ON SELF-ORGANIZING TRANSPORT NETWORKS – AN OUTLINE

J. Göbel[1]
[1]University of Hamburg, Germany

Corresponding author: J. Göbel, Univ. of Hamburg, Dep. of Informatics
22527 Hamburg, Vogt-Kölln-Straße 30, Germany, goebel@informatik.uni-hamburg.de

Abstract. This paper presents an outline of a PhD project in progress. The target is the optimization of *transport networks*, defined as a graph topology where entities are forwarded from node to node following different routes. The progress of such entities is limited by capacity restrictions of both nodes and edges (links). Examples include urban traffic and IP networks. Optimization is conducted using a decentralized approach: Each node (router, traffic light) uses local rules to independently decide about the order in which entities (IP packets, vehicles) are processed. Existing work indicates that such *self-organizing* networks may yield a performance similar to centrally controlled systems in which solutions determined from global knowledge are enforced [5]. However, decentralized networks require less communication and their behaviour is scalable and robust.

A mesoscopic, discrete event-based traffic model to empirically evaluate the performance of a transport network is proposed which allows for efficient computation, while still providing valid results. Microscopic rule components [15] from which efficient network performance on a macroscopic level may emerge primarily rely on selfish behaviour, like (using terminology from urban traffic here) minimizing average local delays subject to a minimum “green” length per direction and to an upper bound to the “red” waiting duration for any link [14] or maximizing the absolute flow or flow per traffic blocked. Adaptivity includes dynamically adjusting phase and lane allocation. Selfish local rules have to be balanced against cooperative behaviour, e.g. permitting *green waves* by retaining platoons of approaching vehicles unless the subsequent links are congested.

1 Introduction

This paper presents an outline of a PhD project in progress. The target is the optimization of *transport networks*, defined in Subsection 2.1 as a graph topology where entities are forwarded from node to node following different routes. The progress of such entities is limited by capacity restrictions of both nodes and edges (links). Examples include networks seemingly as different as urban traffic (vehicles proceeding from one intersection to the next), conveyor-based manufacturing systems (items processed successively by different workstations) and telecommunication networks (e.g. IP packets being forwarded from router to router). Transport networks can be interpreted as complex adaptive systems [11]; attempts to adaptively optimize (i.e. minimize waiting or travel times despite traffic conditions continuously changing) such a network centrally typically imply exponential computational complexity and depend on the availability of a central server and the communication to this authority, see Subsection 2.2.

This motivates applying a decentralized approach: In the absence of a central authority, each node (router, traffic light) uses local rules to independently decide about the order in which entities (IP packets, vehicles) are processed. Existing work indicates that in *self-organizing* (Subsection 3.1) networks, a performance better or at least not significantly worse than in centrally controlled systems where solutions determined from global knowledge are enforced may emerge [5]. However, decentralized transport networks require less communication and their behaviour is adaptive, scalable and robust. In Subsection 3.2 a measurement for the “disorder” of a transport network is proposed, which is sharing typical properties of entropy in Thermodynamics [20] and Information theory. The problem of engineering self-organizing systems is addressed in Subsection 3.3; self-organizing systems can neither be developed top-down (unless suitable local rules are already known) nor bottom-up (undesired emergentes may occur), which facilitates pattern-based approaches in-between micro and macro-level.

A simulation model to empirically evaluate the performance of a transport network is presented in Subsection 4.1. This mesoscopic traffic model allows for efficient computation, while still providing valid results; the model is well-suited for discrete event simulation. Microscopic rule components [15] from which efficient network performance on a macroscopic level may emerge are discussed in Subsection 4.2. Such rule components primarily rely on selfish behaviour, like (using terminology from urban traffic here) minimizing average local delays subject to a minimum “green” length per direction and to an an upper bound to the “red” waiting duration for any link [14]. Other approaches attempt to maximize the absolute flow or flow per traffic blocked; e.g. protected left-turning in right-hand traffic is more expensive in terms of traffic blocked than right-turning. Adaptivity includes dynamically adjusting phase and lane allocation (see Figure 3). Selfish local rules have to be balanced against cooperative behaviour, e.g. facilitating *green waves* by retaining platoons of approaching vehicles (even if side road traffic has to wait longer) unless the subsequent links are congested. The paper is concluded by a short summary (Section 5).

