On the Optimization of Collaborative Kerbside Waste Collection

ALESSANDRO PELLEGRINI
Sapienza, University of Rome
DIAG
Via Ariosto 25 - 00185 - Roma
Italy
pellegrini@dis.uniroma1.it

PIERANGELO DI SANZO
Sapienza, University of Rome
DIAG
Via Ariosto 25 - 00185 - Roma
Italy
disanzo@dis.uniroma1.it

Abstract: In this paper, we target collaborative kerbside collection from a planning and real-time monitoring point of view. This is a non-trivial problem, where several vehicles are set on streets to finish a task—the collection of all waste—by a certain maximum amount of time. While deciding upon a collaborative strategy is a well-studied and complex problem by itself, we focus as well on *re-planning*, whenever live data collected by the vehicles suggest that the current scenario has deviated from the provisional plan, due to changes in external environmental factors. To this end, we propose a global mission-management architecture, which tries to optimize at once the time required to finish the waste collection, the distance traveled by the vehicles, the amount of fuel burnt (accounting as well for idle time at collection points), and the impact of pollutants emissions.

Key–Words: Kerbside Collection, Optimization, Re-Routing, Real-Time, Deep Learning, Scheduling.

1 Introduction

The problem of the management of urban solid waste, which has always been the source of health and public order problems [29], is far from being completely solved. Differentiated causes, such as politics, bad management, illicit interests, or misinformation prevent an intelligent collection and disposal of urban waste as widespread and as efficient as it could be nowadays.

Kerbside collection is an even less trivial problem. Indeed, several different factors can significantly affect the timeliness and efficiency according to which such a system can be effective or not. Among them: i) the *timeliness* according to which waste is collected; ii) the time waste is left on public streets; iii) the *accuracy* of people to dispose waste in proper bags/bins.

In this paper, we propose a global architecture to address the three points above, in the context of Italian kerbside waste collection [3, 21], although our proposal can be easily adapted to other cities/scenarios. The ultimate goal is to devise a collection systems which is at the same time efficient and cost-effective.

To cope with points i) and ii) above, we note that they strongly depend on the efficient management of the vehicles used for collection. Therefore, our architecture relies on an information system to monitor in real time various territorial and environmental elements. This is done by means of vehicles equipped with data capture devices (e.g., cameras, scanners, electronic scales), positioning systems (based on GPS technology), and assisted-driving devices (i.e., navigators). Anyhow, external elements might interfere with the optimality of a certain collection plan, such as traffic and exhausted capacity of vehicles. Therefore, we add to our waste collection architecture the capability to monitor in real-time the position, the state of each vehicle—in terms of current available allowance, current distance from collection points—so as to globally exploit this information to enforce possible *re-routing* and *re-assignment* of tasks, considering all vehicles as a collaborative system, and accounting as well for current traffic conditions. We note that, by itself, this is a non-trivial problem to solve: in fact, finding an optimal routing strategies for vehicles sent out to collect waste is already a complex task, similar to the *traveling salesman problem* [11], which has been shown to be NP-hard [14]. Moreover, we add to the problem's complexity the fact that capacity allocation [15] is taken into account, in a collaborative fashion.

Similarly, we consider the possibility that during its operation, some vehicle might be involved in an accident, or the vehicle itself might break. Each vehicle is therefore equipped with a feedback system that allows the driver to interact with the monitoring system, in order to promptly notify unexpected events, and allow for an immediate replan. This allows to exploit the remainder of the available fleet to take care of the work which the unavailable vehicle cannot complete.

In addition, all the data acquired by vehicles, such as daily waste production, is used for long-term data-mining in order to forecast the possible amount of waste that each vehicle could find along its route in the collection process, and therefore prepare an initial collection plan as accurate as possible. This allows to reduce the likelihood that a vehicle gets filled during the collection, thus demanding for a re-plan by the monitoring system.

Regarding the aforementioned point iii), we note that in order to increase people's accuracy to dispose waste, a form of *incentive* could be put in place. While deciding upon the nature (and possibly the amount) of the incentives is something more related to the government and management of the waste collection network/city, we propose a technical solution to determine the amount of waste actually separated by single people or aggregated groups of people (e.g., apartment blocks).

