Development of the rule based approach to traffic management by the Dutch road authorities

Silvie Spreeuwenberg, LibRT, www.librt.com,
Rolf Krikke, Quovadits, www.quovadits.com

Abstract. Rule based traffic management is a methodology for dynamic traffic management that is a joint development of the different road authorities in the Netherlands. The approach tackles operational issues by disentangling problem detection, problem solution and conflict handling for each element in the road network. Each element in the road network follows the same generic rules and defers the traffic problem to other roads in the network that have less priority. This way traffic is distributed across the road network and congestion is prevented.

Keywords: traffic management, business rules, complex event processing, multi agent systems, decision modelling notation, intelligent agents, network propagation.

1 Introduction

Rule based traffic management (RBTM) is a methodology for dynamic traffic management that is a joint development of the different road authorities in the Nether-lands. The methodology supports collaboration between various traffic management authorities (e.g., counties, provinces, or cities) in diagnosing and improving traffic conditions by offering a common vocabulary and business rules. The approach may be used in similar domains (for example, crowd control). The proposed method to standardize a domain's compliance process may in principle be used to standardize compliance processes in other domains.

1.1 Traffic management

The road networks in the Netherlands are, together with those of the Scandinavian countries, among the safest in the world. One of the factors contributing to this position is the attention paid to road safety in traffic management processes and sys-tems. Significant parts of both networks have automated queue protection. The network has a relatively high traffic density (about 20 million vehicle kilometers per motorway kilometer per year) over the whole highway network, with (much) higher densities in metropolitan areas. Inductive loop detectors to measure traffic passages, matrix traffic signals, traffic signs and dynamic route information panels are used extensively. Distances between junctions and interchanges are short and speed limits are not uniform. There are five regional traffic control centers for managing the highways, tunnels and bridges.

The thirteen provinces have traffic control centers for managing the regional roads. Some larger cities have traffic control centers to man-age the urban roads.

Response plans are used for mutual agreement, communication and control of the desired response to specific traffic events. A response plan translates policy and experience in operational instructions, often in the form of a flow chart. Response plans are intended for human operators to deal with events, incidents and road works. Response plans are typically executed by Traffic Management Systems of multiple authorities.

1.2 Challenges and objectives for road authorities and traffic operators

In the last decade response plans have also been developed for daily traffic management. The broad use of response plans has resulted in a labor-intensive maintenance and operational execution process. Many response plans are needed for different locations and situations (e.g., morning rush, evening rush, forecast event, incident, roadwork). Each response plan details: problem detection, problem solution in terms of the setting of signs and signals, handling of conflicting requests for a traffic control device, and restoring to 'normal'.

The main objective of RBTM is to develop a more efficient process to facilitate and execute daily operational traffic management, thus enhancing quality and effectiveness. Additional benefits of the standardization of traffic management rules are an easier transition to new technologies (like in-car systems) and tightening of the connection between traffic policy and operational execution (compliance).

1.3 Rule based traffic management

The rule based traffic management approach tackles these issues by disentangling problem detection, problem solution and conflict handling. Capacity problems on the road network are detected in a generic way, problem solutions are standardized and may be re-used for different problems, and finally the agreed priority of a road is used to solve conflicting solution requests. The idea is to distribute or move capacity problems on the road network to roads with a lower priority. This can be achieved by the repeating execution of simple and generic business logic by every link in the road network.

1.4 Alternative approaches

The most common practice in traffic management centers around the world is the semi-automated support of response plans. A response plan is a procedural description that could be described by a set of production rules and is often represented as a flow chart. Each response plan deals with a specific situation like an event, rush hour or an incident at a specific location. Some decisions or actions in the response plan are automated and others need manual intervention. The complexity for the traffic operator is dealing with situations when multiple response plans apply and need to be combined 'in real time'.

Many innovations in traffic management deal with innovative road side equipment that work on a local level like an intersection or a combination of ramp metering and traffic lights (Hoogendoorn, Van Kooten, & Adams). Neural network technology is used to optimize the cycle times of traffic lights resulting in new signal plans for different kinds of situations (Saraf, 1994).

The difference between those innovations and this approach is that our approach deals with a large road network including highways, regional roads and urban roads of different regions in a country.

