New structures to solve aggregated queries for trips over public transportation networks^{*}

Nieves R. Brisaboa¹, Antonio Fariña¹, Daniil Galaktionov¹², Tirso V. Rodeiro¹, and M. Andrea Rodríguez³⁴

 1 Universidade da Coruña, Fac. Informática, CITIC, Spain 2 Enxenio S.L., Spain

 3 Universidad de Concepcion, Computer Science Department, Chile 4 Millennium Institute for Foundational Research on Data, Chile

Abstract. Representing the trajectories of mobile objects is a hot topic from the widespread use of smartphones and other GPS devices. However, few works have focused on representing trips over public transportation networks (buses, subway, and trains) where user's trips can be seen as a sequence of stages performed within a vehicle shared with many other users. In this context, representing vehicle journeys reduces the redundancy because all the passengers inside a vehicle share the same arrival time for each stop. In addition, each vehicle journey follows exactly the sequence of stops corresponding to its line, which makes it unnecessary to represent that sequence for each journey.

To solve data management for transportation systems, we designed a conceptual model that gave us a better insight into this data domain and allowed us the definition of relevant terms and the detection of redundancy sources among those data. Then, we designed two compact representations focused on users' trips (TTCTR) and on vehicle trips (AcumM), respectively. Each approach owns some strengths and is able to answer some queries efficiently.

We include experimental results over synthetic trips generated from accurate schedules obtained from a real network description (from the bus transportation system of Madrid) to show the space/time trade-off of both approaches. We considered a wide range of different queries about the use of the transportation network such as counting-based/aggregate queries regarding the load of any line of the network at different times.

^{*} Funded in part by European Union's Horizon 2020 research and innovation programme under the Marie Sklodowska-Curie grant agreement No 690941 (project BIRDS). The Spanish group is also partially funded by Xunta de Galicia/FEDER-UE [CSI: ED431G/01 and GRC: ED431C 2017/58]; by MINECO-AEI/FEDER-UE [Datos 4.0: TIN2016-78011-C4-1-R; Velocity: TIN2016-77158-C4-3-R; and ETOME-RDFD3: TIN2015-69951-R]; by MINECO-CDTI/FEDER-UE [INNTERCONECTA: uForest ITC-20161074]; and by FPI Program [BES-C-2017-0085]. M. A. Rodríguez is partially funded by Fondecyt-Conicyt grant number 1170497 and by the Millennium Institute for Foundational Research on Data.

1 Introduction

The management of public transportation systems is a complex problem that has been typically faced from the point of view of the **offer** (lines, stops, schedules of journeys for each line, ...). In the last decade, the widespread use of new technologies allowing somehow the tracking of users' movements along a network transportation system (mobile phones with GPS, use of RFID technologies, smart cards used to pay and enter buses/trains, ...) brings new opportunities to gather the actual usage of the transportation systems allowing to study the problem from the point of view of **users' demand**. In consequence, it is now possible to develop new applications to exploit those data in order to effectively handle the resources of the transportation system and to give a better service to the users.

The management of the transportation system has become a Big Data problem in many important cities around the world, where millions of passengers use the public transportation network every day. Therefore, even though we can assume the gathered data is reliable (even in the case of depending on the smart cards provided to users that typically gather only the entry point to the network, the end point can commonly be derived using historical data from user trips and transportation models [9]), the problem lies now on how to represent user trips in such a way that not only we provide a compact representation but also we enable performing queries in an efficient way.

While there exist many works that tackle the problem of representing trajectories of mobile objects constrained to networks [12, 6, 5], they typically aim at locating the position of those objects from the underlying trajectories. Others [8, 7] focus on solving *strict and approximate path queries* that permit to find the trajectories that follow a given line pattern within a given time interval. The latter work [7] is, to the best of our knowledge, the first work using a compact data structure to represent the spatial data (a FM-index [4]). Yet, none of them have been designed to tackle the analysis of the usage of the transportation network and would hardly support queries such as *count the number of user trips that went from stop X to stop Y*, or *show the load of the lines at a given hour*.