Note that a short overview of this PhD project has already been presented to the scientific public in [8]; this contribution emphasizes the findings from the last six months and presents a more comprehensive outline.

2 Transport networks

This section introduces the application area of the PhD work, namely transport networks (Subsection 2.1) and briefly discusses the problem of optimizing this class of networks (Subsection 2.2).

2.1 Definition

Proposing a quite general definition, transport networks are graph topologies consisting of nodes and edges (links) in which entities are successively processed by the nodes and forwarded to the next node until eventually reaching the destination node. Both nodes and edges are subject to capacity restrictions. Typical examples of such network environments include IP packet transmission between routers in wired or wireless communication networks as well as urban traffic networks where vehicles are forwarded from one intersection to the next by means of traffic lights, or automated manufacturing systems where items pass through machines connected by conveyors.

The performance of such a network can for example be measured in terms of entity throughput or residence time. Performance is mostly determined by the node's policy of allocating processing capacity to the entities, e.g. which IP packet is processed next by a router or which lanes at an urban intersection receive "green" at a specific instant. Other means of influencing the network's performance may or may not be available, like discarding IP packets if quality of service requirements permit or modifying the routes of the entities (e.g. drivers do not necessarily follow recommended routes).

2.2 Optimization

In [11], the author argues that urban traffic belongs to the class of *complex adaptive systems*(cas) due to the characteristic properties

- *Aggregation of local interactions*. The complexity of the network's behaviour is "greater than the sum of its parts", potentially yielding highly sophisticated behaviour even if each node's actions are determined by simple rules; compare the complex behaviour of an ant colony arising from simple patters of ant behaviour. Similarly, the overall behaviour of the transport network can not necessarily derived directly from the local node policies.
- *Non-Linearity*. A small local change like increasing the ratio of a node's capacity dedicated to vehicles from a certain direction may yield significant changes to the network's performance that are difficult to deduce from similar system configurations, comparable to the "butterfly effect" in weather forecasting.
- *Flows*. The system's performance is determined by implicit interaction based on material or immaterial flows (e.g. vehicles, IP packets) among the nodes, including non-adjacent pairs of nodes. Such implicit interaction may be difficult to identify, e.g. so-called hidden bottlenecks.
- *Diversity*. System performance may require local adaptation, e.g. different cycle configurations in each node of the network, compare specialization in biological evolution.

which prevent acquiring an efficient node policy configuration analytically or using basic search techniques of Operations Research: Observe that the solution space of potential node policy configurations is huge and that the fitness function (mapping node configurations to the resulting network performance) has a positive, yet limited auto-correlation. This permits centralized optimization methods for such "systems at the edge of chaos" [10] performing significantly better than random search or full enumeration, i.e. providing a central authority with all data required and imposing the solution the authority has determined to all nodes; examples of such centralized optimization methods for the special case of urban traffic optimization include commercial systems like SCOOT [21] as well as scientific research like [4]. However, the computational requirements of such centralized solutions typically increase exponentially with respect to the size of the network and the dependence on a central authority and the communication channels required yield robustness and scalability problems. Moreover, the load offered to the network may evolve continuously, thus dynamic adaptation is required.

In contrast, decentralized optimization of the transport network would neither depend on a central authority nor lack scalability: Each node is free to autonomously follow its own local decision rules (and to alter them if need be); overall system behaviour *emerges* from agent interactions. The system is able *self-organize* such that desired properties (e.g. efficiency) are maintained in an equilibrium state as long as external influences (i.e. entity inflow) are constant and dynamically (re-)acquired given varying external influences.

The next section takes a detailed look at concepts of emergence and self-organization mentioned (yet not properly defined) here.

3 Self-organization

This section describes and distinguishes the terms *self-organization* and *emergence* (Subsection 3.1) and proposes how to measure order in transport networks (Subsection 3.2). Subsection 3.3 addresses the problem of how to engineer self-organizing systems.

