

Decentralized real-time control of water distribution networks using self-organizing multi-agent systems

Florian Dotsch, Jorg Denzinger, Holger Kasinger, Bernhard Bauer

Angaben zur Veröffentlichung / Publication details:

Dotsch, Florian, Jorg Denzinger, Holger Kasinger, and Bernhard Bauer. 2010.
"Decentralized real-time control of water distribution networks using self-organizing multi-agent systems." In *2010 Fourth IEEE International Conference on Self-Adaptive and Self-Organizing Systems*, 27 Sept. - 1 Oct. 2010, Budapest, Hungary, edited by Márk Jelasity and Janos Sztipanovits, 223–32. Los Alamitos, CA: IEEE.
<https://doi.org/10.1109/saso.2010.20>.



Decentralized Real-time Control of Water Distribution Networks Using Self-organizing Multi-Agent Systems

Florian Dötsch*, Jörg Denzinger[†], Holger Kasinger* and Bernhard Bauer*

**Department of Computer Science*

University of Augsburg, 86135 Augsburg, Germany

Email: {doetsch, kasinger, bauer}@informatik.uni-augsburg.de

[†]Department of Computer Science

University of Calgary, Calgary Canada T2N 1N4

Email: denzinge@cpsc.ucalgary.ca

Abstract—In this paper we present a self-organizing emergent multi-agent system designed for the decentralized real-time control of water distribution networks. The agents, which are associated with every pump, tank, and water tower of such a distribution network, interact by means of a decentralized coordination mechanism based on digital infochemicals. In particular, they utilize a biologically-inspired coordination strategy, called indirect defense, which is used by certain plants in nature in order to get rid of herbivores. The experiments, which we have conducted on several different water distribution networks including one real-world municipal distribution network, demonstrate that the presented variants of our approach are able to achieve a self-organizing emergent solution comparable to the decisions made by inexperienced as well as experienced human operators, but with a higher flexibility, robustness, scalability, adaptivity, and autonomy.

Keywords—self-organization; multi-agent system; biologically-inspired; water distribution; infochemicals

I. INTRODUCTION

Due to their inherent flexibility, robustness, scalability, adaptivity, and autonomy, self-organizing emergent solutions are predominantly suited for problem classes characterized by high complexity and dynamics, openness, large scale, unpredictability, as well as dynamic disturbances. A problem class exhibiting these characteristics is resource distribution, whereby a resource can be electricity, gas, oil, or water. Because the problem constraints vary depending on the resource in hand, solution approaches can not be transferred between different problem domains without further ado. For instance, whereas huge capacities of gas, oil, and water can be temporarily stored, there exists no such storage for huge capacities of electricity. Also, the distribution networks for each of these resources differ significantly in terms of materials, physics, timing constraints, etc.

The focus of this paper is on water distribution by use of a water distribution network (WDN) [1]. A WDN consists of a network of pipes that connect water supplying components such as reservoirs (lakes, rivers), underground tanks, or water towers, with customers such as private households or commercial facilities. The transportation of water from the

supplying components to the customers within a WDN can be accomplished either by gravity or, due to rolling terrains, more usually by pumps in pumping stations. Because the pumping of water requires a lot of energy, the interest of water suppliers is to minimize pumping and energy costs, respectively. To the contrary, consumers expect their water demands to be fulfilled at an comfortable and satisfactory level even during peak times.

Today, the operation of a WDN is accomplished centrally. Although a number of software tools exist that gather and analyze operational data of a WDN in real-time, operational management decisions, in particular to turn on or off a pump, are made by certified human operators. In order to utilize the beneficial properties mentioned at the beginning of this section for the operation of WDNs, in this paper we present an agent-based, self-organizing emergent system for the real-time control of a WDN. This inherently decentralized approach assumes that each pump and tank within a WDN can be equipped with an agent able to communicate with each other in order to minimize energy costs. Please note that in particular the vast number of consumer interfaces in a WDN do *not* have to be equipped with any agents respectively expensive and fragile sensors, which significantly reduces the capital expenditures of water suppliers.

For the self-organizing, decentralized coordination of the agents, we have designed a mechanism inspired by the so-called *indirect defense* phenomenon [2], a strategy that certain plants apply in biology to get rid of herbivores. Because this strategy is based on chemical stimuli respectively infochemicals [3], we have based this decentralized coordination mechanism (DCM) on the digital infochemical coordination model [4], which adapts the general principles behind infochemical-based coordination (IBC) in nature. We have developed two variants of this DCM: a basic variant, which is intended to correspond to the decision making of an inexperienced human operator, and an enhanced variant, which includes more information within an infochemical and is intended to correspond to the decision making of an experienced human operator.

We have executed experiments with varying consumer demand patterns, different energy consumption curves of the pumps, as well as different WDN topologies. Whereas most of these topologies were constructed especially for evaluation purposes, one topology is based on the municipal WDN of Novato, California. For the experiments we have used and extended the EPANET toolkit [5], a water distribution simulation software released by the US Environmental Protection Agency (EPA) under the public-domain license. The experimental results demonstrate that already the basic variant of the mechanism is able to achieve effective solutions, while the enhanced variant is able to improve these solutions up to 86%, depending on the demand patterns, the energy consumption curves, and the topologies used.

The rest of this paper is organized as follows: Section II provides the background on WDNs as well as on IBC. Section III describes the self-organizing emergent solution approach, starting with the biological inspiration, its adaptations, as well as their application to water distributions networks. Section IV presents the experimental evaluation, whereas Section V overviews related work. Finally, Section VI presents concluding remarks as well as an outlook.

II. BASICS

In this section, we provide brief introductions into WDNs and IBC. The former, obviously, represents the application domain for our DCM proposed in Subsection III-B, which is an instantiation of the latter.