In [1], a representation for user trips along a transportation network referred to as CTR was presented. The different stops from bus lines were given a *node-ID*. Then, each user trip was associated a string composed of the sequence of *node-IDs* traversed. Finally, a CSA-based representation [13] was used to represent the collection with all users' trips, and a Wavelet Matrix (WM) [2] aligned with the CSA represented, for each trip, the time instant when every node from that trip was reached. CTR enabled answering counting-based aggregated queries (number of users that started/ended a trip at a given node within a given time interval, number of users that used a node, top-k most used nodes, etc.). In addition, since it represented the actual trips in a compact self-indexed way, CTR still possessed enough flexibility to support more complex queries. CTR succeeded at providing a compact representation for general trips. Yet, it still represented trips in a redundant way when considering public transportation by bus, train, and subway. This happens because it does not exploit the fact that all the passengers in the same bus/train traverse each stop at the same time, nor the topology of the network (for all the users' trips from a stop X to a stop Y along a given line, all the intermediate stops are always the same).

In this work, we have analyzed the problem of representing both *offer* (stops, lines, schedule for each line) and *users' demand* (user trips, and stages that include stops where users get on/off or switch lines) within a public transportation network. We present a conceptual model that provides valuable insights into this domain and shows both the data needed and the relationships among them.

Then, we present two complementary structures to represent those data and show how they handle some useful queries in this context. The first solution is named Topology&Trip – aware CTR (TTCTR) and is based on a modification of CTR that also represents all the user trips but exploits both the network topology and the fact that all the passengers of the same vehicle journey reach the same stops at the same time, hence temporal information can be related to the vehicle journey rather than to each user trip. Therefore, it still makes up a general representation focused on users' trips. The second solution, named Accumulated – Matrix (AcumM), does not actually represents user's trips. It focuses on the journeys of each line, and accumulates the number of passengers that get on/off in each stop of each journey. Therefore, AcumM is a summarization of the load each line had considering each of its journeys, in the same way a data warehouse is a summarization of the operational data in a database.

The structure of the paper is as follows. In Section 2 we discuss the conceptual model associated to the network transportation problem and provide some definitions. In Section 3, we present our representation of the *offer* (lines, stops and journeys) which is then used in our two solutions. The next sections describe both TTCTR and AcumM and discuss the types of queries they are designed for. Section 7 includes experimental results to show the space/time trade-offs of our proposals. And finally, conclusions and future work are discussed in Section 8.

2 A model to describe a public transportation network

The E-R conceptual model at Figure 1 represents the relevant data of any public transportation system including data related with both the *offer* and the *demand*. We did not include entities such as vehicles or drivers, as they are out of the scope of this work. To create that model we have defined the following concepts:

- Stop (or Stop-place). Places were passengers can get on/off from a vehicle.
- Lines. A line (or route) is a sequence of stops that starts at a given stop X and ends in another stop Y. We consider a line and its *return* line as different lines because they include different sequences of stops.
- Journeys. We define a journey (or vehicle journey, or line trip) as a trip that a vehicle performs. It departs at a given day and time from the first stop of a specific line and follows the complete sequence of stops of that line until the ending stop, allowing passengers to get on and to get off in each stop. For instance, a *journey* is the trip that a bus performed along line L1 departing at 9:00am on 2017/05/05, and stopping at each stop of the line.

In addition to the day and time each journey starts, it could be interesting to have the time at which each stop was reached by each specific journey of the line. Yet, if such an accurate time is not needed, we can save a large amount of space by only storing, for each stop of a line, the average accumulated time needed to reach that stop from the initial stop (as in the examples shown along this paper). Some other solutions, with different trade-off between accuracy of the temporal data and storage space, are possible. For example, we could store the average time to get to each stop of a line in peak and nonpeak hours. In any case, all those strategies enable us to associate temporal data to the user-trips done within a given vehicle journey.



Fig. 1. E-R model for a public transportation network.