Overall, this paper presents an architecture which jointly exhibits the following three innovative aspects:

1. vehicles used to collect waste are equipped with sensors and network elements which allow continuous monitoring and strategy (re-)planning in a *cooperative way* and according to a *global monitoring plan*. This allows to minimize the time required to finish the overall collection of waste, and reducing the length of the path traveled by the vehicles themselves. The minimization is done accounting as well for *fuel consumption* and therefore for pollutants emission;

2. Travel time, and amount of waste collected at each collection point is stored in a history database. This information is exploited through machine-learning and deep-learning techniques to build, at the beginning of each collection day, a provisional collection plan. This is a plan which is expected to be as accurate as possible with a high likelihood. In this way, the probability that the system has to recompute a new plan, due to some errors in the initial plan, are significantly reduced.

3. The collected waste is categorized, by means of optical scanners, into non-separated and separated (per typology) waste. This allows to build "customer merit tables", which are a numerical representation of the goodness of the separation process at each collection point. This information can be used to determine a reward for people (or groups of people), and can be as well complemented with certification of accuracy data at the disposal points, which is anyhow out of the scope of this paper.

Overall, a robust system for territorial governance needs to cope with on-line (re-)evaluation of monitoring activities, while coping with both timing/spatial and resource constrains. This is exactly the objective of this paper, where optimization and dynamic re-evaluation are core aspects underlying design/development activities.

We complement our solution with an experimental assessment, carried out by means of simulation. In particular, we exploit a micro-scale discrete-event simulation model which simulates vehicles moving in the city and the trucks set out to collect the waste. Thanks to the introduction of traffic conditions and random accidents on the vehicles, we show how our collaborative re-routing strategy is able to reduce the time required to collect all waste at all collection points, and the distance traveled by each vehicle, with respect to having them all stick to the initial mission plan.

The remainder of this paper is structured as follows. In section 2 we discuss related work. Section 3 presents the application scenario for our proposal in more details. Our automated kerbside collection monitoring architecture is presented in Section 4. Section 5 presents experimental data.

2 Related Work

As mentioned, waste collection, and more specifically collaborative kerbside collection, is a non-trivial problem which touches several research aspects and fields, from scheduling to capacity allocation, from real-time monitoring to graph visiting in order to determine the optimal collection path.

In the context of *graph visiting*, several works can be used in scenarios similar to the one we target. For example, in [11], a variant to the traditional Traveling Salesman Problem (TSP) is proposed, namely the *TSP with profits*. This is a generalization of TSP, where it is not necessary to visit all vertices. A profit is associated with each vertex. The overall goal is the simultaneous optimization of the collected profit and the travel costs. Our problem shares with TSP with profits the idea that each vertex, namely a collection point, has some profit, which is a function of the amount of waste to be collected, and possibly the time the waste has been left there. Nevertheless, in our approach, we explicitly set the number of vertices that cannot be visited to zero, as a successful collection is able to reach all the waste bins.

The work in [16], on the other hand, explicitly tackles solid waste collection in urban areas, explicitly accounting for routes and the available vehicle fleet. Nevertheless, the solution the work proposes is based on an operational problem, while we complement a nearly-optimal solution with live data to drive the collection in real-time.

In [7], the authors investigate a way to assign stations to vehicles so that constraints are satisfied, and the mileage covered by the fleet is a minimum. While we keep this ability, we exploit communication technology and available computing power to repeatedly recompute the strategy, so that the actions taken by the fleet are resilient to environmental changes.

Several works [22, 26, 8, 9] have studied a problem similar in spirit, namely the vehicle and Trailer Routing Problem (TTRP), which shares several goals and constraints to kerbside waste collection. Most of the proposed solution consider only the distance traveled and time spent in the loading/unloading. On the other hand, we consider as well the expected waste to be found at collection points, thus making our initial plans more reliable. The work in [26] takes into account as well the amount of fuel burned during the detainment at specific places, which we further back with machine learning-based data, to make our estimate more plausible.

In the context of dynamic capacity acquisition and assignment, several works have tried to find optimal solution using different techniques. Beyond the more classical stochastic integer programming [1], the work in [19] relies on discrete-event simulation (DES). We rely as well on the DES paradigm, yet to perform an assessment of our multi-layered proposal, rather than as the core solver for the collection task. In [27], on the other hand, the authors focus on the situation when capacities are uncertain. We consider this latter aspect as orthogonal to our proposal. In fact, uncertainty can be introduced in our system in the form of parameterization of the configuration, and it is therefore not the central goal of our proposal.