1.5 Solution outline

This paper is organized as follows. Section 2 will describe important concepts and terminology of the approach. Section 3 will describe the generic logic of the approach for decisions, process and definitions. Section 4 will describe the evaluation results of the approach. Section 5 will describe the way the methodology has been developed, is disseminated, is adopted and is expected to evolve. Section 6 will relate the approach to AI research, the principles of business rules management and other domains.

2 Building blocks of rule based traffic management

The foundation of rule based traffic management (RBTM) is policy, topology and services. We call these the 'building blocks of RBTM' and we explain each below.

2.1 Standardized policies define our objectives

The road authorities for an area have agreed on a joint vision about traffic management resulting in:

- A managed roads network: the selection of roads that are available for traffic management.
- The preferred routes: the preferred and diversion routes for the most important traffic flows based on analyses of origin-destination flows.
- A road priority map: the priority of a route is based on routes usage and indicates which roads should have good traffic flow conditions, eventually at the expense of other (lower priority) roads in the network – see an example in Fig. 1.
- The traffic management norms: qualitative or quantitative definitions of threshold values for bottlenecks.

2.2 Standardized services define what we do.

RBTM defines a traffic management service as a setting for a traffic control device (eventually a combination of traffic control devices) that is conditionally available for traffic management. There are three different kinds of services:

- Increase outbound flow – a service that controls traffic capacity at a flow control point, thus enhancing the amount of traffic that can exit a link in a given direction.

— Decrease inbound flow – a service that controls traffic capacity at a flow control point, thus restricting the amount of traffic that can enter the link from a specific upstream link.
— Reroute – a service that directs traffic from a destination to a diversion route at a decision point.

The condition under which a service is NOT available is named the restriction of the service.

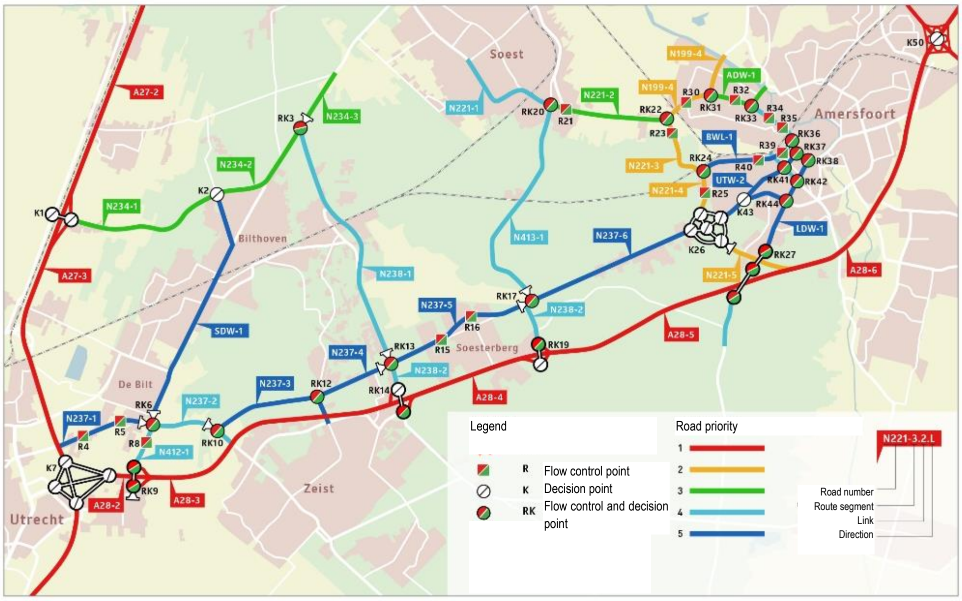

Fig. 1. Example road priority map and managed roads network.

2.3 Standardized topology elements define what we know

The managed roads network (see Fig. 1) is divided by:

— Flow control points: locations on the road network where traffic capacity may be influenced; for example, by ramp metering or a traffic light system.
— Decision points: locations on the road network where traffic may choose a diversion route to a destination. These are the points on crossings or intersections in the managed roads network.
— Route segments: trajectories on the managed roads network between two decision points.
— Links: trajectories on the managed roads network between two flow control points for each driving direction. A downstream link is any link that follows the driving

direction and shares the same flow control- or decision point. Similarly an upstream link is situated in the opposite driving direction.