- Stages. A stage represents the pair of stops where a given user respectively gets on/off to/from a vehicle doing a given journey of a given line.
- User_trip. We define a user trip as a sequence of stages. That is, a user trip can begin at stop A from line 3 and continue up to stop B (first stage), then change to line 2 up to stop C (second stage), and so on. This enables tracking user trips from an origin to a destination. Note that, since stages refer to a given journey, and we can know the time when a journey traverses a stop X, we can also know when a given user trip reached such stop X.

3 Towards a practical representation: common structures to represent the *offer*

In Sections 4 and 5, we present two representations. The first one is focused on the representation of user trips, whereas the second one is focused on the journeys of each line and basically stores the number of users that get on/off at any given stop for each journey of a any line. Both techniques require some common structures that handle the data that represents the *offer* of public transportation the network provides. Such *offer* refers to the structure of the network and includes the representation of the lines, and, for each line, the schedule of its journeys; that is, their departure time from the starting stop. In addition, we use two aligned arrays for each line, one with the sequence of stops, and another with the average accumulated time to reach each stop from the first stop of the line. Note that instead of assigning a unique sequence of estimated times to reach any stop from the line, we could have dealt with several, probably more

4

accurate, estimations for peak/low periods, or even we could have stored the actual time each journey reached each stop. In any case, we can estimate the time when each journey reaches each stop.



Fig. 2. Example of (bus) public transportation network.

Figure 2 includes an example of a bus network with two lines (1 and 2). For each line we show the stops that compose it (e.g. Line 2 contains the sequence of stops $\langle S13, S6, S10, S5, S11, S9, S14 \rangle$) and the accumulated times from the initial stop (e.g. the average time to reach the fourth stop S5 from the starting stop of the line is 433 seconds). We also include the starting times for each journey of each line. In this case, Line 1 has 48 journeys per day, the first one starts at 6:00am, the second one at 6:20am, etc.

Note that given a line X we have direct access to the information related to the *i*-th stop. Yet, given a stop, we do not know the line/s it belongs to. To overcome this issue, we include, for each stop Y, the list of lines that include such stop Y. It is referred to as *inverted index for stops* in the bottom of Figure 2.

To sum up, we saw that to represent the network offer we need: (i) a sequence of stops for each line⁵; (ii) a schedule with the starting times of the journeys of each line; and (iii), an inverted index to mark the lines each stop belongs to.

Apart from the network offer, in Figure 2 we also include (arrows) five user trips done along the network. For example, there is a user trip (dashed arrow from S3) that starts at stop S3 at 06:25am on day-1 (6:20am + 305 sec.), and follows the journey of line 1 until S10, where the user switches to line 2 at time 6:35am (6:30am + 300 sec.) and continues the corresponding journey of line 2 (the one started as 06:30h in S13) up to stop S12. That is, it includes two stages.

 $^{^{5}}$ We also store average estimated times to reach each stop from first stop of the line.

6

4 **Topology**&**Trip** – **aware CTR** (**TTCTR**): a representation focused on user's trips

A previous representation for user trips along a transportation network, named CTR [1], basically associates an integer s_i to each stop in the network, and represents a user trip t_i as the sequence of the stops traversed plus a terminator $(t_i = s_1, s_2, \ldots, s_k)$. Finally, a CSA-based representation is used to represent the collection with all users' trips, and a Wavelet Matrix (WM) aligned with the CSA keeps, for each trip, the (discretized) time instant in which every stop from that trip was reached. For example, the trip from Figure 2 that started at S3 would be represented as $\langle S3, S10, S5, S11, S9, S12, \rangle$ and the times associated to those stops would be discretized into 5-min time periods. CTR exploited the indexing capabilities of the underlying CSA and WM to solve counting-based spatial, temporal, and spatio-temporal queries.

Our TTCTR is an adaptation of CTR that represents a user trip as a sequence of stages rather than as a sequence of stops (hence exploiting the topology of the network). Furthermore, instead of having to represent the time each user trip reaches a stop, we will only store a reference/*id* of the journey (within a vehicle of a line) that the user used. The building process of TTCTR is presented below.

