel.: +44 (0) 20 7149 3750
george.rzevski@magenta-
technology.com

Petr Skobelev
MAGENTA Development
1a Osipenko St.
Samara, 443010, Russia
Tel.: +7 846 270 66 85
skobelev@magenta-
technology.ru

ABSTRACT

The paper describes main features of an Intelligent Scheduler for Road Transportation Applications based on Magenta agent technology and characterized by a number of unique, advanced features such as real-time, incremental, continuous scheduling and schedule improvement, scalability from small to very large enterprise networks, multi-criteria schedule analysis and rich decision support for the user, pro-active agent negotiations.

The scheduler is event driven, which means that it constructs schedules taking into account the occurrence of all schedule affecting events (such as the arrival of a new order, order cancellation, truck failure, road unavailability and driver no-show). As an event occurs the scheduler rapidly re-schedules only the affected part of the schedule thus causing the minimum disruption to already allocated resources.

The Scheduler has been tested on real data supplied by several UK Road Logistics operators.

The Scheduler was developed using protected Magenta technology for agent negotiations in virtual environment and knowledge discovery.

Categories and Subject Descriptors

H.4.2 [Information Systems Applications]: Type of Systems - decision support, logistics

General Terms

Management, Measurement, Design, Economics, Performance

Keywords

Intelligent scheduler, real-time planning, road transportation, multi-agent systems, virtual market, microeconomics, cargo to resource allocation, own fleet utilization improvement.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AAMAS'06, May 8-12, 2006, Hakodate, Hokkaido, Japan.

Copyright 2006 ACM 1-59593-303-4/06/0005...\$5.00.

1. INTRODUCTION

Road transportation logistics is amongst the most complex business problems. The complexity is caused by the exceedingly high variety of possible solutions (large solution space), which rules out traditional combinatorial search algorithms, and uncertainty due to high dynamics and volatility of the operational environment and openness of business networks, which makes optimisation impractical – a single optimisation run is typically an order of magnitude longer than a typical interval between two consecutive changes in operational conditions.

Resource allocation is an ongoing continuous decision making process in real time where criteria are changing “on the fly”. Therefore the two key capabilities currently required from schedulers are supporting complex business networks and planning in continuous mode.

An effective road transportation scheduler must handle transportation instructions (TI) from many different loading points to many different destinations (e.g. customer locations and cross docks where cargoes are offloaded and consolidated) and many different routes by which orders can be delivered. Choosing the best route based on consolidation or other criteria is referred to as dynamic routing. The scheduler must also be able to allocate cargos of many different sizes and weights to many different types of trucks and trailers; take into account preferences of owners, operators and drivers and fit the schedule into numerous constraints imposed by warehouse working hours, driver work rules, safety regulations and enterprise policies, eg, on choosing between own fleet and third-party carriers. Different companies have different critical constraints, e.g. permission to override time or other constraints to achieve a more efficient schedule. The schedule created must be not only feasible but also efficient, i.e. possibilities for backhauls and consolidations should be found.

Complexity is also defined by the number and variety of orders (and other events that affect scheduling) per day and the number and variety of transportation resources such as trucks.

In addition, the scheduler is expected to rapidly reschedule orders and transportation resources affected by unexpected events such as: the arrival of new orders, cancellations, failures, bad weather conditions, road works and no-show of drivers or loading crews.




Fig.1.i-Scheduler Architecture

To enable enterprises to plan and re-plan continuously, reacting to events in real-time, schedulers must support “Planning / Commit / eXecute” (PCX) stages and plan across a multi-day planning horizon.

In the planning stage orders are assigned to a truck and its journey is constructed. During this stage orders can be added or removed and the route planned for the truck can be changed as a result of subsequent events

At some point the scheduler must commit the truck. This will trigger communications to warehouses, driver shift planners, truck servicing etc to make ready the truck for its journey. During this phase changes to the truck schedule are undesirable because there would be knock on effects for the warehouse, driver assignment, etc.

The execute stage starts with the driver performing his pre-journey checks and continues until his debriefing at the end of his shift is completed. During this phase a high level of sophistication would be needed to alter the truck schedule in transit.

Magenta i-Scheduler supports PCX by allowing commit of individual trucks from a rolling schedule.