Similarly, the works in [6, 5, 10] present forecasting models related to waste generation in urban areas when the amount of historical data is very reduced and/or there is a non-minimal lack of sampling. These works rely on grey fuzzy dynamic modeling [6], on analytic models [10], or on time series [5]. Again, we consider all these approaches orthogonal to our proposal, and they can be effectively plugged into our architecture to improve the quality of provisional plans, until the vehicles have collected enough information during the collection missions, which allow data-mining processes to be more accurate in their prediction.

As for the reward system that can be enforced using our architecture, several works [17, 24, 23] have studied the reasons behind an unsuccessful kerbside collection strategy. Among the various reasons, these papers highlight the fact that recycling waste bins are often full and that there are no incentives in participation [17], a lack of social interaction among neighborhoods [24], and the lack of educational and promotional campaigns [23]. We believe that the reward-oriented architecture that can be put in place by relying on our proposal, can significantly address all the discussed reasons, and therefore make kerbside collection more effective.

3 Application Scenarios and System Model

We assume a complex differentiated waste management problem which can cover a variety of real scenarios. We note that different scenarios may arise because of the diversity of the urban contexs, as well as of the different tools that are available to carry out tasks associated with waste collection. Since we consider a complex (more general) problem, the solution that we devise in this paper has a wide applicability. Indeed, it can be used also for any scenario arising from the simplification of one or more of the constraints of our problem.

A list of elements that may lead to different scenarios includes:

- Collected materials, which can be different in the number and the type.

- Collection points, where material may be collected at any time or only during restricted time windows.

- Dumps, which can be different in number and in type of material that can be dumped.

- Fleet of vehicles, which can be of different nature, each one having different weight capacity, volume and/or speed.
- Depots, where different vehicles can start/end a tour.
- Urban streets, which can have different traveling speed, as well as different width, so that some vehicles may be prevented to pass through some street.

In addition to the variety of scenarios arising from the elements we discussed above, we also assume that various events may occur while vehicles are carring out a collection tour. These events include:

- Some vehicle may reach the maximum volume/weight capacity, or may break, before ending a tour.
- Some street may become no longer viable, thus some vehicle may no longer be able to complete a tour according to the assigned schedule.
- Traffic adn/or other events may lead some vehicle to delay execution of collection tasks, thus a vehicle may no longer be able to complete a tour within the schedule time.

To cope with such a complex and general scenario, the system architecture that we propose relies on two main capabilities:

a. Generating an optimal provisional routing schedule for vehicles.

b. Performing real-time monitoring of vehicle tours, and, if necessary, promptly generating an updated routing schedule that has to be sent to vehicle drivers, even in the middle of their tour.

We target a scenario with multiple trucks, and a single centralized *mission control system*. This could be as well replicated in order to increase availability and dependability of the service, but this is out of the scope of this paper. In fact, we are more interested in managing the available resources (the trucks) in a collaborative way in real time.

The mission control system collects data coming from each vehicle, which is equipped with various devices. Some of them are used to collect data about the evolution of the mission, in terms, e.g., of residual available capacity and volume/weight of collected waste per each collection point. This latter