Let us assume a network with n_s stops (S) numbered $s \in [1, n_s]$; n_l lines (L) numbered $[1, n_l]$, and that there are n_j^l journeys (J^l) for each line $l \in L$ numbered $[0, n_j^l - 1]$. Additionally, we have the starting times for each journey and the accumulated average times for the stops of each line as discussed in the previous section. We can define that a user gets on/off from a vehicle following the journey j of line l at a given stop s, as a triple (s, l, j) where $l \in L$, $s \in S$, $j \in J^l$.

Let us define $\mathcal{T} = \{t_1, t_2, \ldots, t_z\}$ as a set of z user trips. Since we want to represent a user trip t_x as a sequence of k stages, but it holds that the final stop of a stage and the starting stop of the next stage are the same (or close in walking distance), it is not necessary to explicitly represent the final stop of each stage, except for the final stop. We define $t_x = \langle (s_1, l_1, j_1), (s_2, l_2, j_2), \ldots, (s_{k+1}, l_{k+1}, j_{k+1}) \rangle$, $k \geq 1$. That is, we have a sequence of k triples that indicate that the user got on a vehicle corresponding to the j_i -th journey of line l_i at stop s_i . The last triple indicates where the user finally got off. Note that, for the last two triples, $l_k = l_{k+1}$, and $j_k = j_{k+1}$ since the beginning of the last stage is represented by the k-th triple, and its end by the (k + 1)-th triple.

Example 1. Assuming that all the user trips depicted in Figure 2 belong to our 1st-day, those user trips can be represented as: $t_1 = \langle (1, 1, \mathbf{0}), (10, 2, \mathbf{1}), (11, 2, \mathbf{1}) \rangle$, $t_2 = \langle (2, 1, \mathbf{1}), (7, 1, \mathbf{1}) \rangle$, $t_3 = \langle (3, 1, \mathbf{1}), (10, 2, \mathbf{2}), (12, 2, \mathbf{2}) \rangle$, $t_4 = \langle (6, 2, \mathbf{0}), (11, 2, \mathbf{0}) \rangle$, and $t_5 = \langle (13, 2, \mathbf{2}), (9, 1, \mathbf{2}), (14, 1, \mathbf{2}) \rangle$. Note that, for example, $(13, 2, \mathbf{2})$ from t_5 indicates that, at stop 13, the user got on a vehicle from line 2, that corresponds to the **2**-nd journey. We know it is the **2**-nd journey because t_5 started at 06:30h, which is the departure time of journey **2**. Note also that the line and journey *ids* of the last triple of each trip are identical to the ones in the previous triple. \Box

In TTCTR, we represent both the spatial (lines and stops) and the temporal component (journeys) of the user trips of our collection of trips \mathcal{T} using respec-

tively a CSA and a WM aligned with the CSA. In the following sections we show how we handle such components, and how we solve some queries of interest.

4.1 Representing the spatial component of TTCTR with a CSA

We use a variant of the CSA [3] for integer alphabets to represent the spatial component, i.e. the sequence of pairs (s_i, l_i) that compose each user trip in \mathcal{T} . However, in order to create a CSA we need to assign each pair (stop, line) a unique integer *id*. This will allow us to create an integer sequence S[1, n] (ended by a \$ terminator) over which our CSA will be built. For this purpose we create a vocabulary V (with $1 + n_s(1 + n_l)$ entries) as follows:

- Entry V[0] is reserved for the terminator symbol \$.
- Entries $\langle V[1], V[2], \ldots V[n_s] \rangle$] are associated to stops $\langle 1, 2, \ldots, n_s \rangle$ and are used to represent the final stops of the trips. That is, when a given stop x ends a user trip, it is given $id \leftarrow x$.
- The last $n_l \times n_s$ entries are associated to the sequence composed of the pairs $(s,l) \in S \times L$ considering that those pairs are sorted by s and l respectively. That is, entry $V[n_s+1]$ is given to (s,l) = (1,1); $V[n_s+2]$ to (1,2); $V[n_s+3]$ to (1,3); ...; $V[n_s+n_l]$ to $(1,n_l)$; $V[n_s+n_l+1]$ to (s,l) = (2,1), $V[n_s+n_l+2]$ to (2,2), and so on. In practice, in this case, the id/pos-in-V for a pair (s,l) is obtained as $id \leftarrow n_s + n_l(s-1) + l$.