A. Water Distribution Networks

Water distribution networks, as they are used in our cities to bring water to both, residential and commercial customers, on a high level can be considered as a graph (V, E) of nodes $v \in V$ of various types connected by edges $e \in E$, also of various types. More precisely, V comprises *junctions* $ju \in Ju$, *reservoirs* $re \in Re$, and *tanks* $ta \in Ta$ (or water towers). Furthermore, E comprises *pipes* $pi \in Pi$ and we also treat *pumps* $pu \in Pu$ and *valves* $va \in Va$ as edges.

Operating a WDN is all about providing/maintaining pressure, which is called in water hydraulics *head*. If the head in the right places in the WDN is sufficient enough, a sufficient *flow* of water towards the demand sites will be present. One basic way of producing sufficient head in a place in the network is by using gravity, i.e. having these places sufficiently elevated compared to the nodes at which a certain flow of water is leaving the network. The obvious nodes in a network that should produce sufficient head are the reservoirs Re , since they are the nodes where water enters the network (in this work we will assume that each reservoir can produce unlimited amounts of water). But given that the role of a reservoir in a WDN is usually played by lakes, rivers, or even wells, having the reservoirs elevated higher than the nodes the water has to flow to is often not an option. This is where tanks and pumps enter the picture.

A tank ta is placed into a WDN to provide head and to act as a buffer allowing the WDN to take some time to react to changes in demand due to the water stored in ta . A tank is characterized by its elevation in the WDN, by the minimum, maximum, and current operation water level as well as its volume-height curve that provides the amount of water stored in the tank given its current water level. Since tanks do not provide unlimited amounts of water, it is necessary to get water from the reservoirs to a tank, which is accomplished using pumps. While in theory using only one pump to connect a tank with one reservoir is enough (which is why we treat a pump as an edge in the graph), in reality some kind of redundancy is required to deal with the need for maintenance and scheduled repairs/upgrades, but also with unforeseen events like a broken pipe. As a consequence, a tank usually can be replenished by several pumps, but often at different costs. These costs are measured either with regard to wear and tear of equipment as well as other criteria, or with regard to energy use, which is what we are interested in in the following.

A pump pu injects mechanical energy into a WDN to move water from lower elevations to nodes with higher elevation, which are the aforementioned tanks, as stated in the last paragraph. A pump, as a directed edge in the graph representing a WDN, creates head at the “receiving” end of the edge dependent on the flow through the pump from the reservoir side. This dependency is described by a so-called *pump curve*, which is also dependent on the energy put into the system by the pump (in case of the commonly used centrifugal pumps this energy depends on the speed of the rotating impeller of the pump [6]). Modern pumps allow for different amounts of energy created by the pump, which means that a pump operator can choose the pump curve he prefers for the current network conditions (achieving the appropriate efficiency of the pump), introducing a second way to influence the energy usage by the WDN (in addition to the choice of pump). Naturally, the mechanical energy put into the WDN by a pump does not equal the electrical energy used by the pump, because of heat loss due to friction.

The influence of a pump and its performance on a particular WDN is governed by a set of complex equations, taking into account the network, the characteristics of all nodes (e.g. elevation, demand at junctions, which are normally used to provide exit points for the water), and the characteristics of the edges (e.g. pump curves, diameter, or pipe lengths). Furthermore, by opening or closing valves, the network structure can be changed. Fortunately, the operation of a WDN does not really require the continuous solving of this set of equations reacting to changes in demand and trying to find a solution that requires minimal energy. A reactive approach looking at the fill levels of the tanks (and using operational minimum and maximum water levels for each tank) is usually sufficient to meet the demands by the customers. Experienced human operators usually also are

able to operate the pumps in reaction to the fill levels of the tanks with a "not so bad" energy consumption. Even designing WDNs is nowadays not using the set of equations directly, instead simulators, such as e.g. the EPANET toolkit, are used to provide feedback about what is happening in a WDN. Such simulators also allow for the evaluation of the performance of the self-organizing multi-agent system (MAS) solution that we present in Section III-B.

While the main purpose of a WDN is to fulfill the demands by the customers at all times, keeping the energy consumption required therefore as low as possible is naturally in the interest of the organization or company responsible for a particular WDN. Since the customer demands are usually not fully predictable, keeping the energy consumption low is a dynamic optimization problem that is well suited for self-organizing MASs.

B. Infochemical-based Coordination

Infochemical-based coordination is the most universally employed communication and coordination model between organisms in biology (cf. [7]). Thus, IBC provides a plethora of inspiring examples that can be adopted as DCMs for self-organizing emergent solutions to (artificial) problem instances in various application domains. For instance, pheromone-based coordination, as e.g. used during foraging in ant colonies, is probably the best-known example for a natural DCM that uses the principles of IBC. It has been successfully used for finding the solutions to problems occurring e.g. in manufacturing control (see [8], [9]). However, in [10] we have shown that in nature not only organisms of the same species coordinate by means of chemical stimuli, but also organisms of different species. These chemical stimuli are then called allelochemicals and are distinguished into allomones, kairomones, and synomones, depending on the fact which organism, i.e. the emitting one or the receiving one, benefits in the interaction. For instance, pollination-inspired coordination (PIC) [11] is inspired by the behavior of honey bees during the pollination of flowers, whereby the latter emit synomones to attract the former. This mechanism has been successfully used for the self-organizing solution to problems occurring e.g. in internal logistics. Moreover, the indirect defense strategy, which serves as a biological inspiration for the solution approach presented in this paper, involves organisms of three different species that coordinate their local behavior by means of infochemicals.

In order to provide a foundation for the systematic identification and specification of DCMs based on IBC, the digital infochemical coordination (DIC) model [4] adopts the general principles behind the communication and coordination by means of chemical stimuli between organisms in biology. In particular, the DIC model allows for the combination of multiple infochemicals having different semantics, dynamics, and functions, which in turn allows for more efficient self-organizing emergent solutions and solution processes.