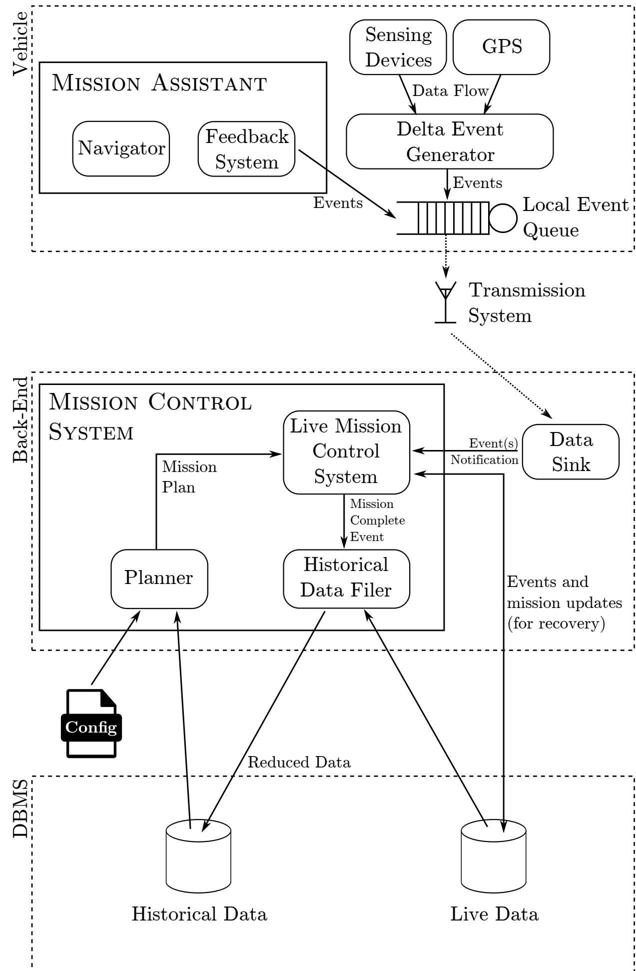

Figure 1: Architecture Organization.

point is achieved in a differentiated way, depending on the type of vehicle. In particular, if the vehicle is set towards the collection of one specific type of waste, the amount of collected waste is measured electronically (e.g., by using electronic scales installed on the vehicles) and directly associated with the proper typology—this works as well in the case of unsorted waste collection. On the other hand, if the vehicle can collect multiple typologies of waste[1], a scanner can be used by the operators to associate the waste with the proper category. Still, the amount of waste is electronically measured.

All the configuration parameters and the constraint to the waste collection problem can be specified by using a configuration file.

[1]We note that some recycling facilities might be able to process different types of waste at once.

4 Automated Kerbside Collection Monitoring Architecture

We now illustrate the organization or our architecture to cope with points **a** and **b** illustrated in the previous section. The overall organization of our architecture is depicted in Figure 1.

Each vehicle is equipped with a *mission assistant device*. This is composed of a navigation system, which directs the driver according to the current *waste collection plan*, and a feedback system which allows the driver to explicitly interact with the central mission control. At the same time, the sensing devices autonomously generate a continuous data flow, which is intercepted by the on-board *Delta Event Generator*.

This is a component which locally (on each truck) stores the data related to a (small) past time window. Only if the variation from previous values exceeds some thresholds (which are configured by the mission control system and can be specified for each measured dimension) the component generates a *variation event*. This event is enqueued into a *local event queue*, which serves as a local buffer to keep events in case no network coverage (e.g., mobile broadband network) is available during a certain portion of the mission.

Each truck therefore sends only delta events to the mission control system. A component, named *data sink*, allows to manage as well the possibility that a batch of events is received from a single vehicle, and combines them so that the processing time can be reduced. The data associated with these events are then notified to the *live mission control system*. This component creates, before processing, a copy on a *live data* database, mostly to allow a safe fail-restart scenario of the mission control system, which therefore does not lose data on the current mission.

Once the waste collection mission is completed, the live mission control system triggers the *historical data filer*, which gathers all the data related to the current mission, and reduces them (by stripping non-necessary information) into the *historical data* database. This data can be used to extract the "customer merit tables", by determining the amount of sorted waste by each customer or group of customers.

The data kept by these two databases, along with the flow of events generated by trucks, are used to setup missions, as described in the following.

4.1 Optimal Routing Schedule

As discussed in Section 2, the literature offers various algorithms to solve a vehicle routing problems and its many variants. Thus, given a specific scenario, the algorithm can be selected among the ones whose underlying problem formulation copes with all requirements of the scenario. Since a wide variety of scenarios and algorithms exists, we do not enter into details of the selection of the algorithms for each possible scenario. However, in this paper we consider an example algorithm that cope with all requirements of our complex scenario. In the case of simpler scenarios, one could use our example algorithm by relaxing some assumption, or could select some other literature algorithm which is based on a simpler problem formulation.

The scenario which we target is the following one. The fleet includes several and different vehicles that can be set towards collection points (bins). They can start their journey from different depots. When a vehicle leaves its depot, it is already assigned a provisional collection schedule, the construction of which will be discussed after in this section. Each vehicle can be associated with a specific type of waste—each collection point could be reached by multiple vehicles. The provisional plan takes into account the vehicles' capacity/maximum weight, type of waste, expected amount of waste to be found at collection points, and distance to be traveled.