Data collected on route segments and links are used to detect one of the three problem phases: saturation, congestion and gridlock. Traffic control devices on or near junctions are used to adjust traffic conditions and traffic flow. The rules will optimize traffic conditions per link, thereby propagating traffic issues to upstream links and moving traffic to roads with a lower priority in the network.

3 Mechanism of rule based traffic management

The rule based traffic management (RBTM) has four principles that determine which service must be requested.

1. Prevent saturation on a link by early detection of bottlenecks and by the request of services to control outbound and inbound traffic.
2. Optimize travel time on route segments by requesting the reroute services of upstream decision points.
3. Turn down a service request when traffic conditions violate the policy constraints set by the traffic management authority.
4. Manage conflicting service requests by turning down the service requested from the least severe traffic situation.

Each principle translates to a set of generic business rules that answer questions like:

- Which traffic management service must be requested for a control point?
- Is a traffic management service available?
- Which traffic management service must be executed when conflicting services are requested?

The business rules are based on the thresholds set by the traffic management norm on a link or route segment and the priority of roads, taking practical considerations, like the availability of traffic data, into account.

3.1 Decision logic

The first and second principles are described by generic business rules that state which kind of service must be requested for a given problem phase of a traffic situation. Table 1 shows this logic presented as a decision table.

The decision tables follow the semantics defined in the standard ‘Decision Model Notation’ (OMG , 2014) as ‘rules as cross-tab’ or ‘rules as columns’. The output values are distinguished from the input values by presenting them in a cell with a white background.
The tables shown are contracted tables. A dash symbol (‘-‘) in an input value cell is used to mean any input value, i.e., the input is irrelevant for the containing rule. A dash

symbol ('-') in an output value cell is used to mean no output value, i.e., the output does not apply for the containing rule. The default DMN semantics is used meaning that the table returns the output of one rule only. It should not contain overlapping rules.

Which kind of service to request?			
Problem phase:	Saturation	Congestion	Gridlock
Promote outbound?	Request	Request	Request
Limit inbound?	-	Request	Request
Reroute?	-	-	Request

Table 1.

This table is a 'rules as columns' table. Each column in this table represents one rule. The condition is based on the problem phase.

For each link and route segment the operational traffic engineer must create business rules that define the problem phase of a traffic situation. Table 2 shows a generic version of this logic.

What is the problem phase?			
Topology:	Link	Link	Route segment
Waiting queue:	>90% of sorting lane	>90% of link	-
Travel time:	-	-	> 90% of travel time norm
Saturation?	Yes	-	-
Congestion?	No	Yes	-
Gridlock?	No	No	Yes

Table 2.

This table is a 'rules as columns' table. Each column in this table represents one rule. The condition is based on the topology and traffic data (in this case waiting queue length or travel time). Preferably these business rules are generically expressed in terms of a deviation by the traffic management norm that has been defined in the policy. However, there may be a need to define rules for a specific link or route segment due to local differences in the availability of traffic data. An example is given in table 5.

The third principle is described by generic business rules that state which kind of service is available given the network capacity.

Is a service available?			
	Promote outbound	**Limit inbound**	**Reroute**
No capacity on conflicting direction.	Not available	-	-
No capacity on upstream link.	-	Not available	-
No capacity on downstream link.	Not available	-	
No capacity on diversion route.	-	-	Not available

Table 3.

This table is a 'rules as crosstab' table. For each link and route segment the operational traffic engineer must create specific business rules that define the network capacity.

The fourth principle is described by generic business rules that state when services are conflicting.

Are two services conflicting?			
	Promote outbound	Limit inbound	Reroute
Promote outbound	Instrument conflict	-	-
Limit inbound	Capacity conflict	No	-
Reroute	Capacity conflict	Service conflict	Instrument conflict

Table 4.

This table is a 'rules as crosstab' table. It derives the kind of conflict detected based on the kinds of services requested for the same control point or using the same link.

An instrument conflict means that the services conflict because the instrument is not able to execute both service requests at the same time. For example a traffic light may only promote outbound traffic in one direction and a dynamic route information panel (DRIP) may only present one reroute message (this is a Dutch policy).