To achieve competitive advantage schedulers must take into account real-time economy where decisions are based not on some average rule but on the detailed analysis of the current situation. For example, a truck loaded by only 10% with a special cargo may be very profitable whilst a rule-based scheduler would not allow a nearly empty truck to start a journey. The requirement is to assess the economy of each truck, each journey, etc., which implies using the activity-based cost model. The current

generation of batch schedulers cannot satisfy these requirements; a fundamentally new approach to the task of allocating resources in real-time is therefore needed.

Designing a scheduler that can cope with such a variety of operating conditions, handle uncertainty related to the occurrence of events and at the same time continuously produce schedules that maximise the specified value (or minimise transportation costs) is a real intellectual challenge. To the best of our knowledge such schedulers have not been proposed in the literature or implemented in practice.

2. SOLUTION

A complex problem can be solved effectively only with a complex tool. The tool must be capable of providing the “requisite variety” of behaviours, as postulated many years ago by Ashby [1]. Only the tools based on distributed (rather than centralised) decision making fully satisfy this criterion.

The scheduling tool described in this paper handles complexity by balancing and resolving conflicts of interests of many active players rather than by following given rules and satisfying specified constraints.

This is achieved primarily by assigning an autonomous agent to every player in the transportation system and tasking agents to obtain the best possible deals for their clients. Players here include the transportation enterprise as a whole (represented by the Enterprise Agent) and all individual transportation demands and resources (including drivers and resource owners). Demands




Fig.2. i-Scheduler Ontology Architecture

and resources that have common interests self-organise into groups represented by a single agent. Agents may decide to compete or co-operate depending on prevailing circumstances.

Human operators are provided with facilities for monitoring the scheduling process and overruling the decisions made by agents. Magenta i-Scheduler supports users in their interaction with the system by providing them with options and evaluating consequences of their decisions.

Conceptual knowledge on scheduling is separated from the resource allocation mechanism, which greatly simplifies updates and increases the reuse of code.

The process of formalising domain knowledge helps in refining it and closing the gaps left due to the empirical nature of knowledge collection.

To construct a schedule Magenta i-Scheduler uses a formal description of business domain knowledge (Ontology), formal description of particular situations (a Scene), formal description of agent goals (e.g. to increase profit, maximise service level, decrease risk, etc) and real-world events. Formal descriptions can be stored, changed and examined by decision maker at any time.

Agents working in the system are goal-driven. They do not perform combinatorial search and instead negotiate to achieve the goal, which allows overcoming the problem of the very large solution space. Each agent searches only for the solution which is the best for its client, for example, to occupy a particular time slot. If this option is already taken by another agent, the conflict of interest is identified and resolved by negotiation aimed at

maximising the specified enterprise value or minimising transportation costs.

Agents are knowledge-driven. They use a Scene for analysis of the situation in the business network represented by the Scene with a view to improving quality and performance of the scheduling process. Scenes can be constructed and analysed dynamically by agents or human operators, forming the basis for Continuous Planning and Execution.

Scenes are the basis for a sophisticated interaction with users. The user is informed about the current situation by the current Scene and is given support for decision making. The fact that a Scene is a formalised description of a business situation helps to recognise patterns and identify gaps in the user knowledge and understanding.

Finally, the Road Transportation i-Scheduler can work in conjunction with Magenta knowledge discovery tools, with a view to discovering useful patterns of behaviour by means of clustering (finding groups of records with similar properties) which amounts to learning from experience.

3. ARCHITECTURE

The Road Transportation i-Scheduler is based on the standard Magenta software architecture as depicted in Figure 1 and described in [2] and [3]. The key features of Magenta technology are protected [4] and [5]. The earlier version was reported in [6].

The software is powered by a Multi-Agent Engine which contains all runtime tools necessary for the agent operation. Formalised knowledge is in Ontology supported by a powerful ontology

editor. Agents negotiate scheduling decisions in the Virtual Market.

The Logistics i-Scheduler is implemented on a J2EE platform.

4. ONTOLOGY

Ontology is constructed as a network whose nodes are objects (order, transportation instruction, cargo, truck, etc) and links are relations (Journey consist of TI, etc). It contains also attributes of objects like truck capacity, etc.

As depicted in Figure 2, Ontology is divided into three main parts. Basic part of ontology includes main logistics concepts. Schedule Ontology consists of schedule owners and schedules and contains references to the operations which constitute schedules, described in operation ontology. This structure allows the introduction of any type of operations into schedules.

