Route Selection for Mobile Sensors with Checkpointing Constraints

Olga Saukh, David Hasenfratz, Abouzar Noori, Tamara Ulrich, and Lothar Thiele
Computer Engineering and Networks Laboratory
ETH Zurich, Switzerland
{saukh, hasenfratz, ulrich, thiele}@tik.ee.ethz.ch, anoori@student.ethz.ch

***Abstract*—The sensing range of a sensor is spatially limited. Thus, achieving a good coverage of a large area of interest requires installation of a huge number of sensors which is cost and labor intensive. For example, monitoring air pollution in a city needs a high density of measurement stations installed throughout streets and courtyards. An alternative is to install a smaller number of mobile stations which traverse the city. The public transport network builds a perfect backbone for this purpose as public transport vehicles follow fixed and regular mobility patterns. In this paper, we consider the problem of selecting a subnetwork of a city's public transport network to achieve a good coverage in the area. Since we are working with low-cost sensors which exhibit failures and drift over time, vehicles selected for sensor installation have to be in each other's vicinity from time to time to allow comparing sensor readings. We refer to such meeting points as checkpoints. Due to high computational complexity of the route selection problem, both with and without checkpointing support, we adapt an evolutionary algorithm solution and evaluate its output based on the tram network of Zurich.**

I. Introduction

Today's big cities suffer from high concentrations of traffic and industrial facilities that heavily impact ecological sustainability and quality of living in the area. Monitoring air pollution, limiting the amount of transit traffic, and emission reduction has been addressed at many levels, but mostly through legislative decisions and standards. Due to high cost, weight, and size of traditional air pollution measurement instruments, there are still no *precise* maps of air pollutant distribution in cities.

In the last several years low-cost gas sensors have become available on the market. We work on combining this technology with wireless sensor networks for air pollution monitoring applications. In particular, we install sensing stations on top of several public transport vehicles to be able to achieve better coverage than in case of static deployment of the same stations. Public transport networks, referred to as *timetable networks* [11] in the research community, form an attractive backbone for performing periodic measurements due to (1) a large number of spatially spread predefined routes, (2) a fixed timetable, and (3) a usually good reliability.

The first problem we face is how to choose a subnetwork of the public transport network to cover the city well given a route plan and a timetable. Additionally, since our measurement stations are mainly equipped with low-cost gas sensors, we demand that the final subset of selected vehicles allows comparing measured sensor values across different sensors, *i.e.*, we require that different sensors periodically take measurements at the same time and location. In this paper, the problem is referred to as *sensor checkpointing*.

Sensor checkpointing is an important property of a distributed sensing system and contributes to the system's fault tolerance. In particular, it allows recognizing sensor failures and provides necessary support for sensor calibration. Informally, a pair of measurements performed by two separate sensors makes a *checkpoint* if these measurements are taken in each other's temporal and spatial vicinity. This naturally implies the simultaneous presence of the corresponding mobile vehicles at the same place. Thus, throughout this work we talk about two vehicles making a checkpoint, since this is a necessary condition to obtain two measurements taken at the same point in time and space. Timetable networks are generally periodic and thus checkpoints also occur periodically when considering them on an infinite timeline.

Related approaches that solve coverage problem do not consider an additional checkpointing support as required by our scenario. The problem of efficient checkpoint design arises when solving route planning problems atop a timetable network. However, the focus here is to provide good coverage of the region rather than to find optimal paths in the timetable network.

When choosing a timetable subnetwork to provide both maximum coverage of the region and to fulfill sensor checkpointing requirements, the solution space is huge even for moderate-sized cities for the following reasons: (1) in an average-sized city, several hundreds to a few thousands of mobile vehicles compose a public transport network; and (2) transportation lines have different route lengths. Therefore, the problem must be considered for the time period equal to the least common multiple of the round trips of all mobile vehicles.

We investigate the problem of pre-deployment route selection based on air pollution monitoring scenario in the city of Zurich, Switzerland as part of the OpenSense project [1]. The long-term goal of OpenSense is (1) to raise community interest in air pollution and (2) to encourage public involvement in the measurement campaign using enhanced cell phones or pocket sensors [4]. To establish an initial coverage of the city, we deploy sensors on top of several public transport vehicles, such as buses and trams. By doing so, we hope to foster community interest and its active involvement in data gathering.