3.1 Definitions

Self-organizing phenomena and emergent behaviour in nature or society have been observed for couple of centuries, consider the work of Adam Smith ("invisible hand of the market") or Charles Darwin ("order in nature") for examples; Jules H. Poincaré discovered that while the behaviour of weakly interacting systems may be predictable, complex systems can become unpredictable in case their parts interact strongly. Most likely, in 1947 William R. Ashby was the first to use the term *self-organization*, referring to an organization (e.g. a company) in which entities are able to adjust functions and responsibilities [1]. The modern, more general notation of self-organization as meta-science is based on Ilya Prigogine's notation of self-organization in Thermodynamics, i.e. non-equilibrium systems able to *increase order* provided energy is supplied, thus avoiding the irreversible convergence towards disorder in case energy (or matter) is dissipated [20]. Recent research goes even further, e.g. Stuart Kauffman proposing a fundamental force towards order counteracting Thermodynamics. This implies that autocatalytic systems like creatures during biological evolution do automatically tend towards self-organized criticality, yielding a promising trade-off between chaos (permitting non-trivial changes) and order (useful structures not systematically destroyed, which is inevitable in a totally chaotic regime) [12].

Literature provides many attempts to define self-organization and emergence; closely following [3, 16], *self-organization* can be understood as a dynamical process in which systems autonomously acquire and maintain order themselves despite external influence and perturbations. A system exhibits *emergence* when there is behaviour at the macro-level that coherently and dynamically arises from the interactions between the components of the system at the micro-level. Emergent behaviour is novel with respect to the individual parts of the system and in return causes feedback from macro to micro-level. Figure 1 summarizes both concepts visually.[1]

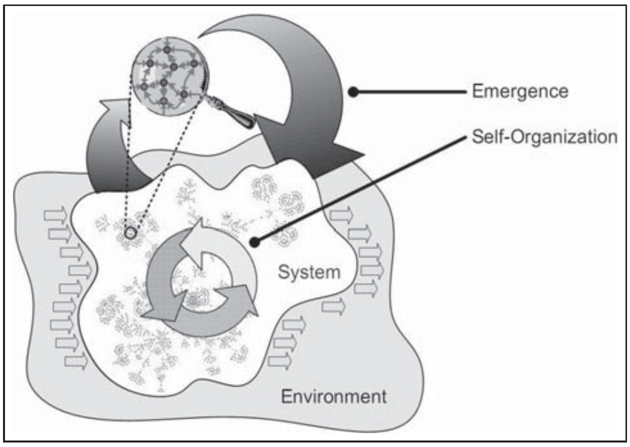

Figure 1: Emergence (micro/macro interdependency) and self-organization (maintenance of order)

3.2 Entropy

Self-organization, as defined in the previous subsection, is supposed to increase the *order* of a system. However, the term "order" has not yet been precisely specified. Again, an analogy to Thermodynamics can be made, where *entropy* can be interpreted as a measurement for the *un*availability of the energy in a system to do work [16]. For example, a snapshot of the locations of all particles of a gas in a box yields a sample from the phase space of all possible particle positions that might look similar to Figure 2 (right), where it is extremely unlikely to "randomly" switch to a distribution like Figure 2 (left) on its own (e.g. due to Brownian motion). The second law of Thermodynamics claims that the entropy or "disorder" of a system never decreases if the system is isolated, i.e. the system is not sustained by the input of energy or matter. Note that it is not necessarily always correct to identify entropy and disorder; consider again Figure 2 (left) as example: An observer interested in the spatial distribution of the particles will agree that they are more ordered than the particles in Figure 2 (right). However, an observer focussing on the particles' directions or absolute velocities might consider both configuration disordered since in both cases particle motion will be approximately uniformly distributed without preferring any direction.

Claude E. Shannon introduced the notion of entropy to Information theory, representing uncertainty in information contents; information entropy $H(X)$ of a discrete random variable X is typically defined as

$$H(X) = -\sum_{i=1}^{n} p(x_i)\log_b p(x_i) \tag{1}$$

[1]Note that an alternative view exists [9, 10] in which self-organization and emergence are considered one phenomenon. Fur the purpose of the transport network application area, it is useful to separate emergence, i.e. micro-macro interdependency, and self-organization, i.e. maintaining order; the self-organizing transport network will mostly (but not completely) rely on emergent node behaviour. Also observe that according to the definitions above, emergence without self-organization (e.g. micro/macro interdependency of gas molecules in an unbounded environment causes gas to evaporate, order not maintained) and self-organization without emergence (e.g. a multi-agent system with one of the agents in turns in charge of system-wide decisions, yielding order that does not involve the micro-level) are possible.