Note that there will be a large number of unused entries (holes) in V. Yet, this can be efficiently handled by a bitvector B with rank/select capabilities that marks the used entries from V. Therefore, once we gather the position (id) corresponding to a pair (s, l) in V, we obtain its final position (id') in a vocabulary without holes (V') as $id' \leftarrow rank_1(B, id)$. Our id' assignment ensures that the used pairs (s_i, l) , corresponding to a given stop s_i , receive contiguous id's. This will be interesting at query time.

The next step processes each user trip $t_i \in \mathcal{T}$, $i \in [1, z]$ replacing all the pairs (stop, line) in t_i (except in the last one where the line is already known and we only need s) by their corresponding id'. After each trip, a 0 (id' for terminator \$) is added. That is, a trip t_i with k stages is regarded as a string $s_1s_2\ldots s_ks_{last}$, where $1 \leq s_{last} \leq rank_1(B, n_s)$ is the id' of ending stop given to stop s.

Once this process has completed, a sequence S[1, n] that contains only values from V' is obtained, and a CSA can be built on S.

In parallel with the construction of S, we create a sequence Jcodes[1, n] aligned to S where we set, for each trip t_i , the *journey-id* corresponding to each stage in t_i . Recall the *journey-id* is the third term from the triples $(s_i, l_i, j_i), i \in [1, k]$ from t_i . In addition, assuming that S[p] contains the 0 corresponding to the terminator f for the trip, we set $Jcodes[p] \leftarrow j_1$ (i.e. the same journey-id as the starting stop of the trip). According to the discussion above, Figure 3 shows, for the user trips in Example 1: (step-1) the sequence of pairs (stop, line) for each trip, and the corresponding Jcodes; (step-2) the vocabularies, including V, B,

sto J	op:line codes	1 1:1 0 Fo	2 10:2 1 or pair	3 11:* 1 s (s:l)	4 \$ 0 where	5 2:1 1 e the	6 7:* 1 ine is	7 \$ 1 marke	8 3:1 1 ed wit	9 10:2 2 h an "	10 12:* 2 *", sto	11 \$ 1 op "s"	12 6:2 0 corres	13 11:* 0 pond:	14 \$ 0 s to th	15 13:2 2 e last	16 9:1 2 stop	17 14:* 2 of an 1	18 \$ 2 user ti	19 \$ 0 ip.	user trips Including the trailing for each trip
V [В[V ^I [0 1 \$ 1: <u>*</u> 1 0 \$ 0	2 3 2:* 3 0 0 7: <u>*</u> 1	3 4 * 4.* 0 0 11:* 2	5 6 5:* 6: 0 0 12:* 7	7 * 7:* 1 14:* 4	8 9 8: <u>*</u> 9: <u>*</u> 0 0 1:1	10 10: <u>*</u> 1 0 2:1	11 12 1: <u>*</u> 12: <u>*</u> 1 1 3:1 1 7	13 1 13: <u>*</u> 14 0 7 6:2 9	4 15 1:1 1 1 9:1 1 9	16 17 1:2 2:1 0 1 10:2 1 10	18 19 2:2 3: 0 1 3:2 1 11	9 20 1 3:2 . 0 . Vocab	24 <u>5:2</u> 0 ulary	25 26 6:1 6:2 0 1 with n	27 28 7:1 7:2 0 0 o hole	3 29 3 2 8:1 8 0 s	30 31 1:2 9:1 0 1 2	32 33 9:2 10: 0 0 Voca	34 1 10:2 1 1 bular	35 39 40 41 42 1:1 13:1 13:2 14:1 14:2 0 0 1 0 0 y of pairs (stop:line)
	s	1	2 10	3	4	5	6	7	8	9 10	10 3	11	12 8	13	14 0	15 11	16 9	17 4	18	19 0]
	A	19	4	7	11	14	18	6	3	13	10	17	1	5	8	12	16	2	9	15	Bassie
	A'	19	18	4	7	11	14	regu	lar no	n-cyc	lical				•						CSA built on S
	D	1	0	0	0	0	0	1	1	0	1	1	1	1	1	1	1	1	0	1	
	Ψ	1	12	13	14	15	19	3	2	5	4	6	17	7	18	9	11	8	10	16	
	Ψ'	12	13	14	15	19	1	regu	ılar no	n-cyc	lical										
Jc	odes [₩]	0	0	1	1	0	2	1	1	0	2	2	0	1	1	0	2	1	2	2	