In this paper, we concentrate on the following two types of sensor checkpointing: *X-checkpointing* and *R-checkpointing* that we introduce in Sec. II. X-checkpointing requires that any two vehicles are able to compare their measurements (possibly in a multi-hop fashion). Given a set of reference stations, R-checkpointing ensures that each vehicle can compare its sensor readings to one of the reference stations. We investigate the conditions necessary for both types of checkpointing and propose in Sec. III an algorithm for solving the following pre-deployment optimization problem: *Select a subnetwork of a timetable network which consists of K mobile vehicles to maximize coverage of the area of interest and enable sensor checkpointing.* We consider the above problem under the assumption that a mobile vehicle is uniquely defined by its track and its timetable. The proposed algorithm is evaluated based on the tram network of the city of Zurich in Sec. IV. Related work in the area is summarized in Sec. V and Sec. VI concludes this paper and discusses future research directions.

II. Model & Problem Statement

This section introduces the required terminology and the model which are used throughout this work to formally state and solve the problem of pre-deployment route selection with sensor checkpointing support, *i.e.*, before the sensors are mounted on top of the vehicles.

A. Timetable Network

Let $\Omega \subset \mathbb{R}^3$ represent an area and a time period of interest. A timetable network $\mathcal{N} = (H, S, C)$ consists of a set of mobile vehicles H, a set of stations S, and a set of elementary connections C. An elementary connection is modeled as a 5-tuple $c = (h, s_1, s_2, t_1, t_2)$ which is interpreted as a vehicle $h \in H$ departing from station $s_1 \in S$ at time t_1 and arriving at the next immediate station $s_2 \in S$ at time t_2. Throughout this work we assume that in most cases mobile vehicles follow their designated tracks. Another assumption is that each vehicle usually follows its timetable with no delays or arrivals ahead of time. The same schedule repeats on a daily basis. All exceptions to these general rules are infrequent and, therefore, negligible. A *timetable subnetwork* $\mathcal{L} \subset \mathcal{N}$ is a timetable network *induced* by a subset of vehicles $H_{\mathcal{L}} \subset H$ of $\mathcal{N}$. The *size of a timetable network* is defined as the number of vehicles $|H_{\mathcal{L}}|$ the network comprises.

B. Area Coverage

Consider a mobile vehicle $h \in H$ with a sensing station installed on top of it. Measurements can be taken by h while it is moving and at the stops with no restrictions. A *measurement* $z \in \Omega$ consists of a location and a timestamp. We consider that a measurement is a *point measurement*, that is, it has no duration. If a measurement z is taken by a vehicle $h \in H$, we use the notation $z \in h$ to express that z belongs to the space-time movement curve of h. A measurement $z \in h$ is valid in a certain area and for a certain time within Ω. Let $w : \Omega \times \Omega \to [0, 1]$ denote the *validity* of a measurement at a point in its vicinity. For a point $x \in \Omega$, $w(z, x)$ represents the level of *coverage* at x provided by a measurement z. Naturally, w monotonically decreases with increasing distance from z. We define the validity of a measurement as a strictly monotonically decreasing symmetric function independent of the actual measurement position or measurement time:

1) $w(z, z) = 1$;
2) $w(z, z + x) = w(z, z - x)$ (point-symmetric);
3) $w(z, x) > w(z, y) \Leftrightarrow \|z - x\| < \|z - y\|$ (decreasing);
4) $w(z, x) = w(z - y, x - y)$ (shift independent).

where $x, y \in \Omega$.

Let a *density requirement function* $\rho : \Omega \to (0, 1]$ represent the measurement density demand in the area Ω. We use ρ to express the fact that some areas might require greater coverage than others. For a timetable subnetwork $\mathcal{L} \subseteq \mathcal{N}$, the *coverage* of Ω achieved by $\mathcal{L}$ is given by

$$C(\mathcal{L}) = \int_{\Omega} \rho(x) \max_{z \in h, \forall h \in H_{\mathcal{L}}} w(z, x)\, \mathrm{d}x \qquad (1)$$

In this paper we are interested in finding a timetable subnetwork $\mathcal{L}$ that maximize the area coverage.