The algorithm that we consider is the one presented in [4]. The objective function of this algorithm aims at minimizing the spatial and temporal costs of the routing schedule. The total cost to minimize is calculated on the basis of input variables that specify a fixed cost for each vehicle, a unit-distance running cost an hourly driver wage rate. Further, the problem formulation includes 22 constraints that ensure that the calculated routing schedule adheres to all requirements of our complex scenario.

When solving the linear program in [4], we set the value of some parameters depending on historical data collected during previous missions. As an example, we set the *service time* of a certain collection point i s_i proportional to a linear combination[2] of the amount of waste that was observed in the same weekdays, in the same week of the months, and in the same month in the past. We similarly set the *travel time* of the edges

[2] By empirical experience, we suggest that the multiplicative factors of the linear combination can be learnt via reinforcement learning [25].

$t_i j$, accounting for the expected amount of traffic. In case of the startup of the system, we rely on grey fuzzy dynamic modeling [6] to setup an acceptably-reliable estimate from a few samples. We note that by measuring and considering the service time s_i, we are actually trying to plan the visits of the collection points in a way such that as well the fuel consumption is taken into account, additionally when the truck is stopped at the collection point.

Since the vehicle routing problem is NP-hard, the proposed solution technique uses an heuristic approach based on local search. The experimental study to evaluate the solution technique has been conducted by simulating a collection scenario with more than 2800 containers located at 820 collection points, positioned over an are of approximately 2000 km^2. All details about the experiment results can be found in the original paper.

Although complex, we can dedicate enough resources and enough time to find a solution to this problem. In fact, this is not a solution to be found in real time, as this is only the provisional plan which will be refined later, in case the scenario deviates to much from the plan. Indeed, this provisional plan could be computed as well while the previous collection mission is running.

4.2 Real-time monitoring and re-scheduling

Obviously, once selected the most appropriated algorithm to solve the specific vehicle routing problem, the calculated routing schedules are suitable in the case of a *static* scenario, i.e. when no events that may alter the initial conditions (based on which the routing schedule has been calculated) occur until all vehicles complete their tour. To cope with unforeseen events, our system architecture relies on a real-time monitoring system that verifies whether or not the tasks carried out by vehicles since the begin of their tour, as well as information received by vehicles while executing the assigned task, may lead the current routing schedule to fail. For example, a routing schedule can fail if some vehicle reaches the maximum weight capacity before reaching all assigned collection points, excessively delays the execution of assigned tasks, breaks before completing the tour, etc.

As mentioned, the overall goal is to minimize costs associated waste collection optimization. At the same time, we strive for automatic re-plan, in case live data collected by the vehicles suggest that the current scenario has deviated from the provisional plan, due to changes in external environmental factors.

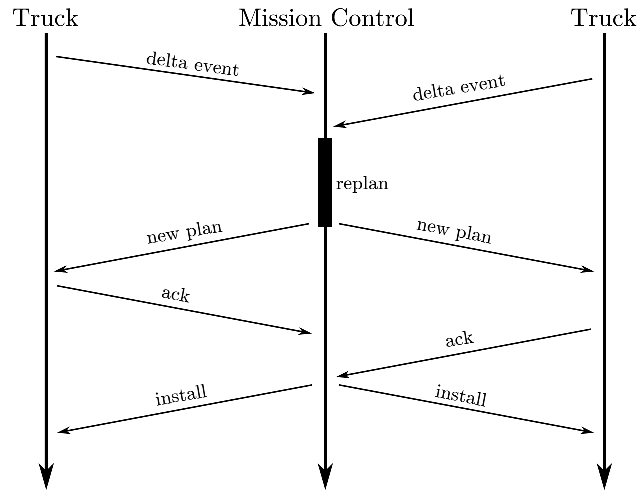

Figure 2: Messages exchanged to install a new plan.

The data and the events collected by the vehicles drive the re-evaluation of the mission by the back-end. In particular, the *live mission control system* is configured using a set of threshold from the configuration file which determine when a new plan should be computed. As an example, if a truck is late at reaching a collection point due to, e.g., traffic conditions, a new recomputation is triggered. Any explicit event, like a truck outage, immediately calls for a replan.