A typical Scene is shown in Figure 3. The scene representation is design to help the user to understand the network situation. Looking at “Client” it is possible to find out where all TIs are at any given moment. Looking at “Location” it is possible to see all Journeys that are planned to go through this location.

Scenes also make decision making logic more visual and understandable to the end-users.




Fig.3. An Example of Scene

5. VIRTUAL MARKET

Virtual Market is a virtual environment in which agents negotiate the allocation of transportation resources to orders.

5.1 Agent Types

The Logistics i-Scheduler has several key types of Agents, including:

- Customer Agent – looks after the required Customer Service Level
- Order Agent – splits orders into transportation instructions and monitors its Key Performance Indicators (KPIs). If the allocation of a Transportation Instruction (TI) is unsatisfactory the Agent will try to improve it by searching for the most appropriate journey
- Transportation Instruction Agent – searches for the best journeys
- Journey Agent – looks for the best TIs, devises good consolidation, looks for the best routes with minimal mileage
- Cross-dock Agent – supports the schedule for cross-docking and allocates Journeys to time slots
- Group Agent – keeps similar TIs together
- Enterprise Agent – minimises fixed costs by decreasing number of trucks used; switches company strategies; monitors schedules and causes disturbances to destroy ‘bad’ journeys

Each type can act as a Demand Agent or a Resource Agent depending on the situation.

5.2 Generic Virtual Market Architecture (Agent Communities)

Magenta i-Scheduler has a generic virtual market architecture, which enables agents to form natural coalitions, e.g. TIs on a Journey or Journeys which together constitute a Driver Shift (all journeys made by a driver during his working day). This means that virtual market is organised in many levels, with nested groups formed by common interests where each group has an agent representative (Group Agent) who monitors what is happening in the group. Agents can join the group and leave it. The whole group can move (e.g. Journey goes to another Driver Shift).

This architecture provides flexibility in developing communities of different types of agents which can work as coalitions. Such approach also gives more options to find a right balance between interests of individual egoistic agents (TI, journey, etc) and interests of communities of agents (fleet, journey, etc). Also, this balance can be changed dynamically in reaction to events (X-Dock under terrorist attack). Finally, this architecture forms the basis for ongoing self-improvement of schedules.

The generic architecture provides stable, well-designed and elegant basis for rapid development of new applications in a wide range of new domains (supply chain, factories, containers, projects, finance, nurses, conference boxes, etc).




Fig.4. Multi-Level Virtual Market Architecture

5.3 Agent Negotiation Strategies and Protocols

Every agent is given the task of obtaining the best possible match for its client and a right to drop a previously agreed match if a better opportunity presents itself as events continue occurring. The abandoned partner pro-actively seeks a new match.

Agents on the Virtual Market use enhanced Contract Net negotiations protocol. Within Contract Net protocol agents pass the following iterations:

- Ask for pre-matching (rough) estimation
- Ask for plan, if possible to allocate without drop
- Ask for complete plan (drop is possible)
- Accept proposal
- Receive notification that plan is applied/failed

At each stage agents can cut or prioritize negotiation branches, thus detecting uninteresting options with fewer calculations.

If a plan fails, the agent can re-send a corrected proposal.

If an agent has two equal options, one of which involves drop and another is without drop – it will prefer the option without drop, as it saves time.

The protocol is decomposed into two parts: the contractor role and coordinator role. The contractor role is responsible for (and fully encapsulates) message exchanges between a particular pair of agents representing demand and resource.

Coordinator role contains decision-making logic of an agent at the protocol level and allows the demand agent to interact with several resources simultaneously (demand agent's coordinator role) and resource agent – with several demands (resource agent's coordinator role).

Magenta agents have multiple objectives, or criteria (quality, price, time, risk) each with a weight coefficient. This can be changed by the user for each individual agent. At any moment of

time the user can increase the relative importance of the service level or maximise profit criterion. If for example orders from VIP customers are allocated and a new order comes from an even more important customer with tighter conditions, this may result in decreasing the service level for already allocated orders within the allowable interval. Changes in weight coefficients represent the change of agent strategy of working in the virtual market.