C. Sensor Checkpointing

Let h_1 and h_2 be two mobile vehicles equipped with air quality measurement stations. Let two measurements $z_1 \in h_1$ and $z_2 \in h_2$ be performed by h_1 and h_2, respectively. Consider a point $p = \frac{z_1 + z_2}{2} \in \Omega$, which represents the middle point between z_1 and z_2. Due to point-symmetry of the validity function w, the *joint validity* of measurements z_1 and z_2 is achieved at p and equals to $w(z_1, p) = w(z_2, p)$. Let α express the *minimum level of joint validity* required in order to compare two measurements. A point p is called a *checkpoint* if $w(z_1, p) > \alpha$. Since timetables are regularly executed, there can be many sequential periodic checkpoints between two mobile vehicles operating on nearby lines. A *sequence of checkpoints* $P_{ij} = \{p\}$ is a time-ordered set of points $p \in \Omega$ in which pollution measurements performed by h_i and h_j can be compared.

Checkpoints are essential in distributed sensing systems, since they allow implementing mechanisms to detect failures of low-cost sensing hardware, identify sensor errors, and provide necessary support for automatic sensor calibration. Leveraging the properties of timetable networks that two vehicles come closer to each other at the stations than during the drive and there are no parallel lines with negligible distance in-between, we reasonably assume that all checkpoints occur at the stations. Compared to the notion of a *transfer* [11], introduced in timetable networks, checkpoints are direction-independent and have no transfer time.

D. Problem Statement

Having a sequence of checkpoints for each pair of mobile vehicles, it is possible to construct a checkpoint graph $\mathcal{G} = (H_{\mathcal{L}}, E_{\mathcal{L}})$ where the set of nodes corresponds to the set of mobile vehicles $H_{\mathcal{L}}$ equipped with sensing stations. There is an edge $e \in E_{\mathcal{L}}$ between any two vehicles $h_i, h_j \in H_{\mathcal{L}}$ if the sequence of checkpoints between them is not empty $P_{ij} \neq \emptyset$.

Connectivity of a checkpoint graph $\mathcal{G}$ can also take into account information on quality and frequency of individual checkpoints. If all sensing stations are of the same kind, *i.e.*, all sensors have the same precision, the quality of a checkpoint depends on: the spatial and temporal distance between the measurements which form a checkpoint and the frequency at which the checkpoint occurs. Consider f_p is the frequency of a checkpoint $p \in P_{ij}$. An edge $e \in E$ in the checkpoint graph exists if

$$\frac{\sum_{p \in P_{ij}, i \neq j} f_p \max_{z \in h_i} w(z,p)}{\sum_{p \in P_{ij}, i \neq j} f_p} > \beta \tag{2}$$

and represents the joint quality of all checkpoints between a pair of mobile vehicles h_i and h_j. The constant β denotes the *minimum level of joint synchronization* in terms of the product of frequency and validity of the checkpoints.

To enable sensor checkpointing we need to ensure that the checkpoint graph $\mathcal{G}$ is connected or k-vertex-connected, meaning that any two sensors can be compared over at least k vertex-independent paths. We refer to this type of checkpointing as X-checkpointing (*cross checkpointing*). In this paper we consider only 1-vertex-connectivity, although requesting k-vertex-connectivity would improve the system's resistance to traffic artifacts such as delays.

X-Checkpointing: *Given a timetable network $\mathcal{N}$. Choose a timetable subnetwork $\mathcal{L}$ of size K to ensure maximum coverage of the area of interest Ω under the condition that the checkpoint graph $\mathcal{G}$ is k-vertex-connected.*

The possibility to synchronize the sensors enabled by checkpointing allows detecting faulty sensors in a similar fashion as majority voting. The quality of checkpointing can be further improved if the sensors can regularly synchronize with a set of *reference stations* $\mathcal{R}$. Reference stations can be static or mobile and are usually capable of performing high-quality sensing. We assume that all reference stations reflect the ground truth and can be used to calibrate low-cost sensors from time to time over one or several hops. This problem is referred to as R-checkpointing (*reference checkpointing*).