In the DIC model, an agent (as a coordination element) of a certain type reflects a living organism of a certain species, able to interact indirectly by emitting and perceiving (digital) infochemicals. A (digital) infochemical ι is defined as

$$\iota = (\gamma, \gamma^{thresh}, \delta, \epsilon, \psi) \quad (1)$$

γ is the current concentration of ι and reflects the dynamically changing concentration of diffusing biological infochemicals. γ^{thresh} is the threshold concentration of ι and reflects the behavioral threshold concentration of living organisms reacting to a specific infochemical. δ is the diffusion coefficient of ι and reflects the chemical diffusion coefficient that allows for a very fine-tuned propagation radius and evaporation time specific to each infochemical. ϵ is the emitter of ι and reflects the biological fact of an infochemical to reveal information on its emitter. ψ is individual information encapsulated by ι and reflects the biological role of an infochemical as a dynamic information carrier. The content of ψ depends on the concrete DCM respectively the domain a solution has been designed for.

The coordination media (space) in the DIC model is represented by an infochemical environment, which enhances the execution infrastructure of the agents, providing them with an active environment where they may share information. Consequently, an agent is situated in a spatial environment consisting of multiple locations the agents may be situated on. Connections between the locations define the possible ways an agent may choose from in order to move between the locations, whereupon the connections may be of different length, depending on the physical conditions.

The coordination laws in the DIC model are defined by a propagation function (controlling the amount and rate of infochemicals propagated from one location to a connected neighbor location), an aggregation function (controlling the aggregation of one or more infochemicals e.g. to an infochemical with higher concentration), and an evaporation function (controlling the amount and rate of the evaporation of infochemicals). Please note, as soon as the concentration of an infochemical is below its threshold concentration, either due to its propagation or its evaporation, the infochemical will be removed from the infochemical environment. All of these functions have to be instantiated by a concrete DCM that instantiates the DIC model.

Based on the perception of an infochemical, an agent may reactively execute one or more actions, which can either be to move to another location, to emit infochemicals by itself as well as to perform an action on its current location, or to interact directly with other agents. However, an agent may also proactively execute actions. Both variants of the DCM presented in the next section instantiate the DIC model.

III. SOLUTION APPROACH

Based on the basics provided by the last section, this section describes on the one side the biological paradigm of

indirect defense and on the other side the indirect defense coordination mechanism that can be used for self-organizing emergent MAS solutions controlling and optimizing the operation of WDNs.

A. Biological Inspiration

In biology, plants are permanently faced with the risk of being attacked and damaged by herbivores or fungi. Thus, over the last millions of years some plants have developed a defense strategy, called *indirect defense* [2]. In order to get rid of herbivores, they emit certain scents that attract natural enemies of these herbivores, which upon their arrival deter or destroy the attacking herbivores but not the plants. From the ecological point of view, these scents are classified as infochemicals, in more detail synomones, as both the emitting organism (plants) and the receiving organism (natural enemies) benefit from this interaction.

B. The Indirect Defense Coordination Mechanism

Similarly as in biology, the tanks in a WDN are permanently faced with the risk of head loss by the customers' demand. By adopting the indirect defense strategy for WDNs (see Figure 1), in case of a problematic head loss, a tank agent associated with a tank will emit digital infochemicals that may be perceived by pump agents associated with pumps, which in turn start pumping water in the WDN, freeing the tanks of their problematic situation. We term this indirect defense coordination (IDC). As we have adopted the general principles behind IBC in the DIC model (see Subsection II-B), we base the IDC mechanism (and concrete variants of it) on this prior work. The main objective of IDC will be to keep the hydraulic integrity and to optimize the control of WDNs with regard to energy costs. To foster energy savings, IDC keeps the head in a WDN as low as possible, but still high enough to satisfy regionally regulated minimum head limits.

In the following, we present two variants (or modes) for IDC that handle the behavior of tank and pump agents. The basic variant, which is called Greedy Mode (GM), will be able to control a WDN similar to an inexperienced operator. As soon as a problematic situation arises in the WDN, GM will just determine the easiest and most obvious solution that fixes the situation, which is to turn on or off (respectively increase or decrease the flow through) the pump that is located closest to the tank that detected the problematic situation. This is the easiest solution from the viewpoint of a human operator observing the network status on monitor walls, which is the current setup for human operators. In the more advanced variant, which is called Coordinated Greedy Mode (CGM), the pump agents compare all available solutions for an adequate reaction to the particular situation and chose the most efficient one. Thus, CGM will be able to control a WDN similar to an experienced operator and perform better than GM.

1) *Greedy Mode*: In GM, a tank agent Ag_{ta_i} is permanently monitoring the head h of the tank $ta_i \in Ta$ it is associated with. To figure out, if h is converging to either the minimum head limit h_{min} or maximum head limit h_{max} , Ag_{ta_i} permanently predicts future head values based on stored historical head values and pressure fluctuations of the past. If h dropped during the latest measurements, Ag_{ta_i} estimates the next head values for a given prediction horizon by linear regression. If any of these predicted head values is below h_{min} , Ag_{ta_i} will detect h converging to h_{min} before h_{min} actually is reached. The detection of h converging to h_{max} works analogously.