Instrument and service conflicts are resolved by selecting the service request from the link with highest road priority or (in case of equal priorities) the service that is requested first.

Capacity conflicts are not resolved. RBTM will act on those and request more and heavier services on upstream roads.

3.2 Process logic

RBTM has process logic that defines the order in which the business rules and traffic situations are evaluated. Typically this logic is executed by a Traffic Management System[1]. The logic is executed in 2 - 5 minute cycles (the control cycle time may be tuned and depends on the average link length in the network) and establishes:

1. The service requests based on the problem phase of each link and route segment.
2. The availability of the services for each flow control or decision point.
3. The most severe traffic situation for conflicting service requests.

When all business rules are evaluated the service requests will be executed and the process restarts[2], see figure 2.

Services may also be started manually by a traffic operator based on a message or point in time. In that case the traffic operator is also responsible to terminate a service. The service availability conditions are still to be respected.

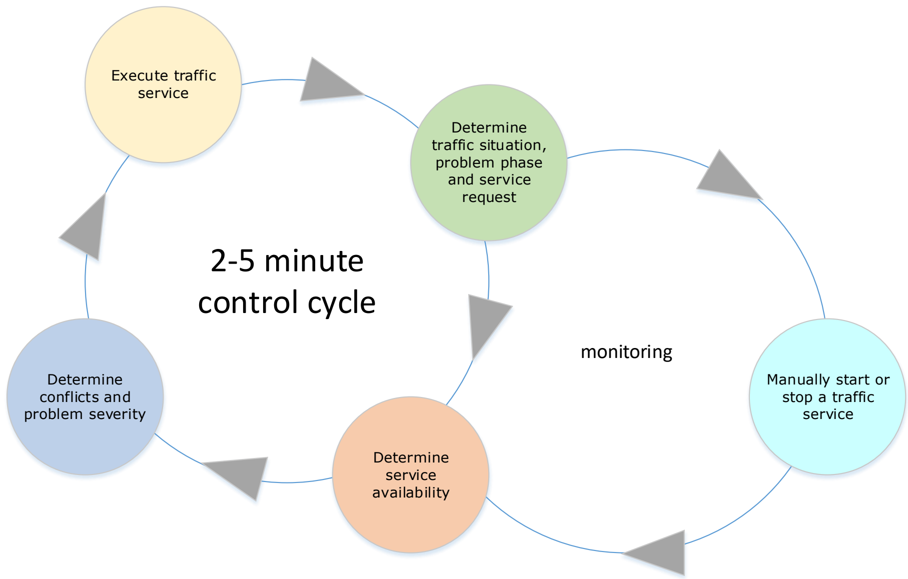

Fig. 2. Visualization of the process logic.

---

1 There are different vendors that offer Advanced Traffic Management Systems. The exact way this process needs to be configured in the system may differ between vendors and extends beyond the topic of this paper.

2 The control time prevents too many service requests due to traffic hysteresis.

3.3 Definition logic

The definition logic is the business rules that the operational traffic engineer creates to tune the system to local variations and feed the generic business rules with information.

For each segment in the managed roads network the definition logic defines the problem phase (saturation, spill back and gridlock) and network capacity in terms of measurable traffic data, preferably presented as a decision table. The example decision in Table 5 combines the definition logic with the generic decision logic of a link and restrictions of the requested services. This helps the traffic engineer to understand how the different kinds of logic interact.

For the convenience of the traffic engineer the formatting of this table is slightly different. The green colour of the output cells with the word ON is to highlight that a service is started. The orange colour of the output cells is to highlight that a service is stopped. The white colour of the input cells is to indicate that the traffic engineer can adjust the parameters.

Service requests by link: A10R S102 > A10R S103?			
Problem phase:	Saturation	Congestion	Gridlock
Speed:	<75 km./h	<65 km./h	<55 km./h
Promote outbound?	ON	ON	ON
If capacity upstream link < 60% of maximum capacity upstream link			
Limit inbound?	OFF	ON	ON
If waiting time TDI < maximum waiting line			
Reroute?	OFF	OFF	ON
If Travel time diversion route < Travel time preferred route			

Table 5.