R-Checkpointing: *Given a timetable network $\mathcal{N}$ and a set of reference nodes $\mathcal{R}$. Choose a subnetwork $\mathcal{L}$ of size K to ensure maximum coverage of the area of interest Ω under the condition that each node in the checkpoint graph $\mathcal{G}$ is k-vertex-connected to the set of reference stations $\mathcal{R}$.*

Many big cities have a very sparse network of highly precise fixed stations, which can be used as references. For example, in Zurich, Switzerland there is one station of the national air quality monitoring network NABEL[1] and four smaller stations part of the cantonal measurement network OstLuft[2]. The availability of this infrastructure allows to considerably improve the system's reliability by selecting a timetable subnetwork which has R-checkpointing property. In contrast, many smaller cities have no reference station installed. In this case, pairwise cross-tests among low-cost sensors are essential for being able to identify sensor faults and recalibration needs.

In the next section we describe our solution approach to the defined pre-deployment timetable subnetwork selection problem with sensor checkpointing.

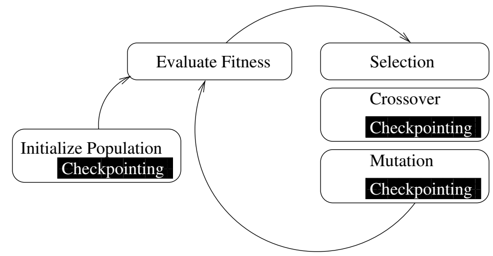

Fig. 1. Structure of the proposed evolutionary algorithm.

City	Network	Routes	Vehicles	Stops	Source
Zurich	tram	13	260	187	zvv.ch'11
Canton Zurich	bus	283	732	2'543	zvv.ch'11
Berlin	bus	149	1'300	2'634	bvg.de'11
Chicago	bus	152	2'000	12'000	wikipedia
NY City	bus	324	5'908	15'226	mta.info'09
Long Island	bus	48	389	n/a	mta.info'10

TABLE I
NUMBER OF ROUTES, VEHICLES, AND STOPS OF PUBLIC TRANSPORT NETWORKS IN DIFFERENT CITIES.

III. ROUTE SELECTION WITH EVOLUTIONARY ALGORITHM

The formulated problem of pre-deployment route selection involves high computational complexity even for moderate-sized cities. The brute force approach would require to go through $\binom{|H|}{K}$ combinations of transport vehicles to compute the optimum timetable subnetwork of size K. In Table I we present the relevant network characteristics of several cities world wide. In case of Zurich, the smallest among the listed cities, the solution space for $K = 10$ is approximately 250^8. Note, neither X- nor R-checkpointing constraints reduce the size of the worst case solution space of the problem. In the worst case all mobile vehicles follow the same track with different speeds and thus any selected subnetwork fulfills checkpointing constraints. However, finding the subnetwork which gives the best coverage requires consideration of all combinations.

For the reason above, we use an evolutionary algorithm to solve the problem. The working principle is schematically depicted in Fig. 1. In the remaining part of this section we describe the representation of a chromosome, the fitness function we use, the selection scheme, and the variation operators (crossover and mutation).

A *chromosome* is a timetable subnetwork of length K. We begin with an initial population composed of a random set of chromosomes. In case of X- and R-checkpointing we only consider chromosomes that fulfill the corresponding connectivity constraints, *i.e.*, we keep generating random subnetworks until we have the desired number of feasible chromosomes.

The *fitness* of a timetable subnetwork represents the coverage of the city Ω provided by the subnetwork. Precise

[1]www.bafu.admin.ch/luft/luftbelastung
[2]www.ostluft.ch

calculation of the coverage as defined by Eq. (1) involves discretization of the region with Monte Carlo sampling and computation of the level of coverage at each of the sampled points. Since the algorithm requires computing coverage for each chromosome in each iteration, we use a simpler coverage metric. We discretize time at each vehicle stop. At each discrete time point we compute a Voronoi diagram [3] in space only by using vehicle positions as a set of sites. The maximum distance between a generator and the points in its associated Voronoi cell describes the coverage at the corresponding discrete point in time. The time complexity of this computation is $O(K \log K)$. The approach is similar to the idea of the point distribution norm h as defined in [7]. Note that the point distribution norm is insensitive to the density requirement ρ and, therefore, can only be used for $\rho \equiv 1$. Finally, we sum up the obtained distances over all initially defined discrete time points. Good coverage corresponds to low values of the fitness function computed this way. The simplified coverage metric assumes only short term validity of measurements in time, which is reasonable for the measurement of air pollutants. If this assumption is not acceptable, the coverage norm in terms of Eq. (1) should be used, which however involves higher computational overhead.