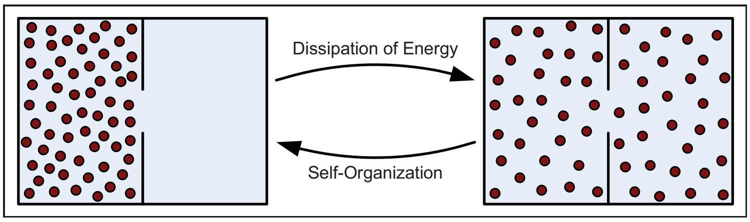

Figure 2: Low entropy (left) and high entropy (right)

where n represents the number of possible states and $p(x_i)$ is the probability of state x_i. Note that this measure corresponds to thermodynamic entropy: The higher entropy, the less information (structure) is available; compare Figure 2 (left), where every particle is located in the left box, thus more details about the spatial position of each particle are at hand than in Figure 2 (right), where a tagged particle may be located in either of the two boxes. In contrast to thermodynamic entropy, however, information entropy $H(X)$ depends on the reference system and the measuring accuracy; e.g. decomposing a state like "particle located in the left box" into sub-states like "particle located in the upper (lower) half of the left box" will in general change the information entropy [7].

Transferring the concept of entropy to transport networks may for instance rely on the lengths $L(Q_i)$ of all queues Q_i at the nodes throughout the network; this is based on the assumption that the travel times the entities spend on the links for the purpose of movement cannot be altered by the nodes, which indicates nodes should focus on reducing entity delay due to waiting for being processed. Given a transport network N this yields an entropy measurement

$$S(N) = \sum_{i=1}^{n} W_i(L(Q_i)) \tag{2}$$

where W_i is any strictly increasing, concave, differentiable weighing function, not necessarily $W_i(x) = \log_b(x)$, compare Eqn. (1); moreover, different W_i may be chosen for each queue i, e.g. assigning higher weights to important nodes (intersections, routers). This measurement of the entropy of a transport network shares important properties of thermodynamic entropy, which provides the basis for analogously applying concepts from self-organization in Thermodynamics to the specific application area of transport networks. Such properties include:

- The entropy of a system is equal to the sum of the entropies of its components.
- In systems without energy influx (isolated system), entropy never decreases (Second law of Thermodynamics); reducing entropy requires work (i.e. nodes forwarding entities).
- The entropy can be determined using the micro-states of the system or (approximately) using older macro-state and the micro-changes in-between; the measurement is self-containing all necessary information[2]:

$$\frac{\partial S(N)}{L(Q_i)} \sim L(Q_i) \tag{3}$$

This measurement of entropy can be further refined (particulary choosing a specific set of weighing functions $W_i(x)$, e.g. not only reflecting the queue lengths but also the duration spent by the relevant entities in the queues). Nevertheless, some differences to the thermodynamic notation of entropy will persist, e.g.

- reversing the relationship between entropy and spatial distribution: all entities concentrated at single queue is considered a disordered state and
- not providing an absolute measurement, similar to information entropy: the entropy of a network is not well-defined (depends on the choice of the weighing functions), while the thermodynamic entropy of any system can in principle be determined (subject to restrictions in measurement accuracy).

3.3 Engineering self-organizing transport networks

Based on the definitions of self-organization and emergence and a means of measuring order in a transport network, this subsection ventures a closer look at how to setup a self-organizing transport network, which most importantly depends on whether or not microscopic rules (i.e. nodes policies) from which the desired macroscopic behaviour (i.e. efficient or ordered network) emerges are already known. If local rules yielding a system behaviour at least sufficiently similar to the target, *engineering* the self-organizing system can be done by straightforwardly bottom-up component assembling the network; from a software development perspective, this yields a *design* approach. To provide an example, in [13] mobile sensor deployment was investigated; optimal wireless coverage was achieved by adopting animal swarming behaviour.

[2] Strictly speaking, Eqn. (3) is correct only if queue lengths are (approximately) continuous, which may be justified for many IP packets of different lengths queued at a router, yet not necessarily for vehicles queued at an intersection.