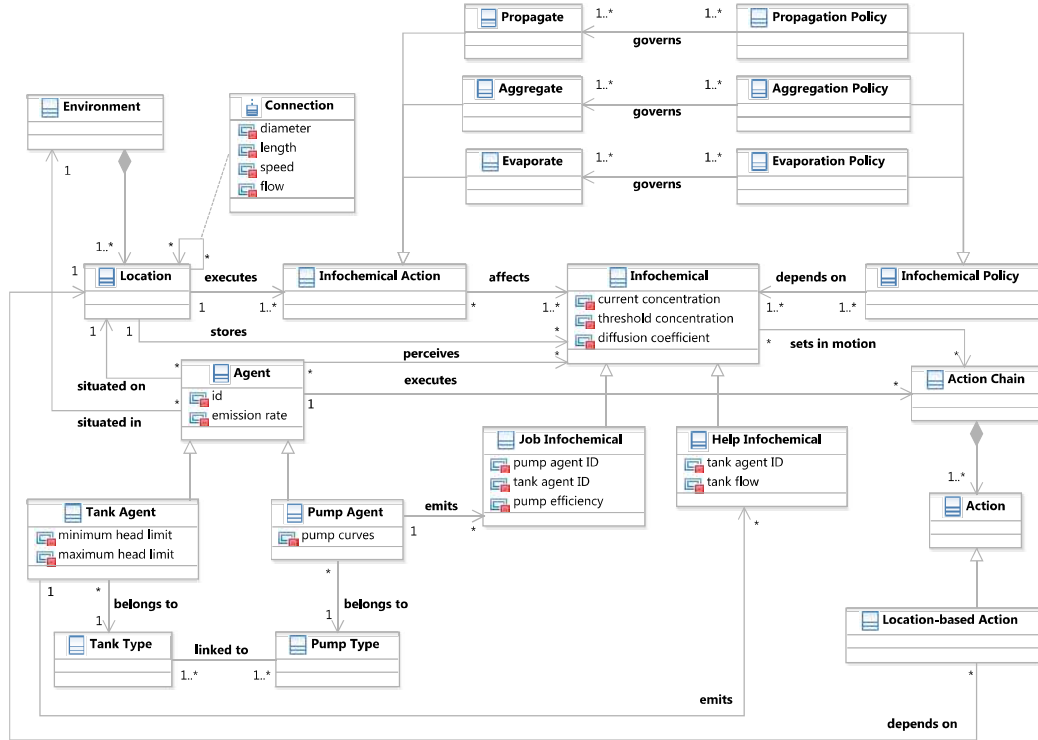
Whenever Ag_{ta_i} detects that h converges to h_{min} or h_{max} , which represents a problematic situation for the agent, it emits a help infochemical τ_{i_l} ($l \in \mathbb{N}$) to let reachable pumps know that it needs help in the form of a pressure change. A help infochemical τ is defined as an instantiation of a digital infochemical of the DIC model:

$$\tau = (\gamma_\tau, \gamma_\tau^{thresh}, \delta_\tau, id_{Ag_{ta}}, q_{ta}) \quad (2)$$

where γ_τ is the current concentration of τ , γ_τ^{thresh} is the threshold concentration of τ , δ_τ is the diffusion coefficient of τ , $id_{Ag_{ta}}$ is the unique id of the emitting tank agent Ag_{ta} , and q_{ta} (as the individual information ψ) is the current inflow of ta , whereby negative values represent an outflow.

The propagation of infochemicals, as one of three specific infochemical actions (see Figure 1), between two locations in the environment is controlled by a propagation function: when an infochemical is propagated, its concentration γ is reduced by the length of the pipe connecting the two locations multiplied by its diffusion coefficient δ . The diameter, speed, and flow of a pipe do not influence the propagation of an infochemical. As soon as $\gamma < \gamma^{thresh}$, the propagation is stopped. In IDC, an aggregation of infochemicals is not required. However, infochemicals evaporate due to a defined evaporation function: in every iteration (time step), the concentration γ is reduced depending on the diffusion coefficient δ . Again, as soon as $\gamma < \gamma^{thresh}$, the infochemical will be removed from the environment.

Due to these infochemical actions, the help infochemical τ_{i_l} is propagated through the environment and finally might reach a pump agent Ag_{pu_j} associated with a pump $pu_j \in Pu$, whose type is linked to the tank type. In order to figure out, if pu_j is able to provide help for the perceived request, Ag_{pu_j} has to determine, to what speed level the pump's impeller would have to be set in order to provide the required help. Therefore, Ag_{pu_j} takes into account the current head at pu_j 's discharge side, the current flow of the pump q_{pu_j} , as well as the information about the inflow/outflow q_{ta_i} included in the perceived help infochemical emitted by Ag_{ta_i} . Based on these values and the given pump curves, Ag_{pu_j} calculates, if pu_j can be set to a speed level that additionally provides the required flow



$|qt_{ai}|$ (in case of a negative value, i.e. an outflow of water in the tank) while maintaining the current head at pu_j 's discharge side. If several speed levels would be possible, Ag_{pu_j} will choose the one that keeps the head in the WDN as low as possible. In case that there is an inflow of water in the tank, i.e. $qt_{ai} > 0$, Ag_{pu_j} consequently might have to reduce the current speed level of pu_j .

In any case, before the speed level of a pump pu_j , which is able to help, is set to a new value, Ag_{pu_j} first emits a job infochemical ρ_{j_k} ($k \in \mathbb{N}$) to inform other pump agents Ag_{pu_s} with $j \neq s$, which might have received the initial help infochemical τ_{i_l} as well, that pu_j is able and willing to help. A job infochemical ρ is defined as an instantiation of a digital infochemical of the DIC model as well:

$$\rho = (\gamma_\rho, \gamma_\rho^{thresh}, \delta_\rho, id_{Ag_{pu}}, id_{Ag_{ta}}, \gamma_\tau^{Ag_{pu}}) \quad (3)$$

where γ_ρ is the current concentration of ρ , γ_ρ^{thresh} is the threshold concentration of ρ , δ_ρ is the diffusion coefficient of ρ , $id_{Ag_{pu}}$ is the unique *id* of the emitting pump agent Ag_{pu} , id_{ta} is the unique *id* of the tank ta the job is referencing to, and $\gamma_\tau^{Ag_{pu}}$ is the concentration of the help infochemical τ at the time Ag_{pu} has perceived it. Again, $id_{Ag_{la}}$ and $\gamma_\tau^{Ag_{pu}}$ together represent the included individual information ψ .