5.4 Microeconomics

The aim of microeconomics in Magenta multi-agent systems is to ensure that at any moment of time each agent knows which journeys are profitable for the company and which are not. There is a possibility to take into account not only "physical costs" but also the cost of planning. For this, the following principles are introduced. Each order finds out the cost of its transportation using price list and finds the most profitable journey. The profit received from transportation is shared between TI, Journey and Truck Agents. The microeconomics includes a system of taxation aimed to further improve scheduling efficiency. It does this by applying a penalty (tax) for making changes to the schedule close to Commit time. This ensures that only changes giving significant benefits are made as planning time runs out.

6. HOW THE SCHEDULING WORKS

The Initial Scene, which is a model of the transportation system at the beginning of scheduling, is constructed by an operator using a user-friendly editor or imported from an XML file. The Scene depicts all transportation resources (trucks, trailers, depots, warehouses, loading points, delivery points) and their locations on the transportation network. Goals, criteria, objectives (KPI), strategies and preferences are specified.

As the first order arrives, agents split the order into transportation instructions and allocate transportation resources to each TI. The Initial Scene is thus modified into the Current Scene, which is being updated after every event. The Current Schedule is extracted from the Current Scene and can take the form of a table or a diagram.

New events are added to the current schedule. They can be imported from third-party systems, from XML-files or added manually. Events are assigned individual objectives and strategies. When new events are entered into the system, the system stops other activities (or delays the incoming event) and starts TI allocation.

If events are entered in a batch, prioritization is carried out. Outstanding TIs can be prioritized with the help of prioritization logic defined in ontology or manually.

Next, new TIs are planned incrementally. TI agents pre-match by making rough option estimation and checking constraints. Next, they start negotiation with Journeys telling them the amount they are ready to pay for allocation. If TI is interesting for the Journey, it tries to allocate this TI using Open Slot, Shift or Drop algorithms. If an option is found, Journey makes the final option estimation and checks constraints. Journey sends message to TI to confirm allocation and informs other TIs belonging to it that new TI was allocated and they can re-calculate their values and costs.

If there are no new events, unsatisfied TIs, Journeys, Driver Shifts and Trucks could start pro-active improvement of their allocation. An important advantage of Magenta software is that this pro-

active improvement takes place in the system perpetually – i-Scheduler continuously works towards the best possible schedule. Any type of agent can be pro-active and agents who are unsatisfied with their current allocation are most likely to get pro-active.

User can review results using a Gantt chart or a map and analyse reports. There is a possibility to manually adjust the schedule asking the system to suggest alternatives. The system can also highlight parts of the schedule that needs improving.

Each Journey has its individual commit time. As commit time approaches, the agents get more and more active to utilise the last chance to find a better option. After the commit time of a Journey (e.g. 3 hours before the driver sets off) no further changes are made to this Journey schedule – Journey is passed to commit stage. In future versions, in order to improve versatility of the system, it will be possible to carry out additional scheduling at commit stage.

The incremental nature of Magenta scheduling and the use of different planning strategies means that the i-Scheduler will always produce a plan for an order, even if it is entered into the system very shortly before its commit time.

7. USER INTERFACES

The main window of the Logistics i-Scheduler, as shown in Figure 5, allows importing orders and placing them in the queue, starting planning and reviewing the schedule using a Gantt chart (Figure 6) and map, including detailed information about journeys and KPIs. Different parts of the window can be easily configured and re-arranged.




Fig. 5: Desktop Interface with Easily Configured Dashboard




Fig. 6: Schedule Gantt Chart

At any moment user can view agent inspector (Figure 7) showing which agents are currently working in the system and analyse a negotiations log (Figure 8).




Fig. 7: Agent Inspector




Fig. 8: Agent Negotiation Log

8. TESTS WITH CLIENT DATA

The first version of the Road Transportation i-Scheduler was tested on two sets of real-time data obtained from Client A, a 3rd Party Logistics Provider, and Client B, a provider of logistics and freight management services.

8.1 Client A Test Results

Client A requirements were to create transportation schedules in real time for 200 transportation instructions and 50 trucks (36 own fleet trucks plus 15 third-party carriers) operating on the UK Business Network. The network included 9 Distribution Centres/factories, cross-dock points for primary/secondary moves consolidation, 3 truck bases doing shared operations. Additional requirements were handling transportation instruction availability windows, backhaul, consolidation, vehicle capacity availability windows and constraint stressing.