4 Evaluation of the rule based traffic management

The approach is evaluated in two pilot projects. The first project evaluated the results of the approach for improved traffic flow on the highway ring-road of Amsterdam. The results showed that traffic flow on the highway improved at the expense of traffic conditions on the urban roads. A second pilot evaluated how the approach could distribute the traffic on the network by extending the approach to more services on the provincial and urban roads. We will describe in this section the results of the second pilot, locally known as Praktijk Proef Amsterdam Noord (PPA-N) and based on the evaluation report (Krikke, 2017).

4.1 Description of road network and evaluation objectives

The evaluation was performed on a network of roads in the North of Amsterdam, consisting of the westbound traffic on N516, starting with the intersection N203 (provincial road), until its junction with the A8 (highway). See Fig. 2.

The network has seven traffic lights, one ramp meter and one bridge. This network was chosen because there is an overloaded intersection that connects with the urban roads to a satellite city (Zaandam) causing many traffic jams. The situation is very typical for other locations in the Netherlands.

The traffic light systems were configured to support the services: ‘increase outbound flow’ and ‘decrease inbound flow’. The length of the waiting queues for traffic lights were measured by radar. Reroute services were not used since they were already evaluated in the first pilot and there were no signs available in this area. A conflict handling strategy based on the road priority was programmed in the traffic light system to handle conflicting service requests.

Fig.2. PPA-noord evaluation network.

4.2 Evaluation method

The new approach was active for 40 days. The results from this time period have been compared to the traffic situation on days where the traffic lights have the ‘typically most optimal’ program. The following metrics have been used:

- Measure delayed traffic conditions by calculating the expected number of vehicle passages based on road capacity and the actual number of vehicle passages; we call this number vehicle loss hours.
- Measure the distribution of the traffic on the network by calculating the length of waiting queues for each traffic light.
- Measure road user satisfaction by following a group of 23 road users using weekly surveys via internet.
- Measure safety conditions by counting the number of red light runners.

– Furthermore the effect, vulnerable road users (cyclist) and the time table of public transportation is measured.

4.3 Results

The Netherlands is a country with a high population density and a road network that is used to its full capacity with little options to change the infrastructure. Therefore spectacular results in the improvement of the traffic situation are very hard to achieve, causing small but significant improvements to be recognized. The results show that the number of vehicle loss hours is slightly improved (4%) during rush hours. More importantly, the traffic problem is distributed to lower priority roads in the network by showing increased vehicle loss hours and longer waiting queues on these roads, resulting in an improved compliance with the policy.

These results are confirmed by showing an expected decrease in rear-end collisions but there is also an increase in red light runners. This is an undesired but natural result of the service 'decrease inbound flow'. The average satisfaction score of the road users is also significantly better on the days where the rule approach was 'on' as shown in Figure 3. The blue line indicates whether the rule approach was 'on' (1) or 'off' (0). The red line shows the average road user satisfaction. It follows the same pattern as the blue line indicating that the road user satisfaction increases when the rule approach was 'on'.

The results of week 52 are not representative due to a holiday in this week. There is no statistically significant other explanation for the increase in the road user satisfaction than the improved traffic conditions due to the rule approach being 'on' (Arcadis).

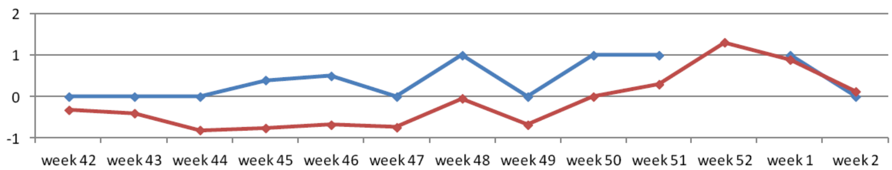

Fig. 3. Road user satisfaction analysis

The traffic engineer and traffic operator have asked to improve their common operational picture by showing the results of the decision tables (problem phase) as a color on a map. None of the domain experts have a background in IT, knowledge representation or business rules and that has not been a disadvantage in any way in developing this methodology and understanding the decision logic represented in decision tables.