As in principle more than one pump might be able and willing to help, all associated pump agents wait for a given amount of time, if they perceive job infochemicals of

other pump agents. In case that another job infochemical is perceived, a comparison function f^{comp} inside each pump agent determines, which of the pump agents will finally claim the job. Therefore, if a pump agent Ag_{pu_j} perceives a job infochemical ρ_s of another pump agent Ag_{pu_s} but with the same id_{ta_i} of the request that it intends to help or is currently helping itself, f^{comp} will compare $\gamma_{\tau}^{Ag_{pu_j}}$ and $\gamma_{\tau}^{Ag_{pu_s}}$. Because due to the propagation function the concentration of any infochemical decreases with the distance to its emitter, the pump agent with the higher perceived concentration $\gamma_{\tau}^{Ag_{pu_i}}$ is closer to the emitting tank agent and will claim the job, i.e. change its speed level, whereas the other pump agents will not start (or even stop) helping. If due to any reason the perceived concentration values are the same, a pseudo random algorithm determines the helping pump agent based on the comparison of the unique *ids*. Because thus always the *nearest* pump capable able to help will be elected, we called this variant “greedy”.

Because Ag_{tta_i} is able to perceive emitted job infochemicals by pump agents as well, it gets automatically informed, if its help request is handled by at least one pump agent. Otherwise, similarly to nature, it will continue emitting help infochemicals, as long as help is under way.

2) *Coordinated Greedy Mode*: Whereas the GM of IDC only focuses on the hydraulic integrity within a WDN in the first instance, it does not pay attention to the efficiency in

terms of energy use. Thus, the CGM of IDC will extend GM and improve the overall energy consumption in a WDN. This means, in contrast to GM, where the distance of a pump to a help requesting tank was of relevance, in CGM only the provided efficiency of a pump is of relevance. The efficiency η_{pu_j} of a pump pu_j with variable speed depends on the flow q_{pu_j} through the pump as well as the chosen impeller speed level and the provided head at the discharge side of the pump. Thus, the efficiency can be obtained from the pump curve (see [6] for more details).

In order to figure out if pu_j is able to provide help for a perceived request, in CGM the pump agent Ag_{pu_j} associated with pu_j has to determine the most efficient speed level of the pump's impeller, which is able to provide the required help. If several speed levels would be possible, Ag_{pu_j} will hence choose the one that provides the highest energy efficiency. For the improvement of the coordinated solution, consequently η_{pu_j} is included in the job infochemical instead of the concentration value $\gamma_{\tau}^{Ag_{pu}}$. This requires a slightly changed definition of a job infochemical:

$$\rho = (\gamma_{\rho}, \gamma_{\rho}^{thresh}, \delta_{\rho}, id_{Ag_{pu}}, id_{Ag_{ta}}, \eta_{pu}) \quad (4)$$

where γ_{ρ} is the current concentration of ρ , γ_{ρ}^{thresh} is the threshold concentration of ρ , δ_{ρ} is the diffusion coefficient of ρ , $id_{Ag_{pu}}$ is the unique *id* of the emitting pump agent Ag_{pu} , id_{ta} is the unique *id* of the tank *ta* the job is referencing to, and η_{pu} is the efficiency that can be achieved by pu .

Consequently, in CGM also the comparison function f^{comp} inside each pump agent has to be changed for comparing the efficiency values included in the job infochemicals. In case that the speed level of a pump has to be increased, the most efficient pump will get the job. In case that the speed level has to be reduced, the least efficient pump will reduce its speed or will even be turned off. As these changes by CGM only affect pump agents, the behavior of tank agents as well as their emitted help infochemicals is not changed compared to GM. The experimental evaluation in the next section will demonstrate the efficiency improvement in terms of energy that can be achieved with CGM.

IV. EXPERIMENTAL EVALUATION

This section presents the general setup of the experiments we have performed as well as describes four series of experiments that evaluate the two IDC variants.

A. Experimental setup

To evaluate the two IDC variants for WDNs, we have modified the EPANET toolkit to allow the agents to receive local information at the node or edge they are representing and to allow the pump agents to control their edges. Like EPANET, we have implemented the MAS in C++.

In each experiment, the simulation within EPANET went through several days of demands, each day having identical

demand profiles. The simulation updated every minute of simulated time. Despite that, the wall clock time used for a whole simulation was between 2 minutes and 2 hours (depending on the particular experiment), which includes the time needed for the agents to make their decisions and for the system to propagate the infochemicals. This means that the overhead produced by the IDC approach is neglectable given the reaction times of real-life WDNs (that are simulated by EPANET).

With regard to the infochemicals, for each experiment we have manually chosen as emitted concentration for each of them a value that allowed the infochemicals to reach each other agent in the network with a sufficient concentration to be above the threshold concentration for reacting to the infochemical. The diffusion coefficient for each infochemical was chosen to keep the concentration above the threshold long enough for the whole decision process to take place.

B. Experimental results

The first experimental series evaluates how the two variants of the IDC approach deal with heterogeneous pumps. The WDN (see Figure 2) consists of three pumps (pu_1 , pu_2 , pu_3), three reservoirs (re_1 , re_2 , re_3), two identical tanks (ta_1 , ta_2), four junctions (ju_2 , ju_3 , ju_4 , ju_6) following a typical residential demand pattern as depicted in Figure 3, one junction (ju_5) using an industrial demand pattern as depicted in Figure 4, three junctions without demand, and the necessary pipes.