However, the absence of microscopic rules from which the desired macroscopic behaviour does result poses the problem of *reverse-engineering* a self-organizing system, for which research this far does not offer an agreed-upon and universally applicable approach [22]: An evolutionary process is required, in which candidate rules have to be designed, tested (i.e. stochastically analyzed by the means of simulation), modified and combined. Such candidate rules may be based on building blocks from microscopic rules sets of self-organizing systems already provided by nature and social sciences; a categorization (that is neither complete nor necessarily disjoint) based on [15, 16] includes:

- *Stochastic exploration and reinforcement learning*, e.g. biological evolution, an ant colony seeking food based on pheromone distribution and diffusion
- *Structure or pattern building on incremental basis*, e.g. nest building, morphogenesis
- *Aggregation*, e.g. flocking, schooling, or herding subject to dynamically different behaviour and specialization
- *Selfish behaviour*, e.g. (partially) pursuing local targets only, like personal prosperity in capitalism
- *Cooperative or altruistic behaviour*, e.g. explicit cooperation (mutual agreements or majority decisions), indirect cooperation (e.g. communication using flows or stigmergy)

Since reverse-engineering self-organizing systems requires determining suitable rule components, software development perspectives purely relying on top-down (observe that macroscopic behaviour cannot necessarily be reduced to a sum or other aggregation of local rules) or bottom-up approaches (emergents that are not desired may arise from candidate rules) are ruled out. Instead, top-down and bottom-up attempts have to be combined: A promising approach seems to be focussing on patterns in-between micro and macro-level

- which on the one hand allow for the identification of local rules inducing such patterns and
- for which on the other hand the positive impact to macro-performance of the network can be asserted.

A typical example for such a pattern are so-called green waves in urban traffic: An efficient network should enable compact platoons of vehicles to traverse successive links subject to a minimum number of stops. Subsection 4.2 will name some local rule components an efficient transport network is based on.

4 Research outline

The target of the PhD project is to recommend local rules that provide a transport network with the ability to self-organize or (using a more abstract view) at least to identify properties of such rules or to establish conditions and principles to evolve them. Subsection 4.1 presents some details about the simulation model developed for the purpose of evaluating the behaviour of the transport network, while Subsection 4.2 ventures a look at the potential building blocks for such rules themselves.

4.1 Traffic model development and simulator implementation

Local node strategies promising to perform well in a self-organizing transport network require a simulation environment for empirical evaluation. In general, traffic simulators can be classified by resolution, ranging from macroscopic approaches that merely determine flow rates to microscopic simulators, in which all entities and their attributes (route, position, velocity, acceleration. . .) are modelled explicitly; in general, complexity and run-time efficiency are traded off against exactness [17].

For the purpose of the PhD project, a *mesoscopic* approach in-between the microscopic and macroscopic extremes has been chosen, targeting at a model that is sufficiently detailed to validly evaluate the performance of a transport network at computational costs that are feasible. This mesoscopic model is an extension of the queue-based traffic model described in [18]. Table 1 shows some core concepts; note that Table 1 should be read in conjunction with Table 2, defining the symbols used in Table 1.

The basic assumption underlying the traffic model is that refraining from continuously determining each entity's position and velocity is feasible without significant expenses in validity; instead, the model keeps track of the earliest possible arrival of each entity at the next node. This ensures a link cannot be left before the entity could have traversed the distance between the incident nodes. This earliest possible arrival, however, is only realized if no slower preceding entity on the link prevents the next entity from arriving at this instant (or if passing is possible, e.g. urban multi-lane traffic) and if no entities are queued in front of the next node. In case such a queue exists, the entity will be delayed at least until the preceding entity has left the queue or even significantly longer, e.g. due to inhomogeneous acceleration and velocity-dependent safety distances in urban traffic. In summary, the model abstracts from any other microscopic entity interaction except queueing in front of the nodes and the entities' inability to pass each other (if appropriate), yielding model performance mostly determined by the allocation of the node processing capacity. The PhD project will argue in more detail that this is a valid approach for an urban network (but not for highway traffic for example) since the behaviour of the traffic light-controlled intersections dominate the advance of each vehicle.