A chromosome satisfies the X-checkpointing constraint if the checkpoint graph $\mathcal{G}$ is connected. We use depth first search to traverse the checkpoint graph until all nodes in $\mathcal{G}$ are reached or until no new node can be found. To test a chromosome for R-checkpointing, we extend the checkpoint graph $\mathcal{G}$ with the set of reference stations $\mathcal{R}$ which are connected among themselves and check the resulting graph for connectivity. Recall that R-checkpointing requires that each mobile vehicle is connected to a reference station, possibly over multiple hops.

To variate the given set of chromosomes, we first randomly assign the chromosomes to pairs. Then, a *crossover operator* is applied to each pair with probability 0.7. We use uniform crossover with a mixing rate of 0.5. Swapping vehicles between two chromosomes is only valid if the vehicle to be swapped is not yet present in the target chromosome. Additionally, in case of X- and R-checkpointing we only consider chromosomes that fulfill the corresponding connectivity constraint.

Each chromosome is *mutated* by selecting one vehicle at random, and replacing it by a randomly selected vehicle which is not yet used in the current chromosome. We keep generating offspring in this way for a maximum of five iterations, until a chromosome that satisfies the constraints is found. No offspring is generated if no feasible chromosome can be constructed.

After calculating the fitness of the newly created offspring, the algorithm uses the restricted tournament [8] *selection operator* to decide which chromosomes of the parents and offsprings are going to survive into the next iteration. In restricted tournament, a variated chromosome can only substitute its more similar parent, and only if its fitness value is better than the parent's fitness. The approach is elitist since the best solution is always kept. Also, the selection operator always preserves checkpointing constraints.

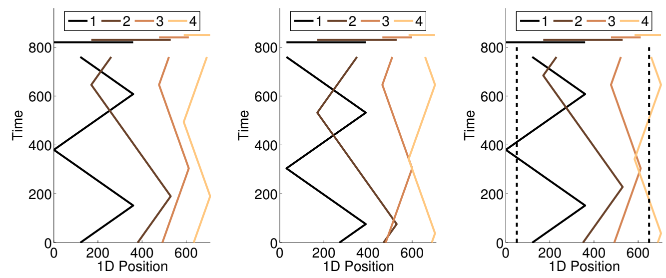

(a) Coverage-only (b) Coverage with X-CP (c) Coverage with R-CP

Fig. 2. Coverage optimization without checkpointing and with X- and R-checkpointing (X-CP and R-CP) in a one dimensional city with K=4 trams out of total 40 trams. The dashed lines denote reference stations.

In the next section we evaluate the proposed algorithm and compare its output to random search. For low values of K the optimal solution is computable and is used as a solution quality benchmark.

IV. Evaluation

Followed by the description of the evaluation setup, we first test our algorithm on "one-dimensional" cities to visualize and intuitively understand the solution. Later we run the algorithm on route and timetable data of a real city. Throughout the evaluation we meet the following three assumptions: (1) To simplify understanding of the algorithm, we calculate fitness values for the uniform constant density requirement $\rho \equiv 1$ using the previously introduced optimization. (2) Furthermore, two trams make a checkpoint if they come closer than 0 m and 200 m for the one-dimensional and two-dimensional cities, respectively. This corresponds to non-zero measurement validity w in all points in space within radius α. The checkpoint frequency in our calculations is $f_p \equiv 1$. (3) Finally, since the optimal solution is only computable for very small timetable networks, we use random search and simulated annealing for comparison if the optimal solution cannot be computed.