Figure 2. WDN for first experimental series

In the experiments, each pump is of a different type, having different peak efficiencies (with regard to head) at 50, 60, and 70 percent of the maximally produced flow. The numbers in the first column of Table I (for the first group of experiments) indicate which pump pu_1 - pu_2 - pu_3 is of what type in the particular experiment. Each experiment repeated the same daily demand pattern for 5 days. As the results show, the GM variant has different energy costs depending on pump pu_1 , while the CGM variant, not surprisingly, results in the same energy cost regardless of pump type configuration. pu_1 is the pump in the network that is the most used by the GM variant, so that a low efficiency pump in that spot is really not good for GM.




Figure 3. Residential demand pattern




Figure 4. Industrial demand pattern

Table I
COMPARISON BETWEEN IDC VARIANTS USING ENERGY USAGE
(IN KWH) PER CUBIC FEET OF PUMPED WATER

Experiment	GM	CGM
50-60-70	0.306	0.153
50-70-60	0.306	0.153
60-50-70	0.158	0.153
60-70-50	0.158	0.153
70-50-60	0.150	0.153
70-60-50	0.150	0.153
demand1	0.475	0.155
demand2	0.264	0.154
demand3	0.215	0.152
demand4	0.159	0.159
many pumps	0.196	0.159
many tanks	0.272	0.227
Navato, broad	2.781	2.498
Navato, medium	3.601	2.491
Navato, narrow	19.430	2.671

In the second experimental series, we have evaluated the two variants with regard to different demands. The WDN (see Figure 5) consists of two pumps (pu_1 with efficiency peaking at 70 percent of the maximal flow and pu_2 with its peak at 100 percent), two reservoirs, two (identical) tanks, six junctions, and the necessary pipes. For this WDN, GM, in general, favors pu_1 to fill ta_1 and pu_2 to fill ta_2 .




Figure 5. WDN for second experimental series

For the four experiments in this series, we opted to create different demand patterns by adding additional demands at different junctions. In *demand1*, ju_6 was the only junction with a demand as indicated by Figure 6. This demand pattern

was repeated over 10 days (as it was done in the three other experiments in this series to evaluate the different patterns over a longer period). For *demand2*, we used the demand pattern at ju_6 from *demand1* and added to it the demand pattern depicted in Figure 7 at ju_3 . *demand3* adds to *demand2* the demand pattern depicted in Figure 8 at junction ju_1 and finally *demand4* adds at ju_4 the daily demand pattern from Figure 9.




Figure 6. Demand pattern for ju_6




Figure 7. Demand pattern for ju_3




Figure 8. Demand pattern for ju_1




Figure 9. Demand pattern for ju_4

As Table I shows, CGM outperforms GM in the first three experiments and in the first two experiments the difference is even substantial. Again, not surprisingly, the fourth experiment that leaves the pump agents with not much choice, since it has high demand nearly all of the time, results in identical quality.

The third experimental series looks at how the MASs deal with competition, more precisely with having more helping components than demanding ones in “many pumps” and more demanding components than helping ones in “many tanks”. While the first experiment is just a more complicated variant of the situations from the first experimental series, the second experiment looks at possibilities for confusion among the active pump components due to various help requests competing for attention. We assumed that the later experiment could show suboptimal behavior of the CGM variant, which, as Table I shows, was however not the case.

In more detail, the WDN for “many pumps” (see Figure 10) uses five different pumps with different peak efficiencies each with its own reservoir to draw from and two different tanks with ta_1 having a bigger volume than ta_2 . There are eleven junctions, where four ($ju_3, ju_4, ju_6, ju_{10}$) use the residential demand pattern from Figure 3, two (ju_5, ju_7) the industrial pattern from Figure 4, and the rest do not produce any demand. The network topology for “many tanks” is nearly the same, except that we replaced the edges (re_5, ju_5), (re_4, ju_4), and (re_3, ju_{10}) by edges (ta_5, ju_5), (ta_4, ju_4), and (ta_3, ju_{10}), which eliminates three pumps from the network. Both simulations ran over 5 days.




Figure 10. WDN for third experimental series

As Table I shows, for both networks CGM clearly outperforms GM with regard to energy used per transported cubic feet of water. Looking more closely into the behaviors during the simulations, we noticed that in the case of “many pumps” GM, nor surprisingly, favors the pumps nearest to the two tanks and only gets the other pumps involved if those nearest pumps are running at full capacity, which is unfortunately not the peak efficiency for these pumps. On the other side, CGM has those other pumps coming online already when the nearest two pumps have reached their peak efficiency, thus resulting in the better behavior. Looking into the behavior of GM and CGM for “many tanks” shows that, as expected, the tanks always start sending help requests at nearly the same time. But, due to having one network, fulfilling the first three requests already also covers the other two requests. Looking at the behavior created by GM, we noticed that after the first two requests are handled by the two pumps, the third request often gets handled by the pump that is already working near peak efficiency, thus straining this pump beyond this peak efficiency. CGM, however, always has the pump working at lower efficiency dealing with the third request, thus resulting in a better overall performance. But this example shows that both variants are not creating optimal solutions, since the third request would be optimally dealt with by having both pumps kind of “partially” react resulting in both running near their peak efficiency.

The final experimental series evaluates different operating ranges for tanks for the WDN for the city of Novato, California. The WDN of this city is the most complex

Table II
OPERATIONAL AREAS FOR THE NOVATO WDN

Experiment	tank	h_{min}	h_{max}
Novato, broad	ta_1	147	153
	ta_2	123	156
	ta_3	156	162
Novato, medium	ta_1	148	152
	ta_2	123	156
	ta_3	157	161
Novato, narrow	ta_1	149	151
	ta_2	123	156
	ta_3	158	160

example WDN included with the EPANET simulator and covers an area of about 150 km². As such, we refer the reader to [5] for the concrete network structure. The WDN contains three tanks, two pumps that connect the network to a lake and a river and 91 junctions. EPANET also includes demand patterns for these junctions, one so-called default pattern that is used for most junctions together with a multiplier specific to each junction and four specific demand patterns for four concrete junctions. Unfortunately, EPANET does not provide efficiency data for the pumps¹, we only know that the pump at the lake is rather weak, whereas the one at the river is rather strong. For the experiments, we had to make up the efficiency for different pump speeds and the peak efficiency for the lake is a little below a flow of 5000 gallons per minute, whereas the river pump reaches its peak efficiency at around 12000 gallons per minute.