The Client A currently employs two operators who require four to five hours to produce the schedule. Because of the current slow scheduling process, the Client plans two days in advance of execution. The Client also experimented with conventional batch schedulers and found that they could schedule only 70% of journeys correctly. The 30% of journeys had to be re-scheduled manually.

The Magenta i-Scheduler has completed scheduling of 200 transportation instructions in 8 minutes, planning 116 journeys of a total of 20790 miles. The result exceeded expectations with 95% of journeys scheduled correctly without any manual rework. The remaining 5% could not be scheduled during the tests because of incomplete information. Using Magenta scheduler Client A could complete planning and execution in the same day.

8.2 Client B Test Results

Client B requirements were to create transportation schedules for 4000 transportation instructions and 200 trucks operating on the UK business network. The network included primary and secondary deliveries between about 600 locations, 3 cross docks, 4 secure trailer swap locations and other types of locations. The network was also characterized by considerable number of very small orders. Special requirements included dynamic routing, cross-docking, handling location availability windows and driver breaks.

It took Magenta i-Scheduler about 4 hours to build a schedule for 4000 orders and 200 trucks with dynamical routing through 3 cross docks. This schedule showed an effective consolidation of small orders onto trucks. It was capable of incrementally planning new orders in near real time (a few seconds was required for each new order). As far as we know, this has not been achieved by any other transportation scheduling system. Before testing Magenta software the Client has spent considerable resources searching for a scheduler capable of handling their requirements and found none.

8.3 Performance Tests

During the tests on client data Magenta conducted an investigation into the effect of the number of orders on scheduler performance. As Fig. 9 shows, in situations in which a new order arrived when previous orders have been allocated, the time

required to complete the schedule is, approximately, linearly dependent on the number of orders.




Fig. 9: The Effect of the Number of Previously Planned Orders on Scheduling Time for a New Order

This result is highly significant. It shows that Magenta schedulers can be easily scaled up to handle any size of order allocation.

A close inspection of the diagram shows that although the average increase in time is linearly dependent on the number of orders there exist significant number of low points and picks. Low points signify a situation where the new order was allocated with a minimum of re-scheduling, and picks are present where the arrival of a new order required a considerable re-allocation of previously assigned resources.

A very strong advantage of Magenta i-Scheduler is that planning is very fast and therefore can be done in real time. The rapid scheduling is achieved by the use of the incremental scheduling process, as opposed to the batch mode. Incremental planning gives an additional benefit of re-scheduling only those parts of the schedule that are affected by the occurrence of a new event, without having to completely destroy the schedule and re-build it "from scratch".

9. CONCLUSIONS

The test results described in Section 8 of this paper show that Magenta multi-agent technology, in general, and Magenta Road Transportation i-Scheduler, in particular, are capable of significantly improving scheduling effectiveness in comparison with conventional rule-based centralised schedulers. The difference is particularly visible in cases where there is a need to match a very large number and wide variety of orders and resources in real time.

The key features that contributed to the effectiveness of scheduling are summarised below:

- The distributed nature of decision making, which enables the swarm of agents to achieve a common goal whilst each agent member of the swarm is allowed to pursue the achievement of its own goal and yet be ready to negotiate when a conflict of interests is discovered.
- The ability to incrementally react to the occurrence of a new event, such as the arrival of a new order, by re-scheduling only affected parts of the previously agreed schedule.
- Pro-activity of agents, which drives them to improve the previously agreed allocations whenever practical.

10. REFERENCES

- [1] Ashby, W.R. An Introduction to Cybernetics. Chapman & Hall, 1956.
- [2] Rzevski, G., Skobelev, P. Magenta Technology: Software Platform v2.1 – Magenta Technology Whitepaper, July 2005
- [3] Wooldridge, M. Responding to Real-World Complexity: Introduction to Multi-Agent Technology – Magenta Technology Whitepaper, August 2004.
- [4] Agent Method and Computer System For Negotiating in a Virtual Environment, Patent Reference No: WO 03/067432 A1. Published 14.08.2003
- [5] Data Mining. Patent Reference No: GB 2 411 015 A. Published 17.08.2005
- [6] Himoff, J., Skobelev, P., Wooldridge, M. MAGENTA Technology: Multi-Agent Systems for Industrial Logistics – Magenta Technology, August 2004.