The replan procedure is carried out by removing from the set of vertices of the graph all the vertices that have already been visited, plus each vertex that is currently being reached by every truck. This latter vertices are removed in order to avoid scenarios in which the driver does not notice the plan change and therefore an additional replan is to be triggered (to avoid multiple trucks reach the same point), and to avoid degenerate scenarios in which a truck which is very close to a destination vertex has to change direction, which could possibly upset the driver towards the system.

The new plan is computed by re-solving the linear program discussed in [4]. Once the new plan is computed, the live mission control system transmits to all mission assistants at each truck the new plan. In any case, the new plan is not immediately installed. In fact, there could be the case that some truck is out of range of the mobile network. In this case, other degenerate scenarios might arise, having multiple trucks reach the same point.

Once a new plan is received, a truck sends back an acknowledgement message to the live mission control system. The system keeps track of which truck has responded, and once a response from each truck is received, the live mission control system sends to all trucks an *install* control message. At this point, each truck installs the new plan. This pattern is depicted in Figure 2.

After a certain amount of time, if no install message is received, the trucks discard the new plan, and so does the live mission control system. This means that some truck cannot be reached by mobile connection. We therefore seek for a sub-optimal solution in the following way. Since the live mission control system knows what are the trucks that did not send an acknowledgement message, a new plan discarding it can be computed.

In particular, the system searches for a new solution to the linear program in[4] by removing first the non-responding trucks and all the vertices that are currently present in each non-responding truck's paths. In this way, we avoid at all that multiple trucks reach the same collection point. Although non-optimal, this replan exploits as well the non-responding trucks to collaborate, as they are anyhow reaching their previously-assigned collection points. Once new delta events are received by all the currently non-responding trucks, a new plan with all trucks can be recomputed.

To avoid network flooding, after a new plan has been installed by trucks, the live mission control system waits at least for a specified amount of time before triggering a new replan.

5 Experimental Results

In order to evaluate the effects of the proposed architecture, we have relied on a Discrete-Event Simulation model, called *traffic* [28], which allows to simulate at a very fine grain the evolution of traffic conditions, according to statistic distributions of traffic at road segments and junctions. The model is able to account for randomized accidents, generated according to some probability distribution. It takes into account as well the expected time before the street is freed by an accident, in order to generate traffic increase in the roads located around an accident point.

We have run the traffic simulation model relying on a Parallel Discrete-Event Simulation engine, namely ROOT-Sim [18, 28], in order to keep tractable the problem. ROOT-Sim is a speculative simulation model, developed according to the Time Warp synchronization protocol [13].

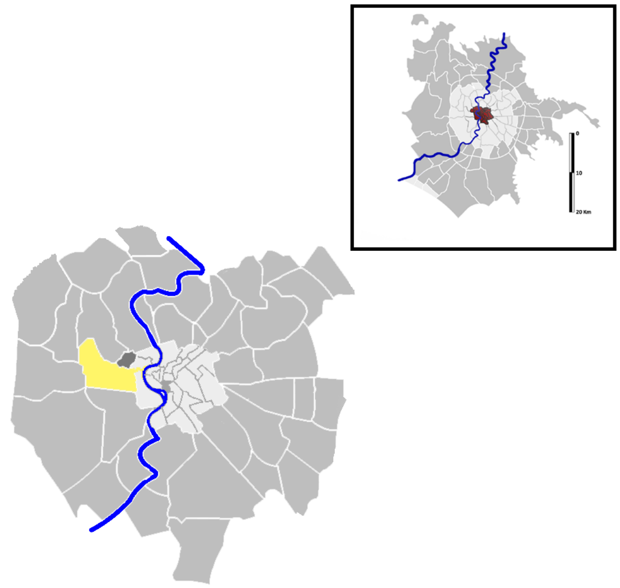

Figure 3: Part of the city which has been micro-simulated (in yellow).