A. One-dimensional Cities

The projection of a city onto a one-dimensional space is a one dimensional city. A timetable network is thus reduced to a set of intervals and a vehicle movement forms a zigzag line in the time-space (see Fig. 2) called a trace. The beauty of one dimensional cities is the ease of visualizing them and understanding the vehicle selection results. The intervals representing vehicle tracks are depicted on top of Figs. 2(a)-(c). A vehicle can start moving from any position within its track in one of the two directions. For simplicity, we assume that all vehicles move with a constant speed which might differ among various mobile vehicles. Fig. 2(a) plots the algorithm output for the case when the fitness function optimizes coverage with no additional checkpointing constraints. Under these settings the traces of individual vehicles almost never cross, since crossing does not generally contribute to the coverage. Note that crossing of the spatial tracks does not imply that the trams following these tracks necessarily meet. In Fig. 2(b), we introduce an additional X-checkpointing constraint, which results in

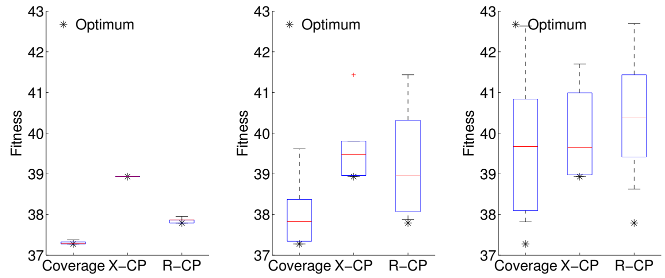

(a) EA (b) SA (c) Random

Fig. 3. Comparison of our evolutionary algorithm (EA) to the optimal solution, simulated annealing (SA), and random search (Random). We simulate a one dimensional city and select K=4 trams out of total 40 trams.

a connected subnetwork. The corresponding checkpoint graph $\mathcal{G}$ has a line topology by sequentially connecting the trams 1, 2, 3 and 4. Fig. 2(c) shows a solution to the coverage problem with R-checkpointing. Dashed vertical lines represent the locations of two static reference stations. In the computed solution the trams 2 and 3 do not make a checkpoint, but connect to the reference point over the trams 1 and 4, respectively. It can be seen in Fig. 3(a) that the achieved coverage is higher than with X-checkpointing. This is because the best possible solution does not necessarily fulfill the constraints introduced by X-checkpointing. Additional reference stations with R-checkpointing relax the constraints and allow for better solutions. The mutation hit rate exceeds 50% for the problem statements with and without checkpointing constraints. For the three discussed settings, we also compare in Fig. 3(b) the quality of the algorithm output to simulated annealing [12] and in Fig. 3(c) to random search. Our implementation of simulated annealing starts with a random feasible solution and uses the same mutation operator as the evolutionary algorithm to generate the new candidate solution. The algorithm probabilistically accepts the new solution depending on the current temperature which falls towards zero with each iteration. The optimum solution, which is computable for the given four tram lines, shows that the proposed evolutionary algorithm always finds a close to optimum solution. Note that the optimum can be computed only for very small networks.

B. City of Zurich

Zurich is one of the target cities in the OpenSense project about to deploy a network of mobile sensors on top of several trams. We run the algorithm on up-to-date data of the Zurich tram network. We obtained the track plan from OpenStreetMap and the timetable from the ZVV information service[3]. The Zurich tram network is depicted in Fig. 4. It serves 13 tram lines (see Table I) with the involvement of maximum 260 individual trams. The algorithm was tested for the subnetwork operating on a business day at 7 o'clock in the morning. We selected a time slice of two hours in order to include the round trip time of all operating trams. For simplicity we consider that there is no difference between business days and weekends or time of the day. The speed of a tram between two stations

[3]www.openstreetmap.org, www.zvv.ch

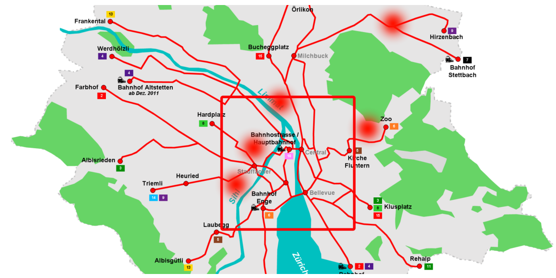

Fig. 4. Tram network of Zurich (source: wikipedia.org) with the positions of OstLuft and NABEL stations. Two left most stations used in the experiments.