Fig. 3. Structures involved in the creation of a TTCTR.

and V' (ending stops s of trips do not need the line, therefore we use "s:*".);⁶ and (step-3) the structures involved in the creation of CSA from which TTCTR uses Ψ , D, (and V', B).

As in CTR [1], we sort the terminators considering that each trip is a cyclical string. For instance in Figure 3, S[18] = 0 would traditionally be followed by S[19] = 0, but for sorting purposes we consider it is instead followed by S[15] = 11. After that, we make Ψ cyclical in the terms of each user trip. That is, let us assume that a user trip lays on $S[i] \dots S[i+k+1]$, i.e. the terminator of the trip is at position e = i + k + 1 in S. Therefore, being A[j] = e, we modify $\Psi[j]$ in such a way that $A[\Psi[j]]$ points not to the initial position e + 1 of the next user trip, but it cyclically points to the beginning of the same trip; that is, we set $\Psi[j] \leftarrow \Psi[A^{-1}[i]]$.⁷ Using a cyclical Ψ will enable searching efficiently for user trips that started at a stop X and ended at a stop Y, as we will see below.

4.2 Representing the temporal component of **TTCTR** with a WM

The temporal component of TTCTR includes the sequence Jcodes described above. Recall *Jcodes* contains *journey-ids* aligned to the values in *S*, and that, for every line *l* there are n_j^l journeys sorted by their starting time and numbered as $0 \dots n_j^l - 1$, and also we have average accumulated times to reach each stop in the line. Therefore, this representation allows us to describe exact times for each stop. In practice, we use $Jcodes^{\Psi}$, which is aligned to Ψ rather than to *S*. See *step-4* in Figure 3. Note that Jcodes[8] = 1 corresponds to $Jcodes^{\Psi}[14] = 1$, since A[14] = 8; Jcodes[9] = 2 corresponds to $Jcodes^{\Psi}[18] = 2$, since A[18] = 9; and so on. We represent the sequence $Jcodes^{\Psi}$ with a WM. This saves space, and provides indexing capabilities to the temporal component.

⁶ In this example, with only 5 trips, we have only 11 used pairs in V, but in a real scenario for each stop of each line (existing pair (s, l)) there will be a 1 in B.

⁷ Note that A^{-1} is the inverse of the suffix array A; i.e. $A^{-1}[i] = j$ iff A[j] = i.

New structures to solve aggregated queries for trips over pub. transp. netws.

5 Dealing with aggregated data: Accumulated – Matrix

We propose AcumM as an intuitive solution for the representation of two-dimensional matrices of integers with support to aggregated queries by row, column, or window/range. In the context of a public transportation network we found queries referred to a given line where data must be aggregated either by *stop* (e.g. number of users that got on a vehicle at stop X); by *time-interval*, hence referred to a sequence of consecutive journeys within that time-interval (e.g. number of users got on at any stop of the line on 2017/03/24); or by *stop and time-interval*.



Fig. 4. Example, for a given line, of get-on matrix (a), accumulated matrix (b), and compact representation with gaps (c). Stop and column 0 are virtual. The values in column/row 0 are set to zero to simplify operations.