As the simulation scenario, we have simulated a neighborhood of Rome, the Italian capital city—called *Quartiere Aurelio*, which is depicted in Figure 3—which has around 1000 waste collection points, differentiated into paper, plastic and metal, wet waste, glass, clothes, and unsorted. Although this is not an extremely large neighborhood (it's only 4.8 Km^2 wide), it's highly populated (around 9,000 people/Km^2), and it serves as the access to the center of the city where a high number of offices is present for a large number of vehicles (a bit less than 1 million vehicles are estimated to enter and leave the city each day [12]). The effect is that the traffic in this area is so high that it can take up to one hour to travel just one kilometer. In this scenario, the waste collection is carried out by an average of 15 trucks [2]. Each collection point should be reached by different trucks, in order to collect the different materials.

We have configured the traffic model so that only major roads are simulated. In fact, by the topology of the portion of the city, these are the roads that are most commonly affected by high levels of traffic, and are the actual roads that are traveled by the waste-collection trucks. Similarly, we have configured the traffic-generation probability according to real statistics [20].

We have additionally varied the probability according to which a truck gets broken, drawn according

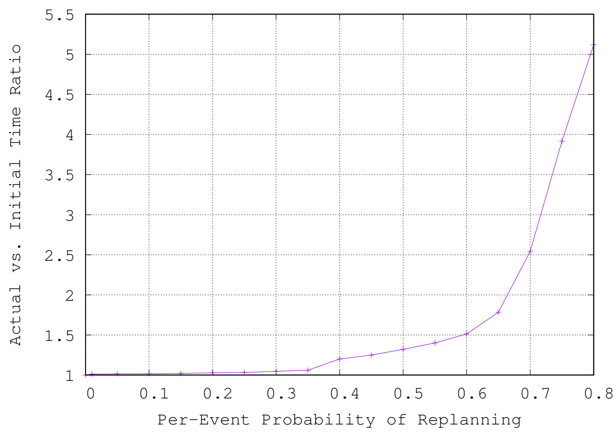

Figure 4: Mission Time when Replanning.

to a Gamma distribution. Overall, this has allowed us to vary the probability according to which a replan is triggered by the simulated vehicles after the execution of each discrete simulation event in the interval [0.01, 0.8]. As discussed, replans might be less accurate, or might be related to outages of trucks, and therefore this is an indication of how much the situation is deviating from the original plan.

In Figure 4, we report the variation of the ratio of actual mission plan vs. the probability of a replan. By the results, we can immediately deduce that, despite the highly-adverse environmental conditions, the system is able to find a suitable replan up to a probability value of around 65%. In particular, until that point, the actual vs. initially-planned time ratio is kept below 1.5, indicating an increase in the time spent (and therefore in the mission cost) of up to 45%.

Additionally, up to a probability value of 35% percent, the variation in the time ratio is so small that it can be considered negligible. This is an indicator of the importance and effectiveness of the collaborative monitoring architecture that we have proposed.

6 Conclusion

In this paper we have discussed an architecture to monitor in real time the waste collection process, carried out by a set of trucks moving in the city. We have illustrated how this organization can effectively help in coping with execution plans of the collection which can deviate from the initial forecast, due to external factors such as traffic, unexpected increase of the amount of waste found at collection points, or truck outages/accidents.

Our experimental evaluation, carried out by means of discrete event simulation, has shown that our proposal can significantly reduce the impact of unexpected/unpredictable events, affecting kerbside waste collection. Moreover, our results show that by relying on an architecture like the one we have discussed in this paper, it is possible to cope with a replan probability of up to 65%, with an increase in the cost of only 45%. Therefore, we conclude that it is fundamental, for waste collection systems of the future, to embed monitoring and replanning capabilities like the proposed ones, in order to be able to offer an efficient system at a controlled cost.

Acknowledgements: The authors are also working with Value Up S.r.l., an InResLab partner. This work is partially supported by the W4R project funded by OMB Roma S.r.l. with the support of MISE research funds.

References:

[1] S. Ahmed and R. Garcia. Dynamic capacity acquisition and assignment under uncertainty. *Annals of Operations Research*, 124(1-4):267–283, 2003.

[2] AMA S.p.A. I nuovi mezzi AMA per la pulizia della città. Technical report, 2009.

[3] N. Belosi. Gestione del Ciclo Integrato dei Rifiuti Urbani Lombardia-Veneto. Technical report, Ecoistituto di Faenza, 2007.

[4] K. Buhrkal, A. Larsen, and S. Ropke. The Waste Collection Vehicle Routing Problem with Time Windows in a City Logistics Context. *Procedia - Social and Behavioral Sciences*, 39:241–254, 2012.