5 Development of the rule based approach for traffic management

After the bill on a kilometer tax to decrease traffic on rush hours was rejected the minister of infrastructure gave the road authorities the assignment to collaborate and use the existing road network capacity in an optimal way. This included the idea of distributing traffic over the whole road network. However, each road authority was used to doing traffic management in their own way, with their own methodology and local experts. The first meetings on collaboration turned out to be a school example of misunderstanding, underestimation and complaints.

5.1 Challenges and obstacles during the development phase

The first challenge was to agree that a common methodology was feasible and desired. We organized 10 workshops with 20 representatives of different road authorities resulting in a description of the current practices and the desired practice in road management. It was far from being a complete methodology and the number of concepts having multiple synonyms was too large.

However the first steps were taken and the following year a second set of workshops with five representatives of different road authorities was organized, resulting in a description of a common vocabulary for network based traffic management. The method was accepted as the standard approach by the national traffic council with representatives of different road authorities.

The evaluation of the approach in the PPA pilot was an important next step in acceptance of the new methodology by road authorities and traffic engineers. The lessons learned in the pilot have been integrated in an update of the methodology.

5.2 Dissemination of the approach

The next challenge is to make sure the approach is well known, accepted and supported by readily available tools and expertise for all road authorities. The following artefacts are used:

- The methodology is published and available as free download (Spreeuwenberg & Krikke, 2017) in the handbook of traffic engineering. The book is well known, used by the traffic industry and integrated in the curriculum of the high school for traffic engineering.
- A one-day training has been developed for traffic managers and traffic operators. As an exercise, an existing non-compliant way of working is transformed to a compliant traffic management process based on the new methodology.
- A game is part of the training. Each trainee represents a link in the road network. Each link receives traffic situations (congestion, spill back) and must follow the standard decision logic to conclude which service is executed. This way, trainees learn and understand how the traffic issues are distributed to lower priority roads and they build trust in the operational execution of the method.

- A two-minute animation has been developed to provide a quick overview of all benefits for managers and operators (Hiroi, 2017).
- Finally we would like to develop a simulation program to test road priority changes on traffic network performance.

5.3 Further developments

The new approach separates the management of the road network (using the generic business rules) and the management of road side equipment (like traffic lights and electronic messaging signs). This paves the road to easily connect to innovations that are changing our world rapidly. Instead of sending a text to an dynamic route information panel to communicate a reroute service we could sent network control information to connected cars or service providers. So by standardizing the network management layer we have also created a way to standardize the communication about our policy with other parties. This is a requirement for adopting the technology innovation since the technology changes, not the policy.

6 Concluding remarks

The authors of this article have a background in artificial intelligence and robotics (respectively) but most importantly they have decades of practical experience in building decision support- and traffic management systems. This paper is the result of looking for something that works in an operational environment while solving practical problems. We did not look for an environment to test a theory; we had no interest in selling a solution without a problem to solve. In fact it has been the other way around: we combined the experiences of the road authorities in a consistent and practical methodology.

By writing this paper we force ourselves to compare our results with the research community and common practice. There are two interesting perspectives: the techniques and methods developed by the research of Artificial Intelligence (AI) and the guidelines that are most known under the name Business Rules Management (BRM). The authors are already looking forward to the third perspective: applicability to other domains.

6.1 The AI perspective

The rule based approach for traffic management is an implementation of a multi-agent system (MAS) (Ferber, 1999) that deploys complex behavior based on simple rules.
Each agent represents a link in the road network and has the goal to optimize the used capacity on its link.

Each agent is a model based reflex agent (Russell & Norvig, 2003). Sensor information from the link (the environment) is translated to an assessment of the traffic situation on the link. Based on that assessment and the rules in the decision tables one or more services are selected for a related point in the network. This service will send an action to an actuator. The actuator will change the environment and the agents on upstream and downstream links will react on those changes with the same strategy.

The methodology is different from earlier agent-based control mechanisms (Wang, 2005) in the domain of traffic management because they optimize multiple traffic lights on one intersection while we optimize a traffic network consisting of different kinds of roads (highways, regional roads and urban roads) using different kinds of actuators (traffic lights, ramp metering, dynamic speed, and rerouting). To our knowledge this is the first large scale implementation of such a multi agent strategy in traffic management.