Let us assume that, for a given line, we have a matrix M^+ that stores the number of users that got-on at each stop (column) from each journey (row). Figure 4.(a) includes an example. To efficiently support aggregated queries, we create, for each line, the accumulated get-on matrix M for M^+ . We compute the value of a cell $M(r,c) \leftarrow \sum_{i=1}^r \sum_{j=1}^c M^+(i,j)$. That is, each cell contains the sum of all the values from position (1, 1) to position (r,c). M is depicted in Figure 4.(b). The accumulated matrix M allows us to solve a range count query over M^+ in O(1) time by computing: $\operatorname{countRange}((x_1, y_1), (x_2, y_2)) \leftarrow$ $M(x_2, y_2) - M(x_2, y_1 - 1) - M(x_1 - 1, y_2) + M(x_1 - 1, y_1 - 1)$. In AcumM, we actually represent, for each line, two accumulated matrices that count, respectively, the passengers that get on and get off to/from a journey in each stop.

countRange allows us to add: (i) consecutive values of a column (e.g. users that got on in a stop X in consecutive journeys, such as those in one day); (ii) consecutive values in a row (e.g. users that, for a given journey, got on along a consecutive sequence of stops, such as those in a neighborhood); and (iii) values in a window (e.g. users that got on in a consecutive sequence of stops in consecutive journeys).

Being C the capacity of a vehicle, a simple way to decrease space usage on M (it has $n_s \times n_j^l$ integers) consists in keeping the middle column $m \leftarrow (n_s + 1)/2$ explicitly, and representing the values in the other columns $m \pm k$ as the difference with respect to column m. This is depicted in Figure 4.(c). The differences in columns $m \pm k$ require at most $\lceil \log_2 kC \rceil$ bits, while retaining direct access.

6 Performing queries on **TTCTR** and **AcumM**

AcumM and TTCTR are designed for different purposes and therefore each highlight in different types of queries.

Queries for Accumulated – Matrix: This data structure resembles a datawarehouse, i.e. it stores aggregated values (rather than individual trips) to efficiently answer aggregated queries about the number of users (load of the network) in a given stop or group of stops over one or more journeys. Recall countRange efficiently sums the values of any submatrix. Yet, some useful queries could need more than one countRange operation, and then to either aggregate or compute the average of those results. For example, if we want to know the *average number of users that got on in line L in stops of a neighborhood (consecutive stops) between 8:00 and 9:00 along the last month*, we will need a countRange operation for each window including the consecutive stops and the consecutive journeys inside that period for each day. Finally, we add the results of those countRange operations (one per day) and divide the result by the number of days in the month.

Since we have both the accumulated matrix for users getting on and getting off, we can easily compute queries about the load of the journeys. For example, to know how many users were inside the vehicle of journey j from line L, between the stops X and X+1, we compute: $tot_up \leftarrow how$ many users got on in such journey j between stops 1 and X (inclusive), using the accumulated get-on matrix for line L (countRange of a row); in the same way, using the accumulated get-off matrix, we compute $tot_down \leftarrow how$ many users got off in the same journey and range of stops; and finally, we return the value $total \leftarrow tot_up - tot_down$.

Queries for TTCTR: Recall that TTCTR actually stores all the individual trips. This allows it to answer queries about the patterns users follow when using the transportation network. For example, queries about how many users start their trips in stop X, or end their trips in stop Y, or even started their trips in stop X and ended in stop Y, can be efficiently answered because CSA easily locates the subsection devoted to each stop, and the cyclic encoding of the trips allows to ask for patters such as X or Y or even Y. Note also that our way to encode the pairs stop:line guaranties that the occurrences of a stop for different lines are consecutive in the CSA, therefore, we can ask both how many users start their trips in stop X or how many users start their trips in stop X of line L. Finally, using the WM we can filter out those queries by time using the journeys.