	IP networking (packets)	Urban traffic (vehicles)
Entity data length ℓ [MBit] and physical length l [m]	$\ell > 0$ $l = 0$	$\ell = 0$ $l > 0$
Min. **space occupied** while traversing a link (o_L) and while queued at a node (o_N)	$o_L = d_L$ (assuming $\frac{\ell \cdot v}{b_L} \gg d_L$) $o_N = 0$	$o_L = l$ $o_N = l$
Condition **link blocked**	if $\sum^L o_L \geq d_L \cdot \#_L$ i.e. max. 1 packet per lane (channel)	if $\sum^L o_L + \sum^N o_N \geq d_L \cdot \#_L$ i.e. space occupied
Constraints for **nodes processing and forwarding entities**		
- Earliest arrival of the entity at the next node after $\Delta_{\text{Arr}} = t_C + t_P + t_T$, but at least $\Delta_{\text{Inter}} = \frac{\ell}{b_L} + \frac{l}{v} + s_t$ later than the previous entity on the relevant link (unless passing permitted).		
- An entity clearing its position in a queue permits the next entity to clear position after $\Delta_{\text{Next}} = t_C + t_P + s_t$ unless link out blocked or lane received "red".		
$\hookrightarrow$ Time to clear position (t_C)	$t_C = 0$	$t_C = \frac{l}{v}$ if not queued $t_C = d_v^a(q,l)$ if queued
$\hookrightarrow$ Time for processing (t_P)	$t_P = \frac{\ell}{\min b_N, b_{Lin}} + \frac{\ell}{\min b_N, b_{Lout}}$	$t_P = 0$
$\hookrightarrow$ Time to traverse link (t_T)	$t_T = \frac{d_L}{v}$	$t_T = \frac{d_L}{v}$ if not queued $t_T = d_v^a(q,d_L)$ if queued
- Note that in wireless IP networks, power available per unit of time may constrain when the next packet can be sent. Assume e.g. the node's battery can supply the maximum signal strength (i.e. transmit over the maximum distance d_{max}) during 100 ms each second (i.e. 10 % of the time); this yields the subsequent entity processed after $\Delta_{\text{Wifi}} = \max(t_P, \frac{\ell \hat{p} \cdot 0.01}{p_N} + \frac{\ell \hat{p}(d_L/d_{max})^2}{p_N})$ seconds.		

Table 1: Summary of the model logic of the transport network simulator.

a	Acceleration	#	Number of lanes
b	Bandwidth [MBit/s]	$d_v^a(s^*,s)$	Duration of traversing a distance of s after already accelerating over a distance of s^* up to maximum speed v, $d_v^a(s^*,s) =$
d	Distance		$\begin{cases} \sqrt{\frac{2(s^*+s)}{a}} - \sqrt{\frac{2s^*}{a}}, & \text{if } s^*+s \leq \frac{v^2}{2a} \\ \frac{s}{v}, & \text{if } s^* \geq \frac{v^2}{2a} \\ \frac{s^*+s}{v} + \frac{v}{2a} - \sqrt{\frac{2s^*}{a}}, & \text{otherwise} \end{cases}$
p	Power available [W/s]		
$\hat{p}$	Power needed to transmit over the maximum link length d_{max} [W/MBit]		
q	Tailback length (i.e. sum of lengths of preceding entities enqueued)		
s_d	Safety distance at max. speed, $s_d = v \cdot s_t$		
s_t	Safety time		
v	Velocity (depends on link and entity)	Indices L and N refer to links and nodes	

Table 2: Symbols used in Table 1.

Since entity state changes do only occur at well-defined instants in time, e.g. an entity arriving at or departing from a node, discrete event simulation is considered a suitable approach to implement the model logic. For the implmentation, the Java-based DESMO-J framework ("Discrete-Event Simulation and Modelling in Java"), was used. DESMO-J [19] is developed at the University of Hamburg (Germany) and can be obtained free of charge from `http://www.desmoj.de`. Key features of the transport network simulator implemented this far include strict object-oriented design which facilitates exchanging model components, e.g. local node rules, see next subsection, separation of model and experiment (using XML input files to parameterize network type, topology and traffic offered) and a prototypical visualization.

4.2 Local decision rules

The PhD project does by no means claim that exactly the same microscopic rules can be applied to networks as different as urban traffic and IP communication; on that contrary, this is unlikely: Observe e.g. that urban intersections and manufacturing systems typically require some setup time (e.g. a traffic light mutually red for a short period due to security reasons), while IP routers do not: The processing duration of an entity does not depend on the origin, route, or destination of the previous entity. Therefore, local rules from which efficient behaviour in urban traffic will emerge will have to the address trade-off of on the one hand minimizing the waiting duration of vehicles from each incident road yet on the other hand not switching too frequently, thus wasting to much processing capacity for setup purposes.