6.2 The BRM perspective

The business rules community has a couple of mantras. One of them is that 'business rules need to be motivated', expressed in Article 8.1 of the Business Rules Manifesto (Business rules group, 2003). By using the quantitative norms expressed in the policy directly in our operational control mechanism, we believe this mantra is being served. When the business (being the road owners) changes the priority of the roads or norms for the roads the operation will automatically follow.

Another mantra of the business rules community is that we should 'manage the business logic, not hardware or software platforms'. Therefore the rule based approach to traffic management expresses both conditions and actions of the decision rules independent from the sensor or actuator technology. Given the fast development of new technology in this domain (internet of things [IoT], in-car systems, self-driving cars, navigation systems) this mantra is of extreme importance. Technology may change quickly but policy must still be enforced; therefore policy should be described in a technology independent way.

Business involvement is another pillar of the business rules community. RBTM has been developed together with domain experts (being traffic engineers and traffic operators) that do not have a background in AI or rule based technology. This is an example of the Mantra 'by and for the business, not IT'. The DMN notation that has been introduced as a condensed way of representing a set of rules has been very intuitive and never been the source of questions or debate.

6.3 Applicability of the approach to other domains

We believe the multi agent systems approach is applicable to similar domains that deal with flows in time; for example, crowd management during an event or passenger flows in an airport. The terminology for traffic situations and services may need to be slightly adjusted for those domains. Also the sensor data and actuator actions will be different. But the general idea of the standard decision logic for service request, service availability and service conflicts is likely to hold.

The applied methodology to standardize the process of policy compliance may also be extended to other domains. The general idea is to create a high level vocabulary in terms of events and actions. These are combined using standard decision logic. A domain expert uses this as a framework to configure domain definitions. This way the domain expert does not need to work on algorithmic logic himself and instead finds his way more easily in a system of rules. This strategy 1) avoids the need for verification requirements (Spreeuwenberg & Gerrits, Requirements for Successful Verification in Practice) and 2) solves the well-known 'knowledge acquisition bottleneck' in expert systems development described by (Cullen & Bryman, 1988).

Recent presentations of the methodology to a broad audience working in different domains taught us that the methodology is inspiring for processes in retail (logistics) and the finance industry (money flows). We are looking forward to applying this strategy to other domains and generalizing the methodology applied.

References

Arcadis. (n.d.). *Evaluatie rapport praktijkproef amsterdam.* Retrieved from http://www.praktijkproefamsterdam.nl

Business rules group. (2003). *The business rules manifesto*. Retrieved from http://www.businessrulesgroup.org/brmanifesto.htm

Cullen, J., & Bryman, A. (1988). The Knowledge Acquisition Bottleneck: Time for Reassessment? *Expert Systems*.

Ferber, J. (1999). *Multi-Agent System: An Introduction to Distributed Artificial Intelligence.* Harlow: Addison Wesley Longman.

Hiroi, K. (Director). (2017). *Animation Landelijke Regelaanpak* [Motion Picture]. Retrieved from https://youtu.be/8twSS0mJD9M

Hoogendoorn, S., Van Kooten, J., & Adams, R. (n.d.). Lessons Learned from Field Operational Test of Integrated Network Management in Amsterdam. *Transportation Research Record: Journal of the Transportation Research Board, 2554*.

Krikke, R. (2017). *Evaluatierapport PPA noord.* Retrieved from http://www.praktijkproefamsterdam.nl

OMG . (2014). *Decision Model and Notation standard* .

Russell, S., & Norvig, P. (2003). *Artificial Intelligence: A Modern Approach.* New Jersey: Prentice Hall.

Saraf, R. K. (1994). *Adaptive traffic control using neural networks.* Vanderbilt University Nashville.

Spreeuwenberg, S., & Gerrits, R. (n.d.). Requirements for Successful Verification in Practice. *Proceedings Fifteenth Flairs conference,* p 221, 2002.

Spreeuwenberg, S., & Krikke, R. (2017). *Handboek verkeersmanagement, module regelaanpak.* CROW. Retrieved from http://www.crow.nl/regelaanpak

Wang, F.-Y. (2005). Agent-based control for networked traffic management systems.