For using the IDC mechanism by the agents, each tank agent uses a value h_{min} as operational minimum head and a value h_{max} as operational maximum head with the goal to keep the actual head of the tank between those two values. In this experimental series, we looked at the effect of choosing these two values on the performance of the two variants of IDC. Table II shows the three different operational areas for the three tanks (in feet). Note that we did not change the values for ta_2 , since this tank is used in the experiments only to handle emergency head values within the network as indicated by the very large operational area.

As Table I shows, even the small changes on the operational area can cause serious efficiency problems for GM. The smaller the operational area the more requests are sent out by the tank agents and the more potential for inefficiency in fulfilling these requests exists. CGM is clearly better at this, although the narrow operational area comes with a slightly higher cost.

In general, both variants of the IDC approach can handle the broad variety of scenarios that we tested, without creating any serious problems like tanks running empty or not having enough head in the WDN to meet the demands. But the quality of the performance of GM and CGM can be quite different. While for a few scenarios the two variants

¹This is the reason why we cannot compare our results to real solutions.

perform at the same quality, most scenarios show that the added coordination by additional information in case of the CGM variant pays off in achieving lower costs for transporting the required water. Due to the dynamic nature of this optimization problem (naturally, the agents were not given the demand patterns used in the experiments) we can not expect any system to consistently produce optimal solutions, so that also CGM is not guaranteed to always produce good solutions, but it is clearly preferable over GM.

V. RELATED WORK

Controlling and optimizing the operation of resource distribution networks by self-organizing emergent MAS solutions is not a novel approach, of course. For all types of resources corresponding solution approaches can be found. For instance, in [12] a MAS approach has been presented for the control of distributed energy (electricity) management. However, as indicated in the introduction, because the problem constraints in an electricity network are quite different compared to the constraints present in a WDN, such approaches cannot be easily transferred. Furthermore, in this approach a demand-side management technology has been used, i.e. the customer interfaces were equipped with agents in order to manage the electricity consumption at the point of use, which generates high capital expenditures.

But also for controlling and optimizing the operation of WDNs several approaches already exist. Most of them, however, are centralized solution approaches that assume well-known water demand profiles, which turns the dynamic problem to solve in the real world into a static one. For instance, the approach presented in [13] uses simulated annealing to find the global maximum in all the permutations of the interactions between pumps. To keep this set of permutations limited, they assume to have only pumps with fixed speeds that can either be turned on or off. Furthermore, changes in pump operation are only able to be applied in steps of one hour. Similarly, the approaches presented in [14] and [15] use various non-linear programming techniques for optimizing pumping cost in WDNs. In [16], genetic algorithms are used for optimizing the operation of a WDN. Apparently, in all of these approaches the operation of a certain pump is centrally pre-calculated, which makes these approaches unable to react to unpredictable, dynamic events.

In this connection it should be noted that there exist also general approaches that use multi-agent systems to perform such pre-calculated optimizations, which, however, should not be confused with using multi-agent systems in the way the IDC mechanism does. For example, [17] presents an approach for optimizing the operations of gas pipelines using several agents, each performing a search for the optimal operation parameters for the various components of a given gas pipeline network for given demands. The agents cooperate by exchanging solution candidates, which results in some kind of emergent behavior of the whole

system. But this behavior is emergent with regard to the search, not emergent with regard to the produced behavior of the network.

However, there also exist centralized approaches for dynamic problem solving in WDNs, mainly based on some kind of demand prediction. For instance, in [18] an approach is presented for an adaptive optimization system for dynamic WDN operation, based on the combined use of an artificial neural network for predicting the consequences of different pump settings and a genetic algorithm for optimization. The system receives its data via an online monitoring system and is designed to assist WDN operators to select optimal operating control settings. Similarly, in [19], fuzzy inductive reasoning is used for demand prediction. However, these approaches are not able to react appropriately to unpredictable, dynamic events as well. A sole exception is made by the approach presented in [20], which uses a numerical optimization model for the real-time control and optimization of the operation of a WDN. Even though the approach is able to react on dynamic events very quickly, it is still centralized and hence does not provide the required scalability and robustness.

The only decentralized MAS approach for controlling and optimizing the operation of a WDN that exists to our knowledge so far is presented in [21] and evaluated based on direct agent communication in [22]. The proposed MAS similarly is in charge of optimizing the pumping costs in a WDN. However, this approach only considers a simple WDN with one tank and one pump. Moreover, the water demand profiles by customers were again supposed to be well-known and predictable, what turned the problem to solve again into a static one.

The IDC mechanism presented in this paper also differs from existing DCMs based on IBC, such as pheromone-based coordination [8] or pollination-inspired coordination [11], because the local behavior of the agents as well as the used types of infochemicals differ significantly from these existing approaches. However, all of these approaches instantiate the DIC model.

VI. CONCLUSIONS

In this paper, we presented a self-organizing emergent MAS approach designed for the decentralized real-time control of WDNs. The approach is inspired by the indirect defense phenomenon in nature and is an instance of the DIC model. In the experiments, we have compared two concrete variants of the approach, one solely based on a greedy selection of the nearest pump, capable of fulfilling requests from tanks, and one based on including additionally anticipated costs in the information conveyed by the infochemicals.

The experiments showed that both variants fulfill the main objective of a WDN, i.e. servicing customer demands. But they also showed that the additional information yielding a

better coordination of the second variant often results in a more cost efficient control of the operations of the network.