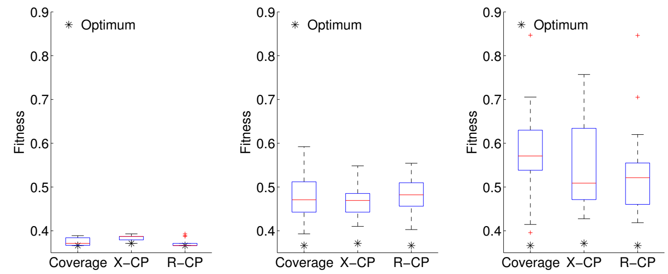

(a) EA (b) SA (c) Random

Fig. 5. Comparison of our evolutionary algorithm (EA) to the optimal solution, simulated annealing (SA), and random search (random). We use the tram network of Zurich and select K=2 trams out of total 260 trams.

is linearly interpolated, although it might differ between two station pairs depending on the timetable. The population size is set to 60 chromosomes. For the area of interest Ω we picked the Zurich city center, represented with a red square in Fig. 4.

We evaluate the quality of the algorithm output for three cases: coverage-only fitness function, coverage with X- and R-checkpointing constraints. The set of references $\mathcal{R}$ includes locations of two stations of the cantonal network OstLuft. The locations of the reference stations are visualized with large red dots in Fig. 4. In Fig. 5 we show results for the case when the selected subnetwork must contain $K = 2$ trams for EA, SA, and random search. The optimum solution for the coverage-only fitness corresponds to the setting with two trams (line Nr. 9 and Nr. 3) which start on the opposite corners of the city and never meet each other. X-checkpointing with zero temporal measurement validity gives two trams following the same line Nr. 9 from the opposite directions. The optimum solution for an additional R-checkpointing is composed of the lines Nr. 9 and Nr. 3, same as for the coverage only solution. Despite that the two trams never meet

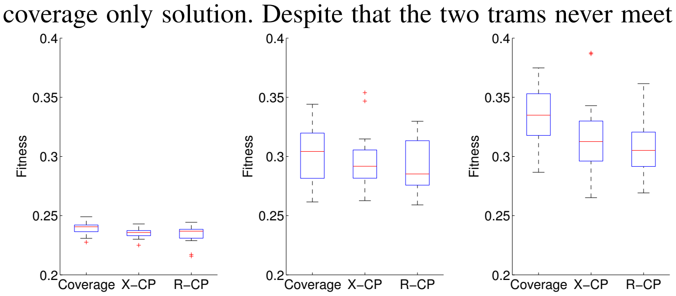

(a) EA (b) SA (c) Random

Fig. 6. Comparison of our evolutionary algorithm (EA) to simulated annealing (SA) and random search (random). We use the tram network of Zurich and select K=10 trams out of total 260 trams.

R-checkpointing is satisfied due to the available reference stations. The proposed EA finds a solution close to optimum and considerably outperforms SA and random search. In Fig. 6 we compute a subnetwork of size $K = 10$, which corresponds to the target number of stations we plan to place on top of trams in Zurich as part of the OpenSense project. The optimum solution is not computable in this case. However, the solution found with the EA provides better coverage than SA and random search. The mutation hit rate in EA exceeds 60%. Note that due to general connectivity of the public transportation network of a city and short temporal validity of a measurement, the values of coverage achieved do not change significantly if X-checkpointing constraints are applied. The execution of the proposed EA is limited to 20 iterations and lasts less than 30s for the case of ten trams on a ThinkPad T410 laptop, 2.6GHz.

V. Related Work

The work related to our approach can be divided in three groups: solutions to the area coverage problem with static sensors, route selection and planning algorithms in the context of area coverage, and route finding in timetable networks.

The problem of area coverage with sensors is often considered without sensor checkpointing in static settings. Related work on the topic includes solutions to the coverage problem as is [13], combined with event detection [5], and motion planning of mobile agents to achieve area coverage [2], [9], [15]. There is little work on checkpoint design in this context.

Towards checkpoint design, in [16] the authors present an approach for saving energy in wireless sensor networks by introducing a mobile base station and designing a set of rendezvous points for data collection. Further related approaches on the design of a movement pattern for the base station can be found in [10], [14], [6]. The following two properties distinguish our approach from the above solutions: (1) the underlying timetable network provides a *fixed* backbone and is a considerable limitation in terms of coverage; (2) all nodes in the network are mobile and thus time-dependent.

Similarly to our scenario, timetable networks consist of plenty of mobile nodes. Routing on timetable networks is currently a hot topic in the respecting community. In particular, the interesting problems are earliest arrival and minimum number of transfers when planing a route from A to B [11]. The main difficulty here is the computational overhead due to lack of hierarchical structure in timetable networks. Both problems are concerned with an efficient design of checkpoints. This closely reflects the problem we face when designing a connected timetable subnetwork. In contrast to route planning, we are rather interested in selecting a subnetwork with very short transfer times, and name these interchange points checkpoints.