[5] N. B. Chang and Y. T. Lin. An analysis of recycling impacts on solid waste generation by time series intervention modeling. *Resources, Conservation and Recycling*, 19(3):165–186, 1997.

[6] H. Chen and N.-B. Chang. Prediction analysis of solid waste generation based on grey fuzzy dynamic modeling. *Resources, Conservation and Recycling*, 29(1-2):1–18, 2000.

[7] G. B. Dantzing and J. H. Ramser. The Truck Dispatching Problem. *Management Science*, 6(1):80–91, 1959.

[8] U. Derigs, M. Pullmann, and U. Vogel. Truck and trailer routing - Problems, heuristics and computational experience. *Computers and Operations Research*, 40(2):536–546, 2013.

[9] R. Dondo and J. Cerdá. The heterogeneous vehicle routing and truck scheduling problem in a multi-door cross-dock system. *Computers and Chemical Engineering*, 76:42–62, 2015.

[10] B. Dyson and N.-B. Chang. Forecasting municipal solid waste generation in a fast-growing urban region with system dynamics modeling. *Waste Management*, 25(7):669–679, 2005.

[11] D. Feillet, P. Dejax, and M. Gendreau. Traveling Salesman Problems with Profits. *Transportation Science*, 39(2):188–205, 2005.

[12] ISTAT. Popolazione Roma 2001-2015. Technical report, 2015.

[13] D. R. Jefferson. Virtual Time. *ACM Transactions on Programming Languages and System*, 7(3):404–425, 1985.

[14] R. M. Karp. Reducibility among Combinatorial Problems. In *Complexity of Computer Computations*, pages 85–103. Springer US, Boston, MA, 1972.

[15] H. W. Kuhn. The Hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, 2(1-2):83–97, mar 1955.

[16] V. Maniezzo. Algorithms for large directed CARP instances: urban solid waste collection operational support. Technical report, 2004.

[17] S. McDonald and C. Oates. Reasons for non-participation in a kerbside recycling scheme. *Resources, Conservation and Recycling*, 39(4):369–385, 2003.

[18] A. Pellegrini. Hijacker: Efficient static software instrumentation with applications in high performance computing. In *Proceedings of the 2013 International Conference on High Performance Computing and Simulation*, HPCS, pages 650–655, Helsinki, Finland, 2013.

[19] M. E. H. Petering. Decision support for yard capacity, fleet composition, truck substitutability, and scalability issues at seaport container terminals. *Transportation Research Part E: Logistics and Transportation Review*, 47(1):85–103, 2011.

[20] Provincia di Roma. Flussi di traffico. Technical report, 2014.

[21] M. Ricci, A. Tornavacca, and C. Francia. Gestione integrata dei rifiuti urbani: analisi comparata dei sistemi di raccolta. Technical report, FederAmbiente, 2003.

[22] S. Scheuerer. A tabu search heuristic for the truck and trailer routing problem. *Computers and Operations Research*, 33(4):894–909, 2006.

[23] S. Shah and M. Bull. OpenMP. In *Proceedings of the 2006 ACM/IEEE conference on Supercomputing*, SC. ACM, 2006.

[24] P. J. Shaw. Nearest neighbour effects in kerbside household waste recycling. *Resources, Conservation and Recycling*, 52(5):775–784, 2008.

[25] R. S. Sutton and A. G. Barto. *Reinforcement Learning I: Introduction*. MIT Press, 1998.

[26] Y. Suzuki. A new truck-routing approach for reducing fuel consumption and pollutants emission. *Transportation Research Part D: Transport and Environment*, 16(1):73–77, 2011.

[27] B. Toktas, J. W. Yen, and Z. B. Zabinsky. Addressing capacity uncertainty in resource-constrained assignment problems. *Computers and Operations Research*, 33(3):724–745, 2006.

[28] R. Vitali, A. Pellegrini, and F. Quaglia. Towards symmetric multi-threaded optimistic simulation kernels. In *Proceedings of the 26th Workshop on Principles of Advanced and Distributed Simulation*, PADS, pages 211–220. IEEE Computer Society, jul 2012.

[29] D. Q. Zhang, S. K. Tan, and R. M. Gersberg. Municipal solid waste management in China: Status, problems and challenges, 2010.