There are two major directions for possible future work, namely improvements of the approach to increase the efficiency of the demand fulfillment and extensions of the approach to deal with other conditions around the delivery of water. While the dynamic nature of the control problem for WDNs prevents us from developing any control that will always produce optimal solutions (that would require to be able to look into the future to know the exact demand before it is put onto the network), there are definitely periods where the demands are stable and/or predictable. For such periods, systems like the one of [16] can be used to find optimal solutions. Using the basic idea of [23] to combine the approach with optimal solutions for recurring demand patterns might further result in better cost efficiency.

We have concentrated in this work on achieving energy cost efficiency. But there are some other considerations in the operation of a WDN, such as water quality. For example, if water stays too long in tanks this may endanger water quality and might require additional measures. Extending the IDC approach to either also control these measures or to avoid having to take them is also future work we plan to look into.

REFERENCES

- [1] L. Mays, *Water Distribution System Handbook*. McGraw-Hill Professional, 1999.
- [2] M. Dicke, R. M. van Poeckea, and J. G. de Boera, "Inducible Indirect Defence of Plants: From Mechanisms to Ecological Functions," *Basic and Applied Ecology*, vol. 4, no. 1, pp. 27–42, 2003.
- [3] M. Dicke and M. W. Sabelis, "Infochemical Terminology: Based on Cost-Benefit Analysis Rather than Origin of Compounds?" *Functional Ecology*, vol. 2, no. 2, pp. 131–139, 1988.
- [4] H. Kasinger, B. Bauer, and J. Denzinger, "Design Pattern for Self-Organizing Emergent Systems Based on Digital Infochemicals," in *Proc. of EASe 2009*. IEEE, 2009, pp. 45–55.
- [5] EPA, "EPANET 2.0," <http://www.epa.gov/nrmrl/wswrd/dw/epanet.html>, 2008.
- [6] P. Cooper and G. Tchobanoglous, "Performance of Centrifugal Pumps," in *Pumpin Station Design*, 3rd ed. Elsevier, 2008.
- [7] T. Lewis, "The elements and frontiers of insect communication," in *Insect Communication*, T. Lewis, Ed. Academic Press, 1984, pp. 1–27.
- [8] S. Brückner, "Return from the Ant - Synthetic Ecosystems for Manufacturing Control," PhD thesis, Humboldt-Universität, Berlin, 2000.
- [9] P. Valckenaers, K. Hadeli, B. S. Germain, P. Verstraete, and H. V. Brussel, "MAS coordination and control based on stigmergy," *Computers in Industry*, vol. 58, no. 7, pp. 621–629, 2007.
- [10] H. Kasinger, J. Denzinger, and B. Bauer, "Digital Semiochemical Coordination," *Communications of SIWN*, vol. 4, pp. 133–139, June 2008.
- [11] H. Kasinger, B. Bauer, and J. Denzinger, "The Meaning of Semiochemicals to the Design of Self-Organizing Systems," in *Proceedings of SASO 2008*. IEEE, 2008, pp. 139–148.
- [12] A. Zeman, M. Prokopenko, Y. Guo, and R. Li, "Adaptive Control of Distributed Energy Management: A Comparative Study," in *Proceedings of SASO 2008*. IEEE, 2008, pp. 84–93.
- [13] F. E. Goldman and L. W. Mays, "Water Distribution System Operation: Application of Simulated Annealing," in *Water Resource Systems Management Tools*. McGraw-Hill Professional, 2004.
- [14] J. Burgschweiger, B. Gnädig, and M. C. Steinbach, "Optimization Models for Operative Planning in Drinking Water Networks," *Optimization and Engineering*, vol. 10, no. 1, pp. 43–73, 2009.
- [15] —, "Nonlinear programming techniques for operative planning in large drinking water networks," *Open Applied Mathematics Journal*, vol. 3, pp. 14–28, 2009.
- [16] P. F. Boulous, Z. Wu, C. H. Orr, M. Moore, P. Hsiung, and D. Thomas, "Optimal Pump Operation of Water Distribution Systems Using Genetic Algorithms," http://www.rbfcconsulting.com/papers/genetic_algo.pdf.
- [17] T. E. Mora, A. B. Sesay, J. Denzinger, H. Golshan, G. Poissant, and C. Konecnik, "Fuel Optimization Using Biologically-inspired Computational Models," in *Proceedings of IPC 2008*, 2008.
- [18] Z. F. Rao, J. Wicks, and S. West, "Optimising Water Supply and Distribution Operations," *Proceedings of the ICE - Water Management*, vol. 160, no. 2, pp. 95–101, 2007.
- [19] G. Cembrano, G. Wells, J. Quevedo, R. Perez, and R. Arge-laguet, "Optimal Control of a Water Distribution Network in a Supervisory Control System," *Control Engineering Practice*, vol. 8, no. 10, pp. 1177–1188, 2000.
- [20] C. Hähnlein, "Numerische Modellierung zur Betriebsoptimierung von Wasserverteilnetzen," PhD Thesis, TU Darmstadt, 2008, (in German).
- [21] L. Giannetti, F. P. Maturana, and F. M. Discenzo, "Agent-Based Control of a Municipal Water System," in *Multi-Agent Systems and Applications IV*, ser. LNCS. Springer, 2005, vol. 3690, pp. 500–510.
- [22] F. Maturana, R. Kotina, R. Staron, P. Tichý, and P. Vrba, "Agent-based Water/Waste Water Control System Architecture," in *IADIS International Conference Applied Computing*, 2006, pp. 661–666.
- [23] J.-P. Steghöfer, J. Denzinger, H. Kasinger, and B. Bauer, "Improving the Efficiency of Self-Organizing Emergent Systems by an Advisor," in *Proc. of EASe 2010*. IEEE, 2010, pp. 63–72.