Nevetheless, this subsection attempts a look at potential building blocks for rules (within bounds) applicable to the different types of transport networks. Terminology will again be based on urban traffic. Such local rules of intersection behaviour are not allowed to take any special feature of the network topology or flows into account, like fixed lattice lengths in a Manhatten-like grid [2], one way traffic and turning forbidden [6] or all flows mutually exclusive [14]. The ambitious target of the PhD project is to develop node behaviour such that any information used by the nodes must completely be obtained locally (e.g. by the means of cameras, induction loops, microwave-based vehicle detection), so that no *explicit* communication with other intersections or a central server is required; should this turn out too optimistic, at least partially removing the restricting assumptions of the work cited in this paragraph still provides research potential.

Basic building blocks for the node policies might include selfish as well as cooperative and altruistic behaviour (compare categorization in Subsection 3.3). The nodes do autonomously determine their configuration, i.e. the flows receiving "green" and the lane allocation, compare Figure 3. Nodes (intersections) may *selfishly* minimize entropy (compare Subsection 3.2), thus minimize average (and maximum) local vehicle delays. For the trivial case of very low traffic, that can simply be achieved by giving green light to any approaching vehicle immediately or as soon as possible [5]; vehicles do hardly ever wait, entropy stays close to zero. On the contrary, if traffic demand exceeds the system's capacity, it becomes overloaded; deadlocks are possible. The entropy is irreducibly large, no matter how the traffic lights are programmed. In between these extreme cases, spontaneous switching could e.g. favour links from which the highest flow rate of vehicles would traverse the intersection, subject to a given minimum and maximum duration of each period as well as ensuring any link granted green light after a certain maximum "red" period. Dynamic lane allocation (see Figure 3, right) should not necessarily solely be based on demand (i.e. share of right/left turns) but also on intersection capacity blocked by vehicles headed into different directions; e.g. protected left-turning without need to give way to opposing traffic (assuming right-hand road traffic here) is "more expensive" than green light to vehicles turning right since it is more likely that other vehicles are blocked (e.g. left-turning implying a red signal to opposing traffic targeting straight ahead or right); allocating more lanes to left-turning vehicles reduces the duration of a dedicated green light necessary to e.g. clear the queue of such vehicles so that opposing traffic needs only be blocked for a shorter period.

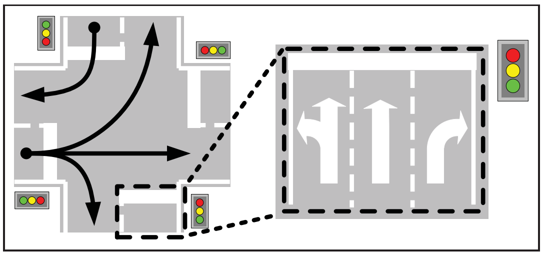

Figure 3: Potential configuration of phases (left) and lanes (right) at an instant.

On the other hand, these rules could be augmented by cooperative behaviour based on *indirect* communication: For example, traffic lights may retain platoons of approaching vehicles by e.g. delaying switching to "red", even if this means that local traffic from other links has to wait longer than selfish node behaviour would require ("inbound planning"). However, it makes sense to disrupt a platoon in case the (main) subsequent link is already congested ("outbound planning"); better allocate the intersection's capacity to other links, even if traffic flow is smaller, than waste it to a link where demand is larger yet unable to traverse the intersection.

5 Summary

This paper has described a PhD project currently in progress. The target is the decentralized and self-organizing optimization of transport networks; advantages in comparison to centralized control include scalability (no central authority required), adaptivity (no dependency on unreliable information gathered in advance, resiliently coping with perturbations), and robustness ("graceful degradation" in case of partial failures). Decentralized (local) mechanisms can be viewed as natural approach to transport network optimization since complexity *arises locally*: The output of one node becomes the input of adjacent nodes. The paper has proposed a measurement of order (entropy) applicable to transport networks and, more generally, discussed the problem of engineering self-organizing systems; evolutionary software development based on patterns in-between microscopic and macroscopic level (e.g. green waves) was identified as a promising approach. A simulation model to empirically evaluate the performance of a transport network was presented. This model is queue-based and well-suited for discrete event simulation. The underlying idea is to keep track of the earliest possible arrival at the next node for every entity; depending on queue lengths and whether or not passing slower entities is permitted, actual arrival may by significantly later than the earliest possible arrival. Finally, the paper has named building blocks for local node rules of allocating processing capacity to entities, e.g. selfishly optimizing the local relative flow (traffic receiving "green" per conflicting traffic blocked) balanced against cooperative behaviour like retaining green waves.