We are not aware of any approach solving the coverage problem atop a timetable network with an additional checkpointing requirement.

VI. Conclusion and Future Work

In this paper we consider the problem of selecting a subnetwork of a public transport network to maximize coverage of a city with an additional support for sensor checkpointing. For example, when designing a network of K measurement stations on top of a tram network in a city, it is essential to have meeting points of different mobile vehicles to be able to detect sensor malfunction or recalibration needs. Upon the availability of a set of reference stations, the quality of the selected subnetwork can be further improved by considering the closeness to these reference stations. We solve both problems with an evolutionary algorithm since the computation of the optimal solution is only possible for very small timetable networks.

The results of this paper rely on the strong assumption that each mobile vehicle is uniquely identified by its track and its timetable. However, public transport vehicles can usually follow different timetables on different days and are not bounded to a specific route, but rather a specific depot which usually serves a subnetwork of the public transportation network. We plan to continue our work on the above problem statement incorporating uncertainties concerning the timetables and the routes taken. In this case, we have to reason about probabilistic checkpoints, their expected frequencies, and qualities. The results obtained in this paper will be used for comparison as the best achievable solution.

Acknowledgements: This work was funded by NanoTera.ch with Swiss Confederation financing.

References

[1] K. Aberer, S. Sathe, D. Chakraborty, A. Martinoli, G. Barrenetxea, B. Faltings, and L. Thiele. Opensense: open community driven sensing of environment. In *ACM IWGS*, 2010.
[2] J. Cortes, S. Martinez, T. Karatas, and F. Bullo. Coverage control for mobile sensing networks. In *IEEE RAS*, 2004.
[3] Q. Du, V. Faber, and M. Gunzburger. Centroidal voronoi tessellations: Applications and algorithms. In *JSTOR SIAM*, 1999.
[4] P. Dutta, P. M. Aoki, N. Kumar, A. Mainwaring, C. Myers, W. Willett, and A. Woodruff. Common sense: participatory urban sensing using a network of handheld air quality monitors. In *ACM SenSys*, 2009.
[5] A. Ganguli, S. Susca, S. Mart, F. Bullo, and J. Cort. On collective motion in sensor networks: sample problems and distributed algorithms. In *IEEE CDC*, 2005.
[6] Y. Gu, D. Bozdag, and E. Ekici. Mobile element based differentiated message delivery in wireless sensor networks. In *WOWMOM*, 2006.
[7] M. Gunzburger and J. Burkardt. Uniformity measures for point sample in hypercubes. In *Technical Report Florida State University*, 2004.
[8] G. Harik. Finding multimodal solutions using restricted tournament selection. In *Morgan Kaufmann ICGA*, 1995.
[9] I. I. Hussein and D. M. Stipanovic. Effective coverage control using dynamic sensor networks. In *IEEE CDC*, 2006.
[10] D. Jea, A. Somasundara, and M. Srivastava. Multiple controlled mobile elements (data mules) for data collection in sensor networks. In *Springer DCOSS*, 2005.
[11] E. Pyrga, F. Schulz, D. Wagner, and C. Zaroliagis. Efficient models for timetable information in public transportation systems. In *ACM JEA*, 2008.
[12] C. Skiścim and B. Golden. Optimization by simulated annealing: A preliminary computational study for the tsp. In *IEEE WSC*, 1983.
[13] A. So and Y. Ye. On solving coverage problems in a wireless sensor network using voronoi diagrams. In *Springer WINE*, 2005.
[14] A.A. Somasundara, A. Kansal, D.D. Jea, D. Estrin, and M.B. Srivastava. Controllably mobile infrastructure for low energy embedded networks. In *IEEE TMC*, 2006.
[15] Y. Wang and I. I. Hussein. Awareness coverage control over large scale domains with intermittent communications. In *IEEE ACC*, 2008.
[16] G. Xing, T. Wang, W. Jia, and M. Li. Rendezvous design algorithms for wireless sensor networks with a mobile base station. In *ACM MobiHoc*, 2008.