6 References

[1] W. R. Ashby. Principles of the self-organizing system. In H. v. Foerster and G. W. Zopf jr., editors, *Transactions of the University of Illinois Symposium*, pages 255–278, London, UK, 1962. Pergamon.

[2] A. Bazzan, D. de Oliveira, and V. Lesser. Using Cooperative Mediation to Coordinate Traffic Lights: A Case Study. In *Proceedings of Fourth International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 463–469, July 2005.

[3] T. de Wolf and T. Holvoet. Emergence Versus Self-Organisation. In S. A. Brueckner et al., editors, *Engineering Self-Organising Systems*, pages 1–15, Berlin, Germany, 2005. Springer.

[4] C. Diakaki, V. Dinopoulou, K. Aboudolas, M. Papageorgiou, E. Ben-Shabat, E. Seider, and A. Leibov. Ex-

tensions and new applications of the traffic signal control strategy TUC (Traffic-responsive Urban Control). *Transportation Research Board*, 1856:202–211, 2003.

[5] C. Gershenson. Self-Organizing Traffic Lights. *Complex Systems*, 16(1):29–53, 2005.

[6] C. Gershenson. Design and control of self-organizing systems. PhD Thesis, Vrije Universiteit Brussel, 2007.

[7] C. Gershenson and F. Heylighen. When Can we Call a System Self-organizing? In W. Banzhaf, T. Christaller, P. Dittrich, J. T.Kim, and J. Ziegler, editors, *Advances in Artificial Life, 7th European Conference, ECAL 2003*, pages 606–614, Dortmund, Germany, 2003. Springer.

[8] J. Göbel. Self-Organising Transport Networks. In T. Schulze, editor, *Simulation and Visualization 2008 – Proceedings of the Postgradute Forum*, pages 1–10, Magdeburg, Germany, Feb 2008.

[9] D. Helbing and T. Vicsek. Optimal Self-Organization. *New Journal of Physics*, 1:13.1–13.17, 1999.

[10] F. Heylighen. The Science of Self-organization and Adaptivity. In L. D. Kiel, editor, *Knowledge Management, Organizational Intelligence and Learning, and Complexity*, Oxford, UK, 2001. Eolss.

[11] J. H. Holland. *Hidden Order – How Adaption builds complexity*. Basic Books, New York, USA, 1995.

[12] S. Kauffman. *At Home in the Universe – The search for the Laws of Self-Organization and Complexity*. Oxford University Press, Oxford, UK, 1995.

[13] A. E. Krzesinski and J. Göbel. A model of autonomous motion in ad hoc networks to maximise coverage. Unpublished, 2007.

[14] S. Lämmer. *Reglerentwurf zur dezentralen Online-Steuerung von Lichtsignalanlagen in Straßennetzwerken*. PhD Thesis, Technical University of Dresden, 2007.

[15] M. Mamei, R. Menezes, R. Tolksdorf, and F. Zambonelli. Case Studies for Self-Organization in Computer Science – Nature-Inspired Applications and Systems. *Journal of Systems Architecture*, 52(8–9):443–460, 2006.

[16] G. Di Marzo-Serugendo, M.-P. Gleizes, and A. Karageorgos. Self-Organisation and Emergence in MAS: An Overview. *Informatica*, 30:45–54, 2006.

[17] A. D. May. *Traffic Flow Fundamentals*. Prentice-Hall, Englewood Cliffs, USA, 1990.

[18] K. Nagel. Traffic networks. In S. Bornholdt and H. G. Schuster, editors, *Handbook on networks*, New York, USA, July 2003. Wiley.

[19] B. Page and W. Kreutzer. *The Java Simulation Handbook - Simulating Discrete Event Systems with UML and Java*. Shaker, Aachen, Germany, 2005.

[20] I. Prigogine. *Order Out of Chaos – Man's new Dialog with Nature*. Bantam, New York, USA, 1984.

[21] United Kingdom Department for Transport. *The "SCOOT" Urban Traffic Control System*. Traffic Advisory Leaflet 04, 1995.

[22] F. Zambonelli, M.-P. Gleizes, M. Mamei, and R. Tolksdorf. Spray Computers: Frontiers of Self-Organization. In *First International Conference on Autonomic Computing (ICAC'04)*, pages 268–269, 2004.