Note that none of those queries can be answered by AcumM, which stores the number of passengers getting on to (off from) a journey in each stop but cannot track individual trips. On the other hand, in TTCTR, queries about the load of the transportation network, such as *number of passengers into the vehicle in journey j between stops X and X*+1, would become very time consuming.

New structures to solve aggregated queries for trips over pub. transp. netws.

11

7 Experimental evaluation

We created a synthetic collection of user trips generated from a GTFS⁸ description of urban and medium-distance buses⁹ from Madrid, with 1049 different lines and 10913 stop locations. We used real stop times from the journeys provided by the GTFS to generate ten million user trips over the span of a month. Each user trip created had one or more stages, defined as pairs (*stop_in, stop_out*), being *stop_in* and *stop_out* respectively triples (*stop, line, journey*) that determine the stop, line, and the journey where the user got on and got off.

The created user trips started from a random stop on a random journey, and followed the stops on that journey. After at least two traversed stops, we used a probability table to determine if the stage ends and user switch lines. In such case, we simulate the user getting off from that journey and either waiting on the same stop (at most 30 minutes) or walking to a nearby (100 meters) stop to get on to a new journey. We ensured there were no inconsistencies in our generated trips (i.e. users getting on the same line from which they just got off). Trip lengths were limited to 100 stops. Yet, after each traversed stop, the probability for ending the trip was 0.01λ , where λ is the number of previously traversed stops.

Finally, we represented all those trips using TTCTR and AcumM.

Implementation details: Due to the relatively small size of the network, the common structures were built using plain arrays of fixed size integers. Table 1 shows the space occupied by these structures.

For TTCTR, we used the CSA from [3] tuned the sampling rate for Ψ (t_{Ψ}) to the values $t_{\Psi} = \{32, 128, 512\}$. To represent bitvector D we used a *SDArray* [10]. In Table 2 (left) we show the space required by Ψ , D, B, and V' for an input of 35,702,981 entries in S and *Jcodes*, when compared to a baseline that uses fixed width integers to represent the pairs (*line*, *stop*) in the trips (that is, of $\lceil \log_2 |V| \rceil$ bits/entry, where V is the vocabulary defined in Section 4.1). In the WM of the TTCTR we used a *RRR* bitvector to compress the bitmaps [11]. We set sampling parameter $s \in \{32, 64, 128\}$, as shown in the Table 2 (right), where we compare the space of WM with a plain representation of the *journey-ids* that uses the number of bits needed to represent the maximum number of journeys on any line, $\arg \max_{l \in L} \lceil \log_2 J^l \rceil$.

	(i) Lines	(ii) Schedules	(iii) Inverted index of stops	Overall size			
Size (KiB)	176.54	7299.66	139.96	7616.17			
Table 1. Size of common structures (in KiB).							

⁸ https://developers.google.com/transit/gtfs/reference

⁹ Provided by CRTM (http://www.crtm.es)

			t_{Ψ}				RRR sampling		
		32	128	512			32	64	128
	iCSA	27.13 MB	21.91 MB	20.58 MB		WM	42.35 MB	39.43 MB	37.98 MB
		(39.84)%	(32.17)%	(30.22)%			(71.07)%	(66.18)%	(63.73)%
Ta	able 2	2. Compre	ssion of th	e CSA (lef	t) and '	WM ((right) con	ponents f	from TTCT

Percentages refer to the sizes with respect to the size of the uncompressed baseline.

	Accumulated matrix	Differential Matrix
get-on matrix	11189 KB (100%)	5596 KB (50.01%)
get-off matrix	11189 KB (100%)	5596 KB (50.01%)
Table 3. Siz	zes of the different	AcumM variants

For AcumM, we consider both the simple *accumulated-matrix* and the version using differential encoding. In the former case we have a simple matrix of integers. In the latter one, the middle column is kept apart as an integer array. For the rest of columns in the *differential matrix* values in each cell are encoded with $\lceil \log_2 N \rceil$ bits, being N the maximum difference (i.e. value on those columns).

Query execution times: We run experiments to show the query execution times of our proposals. An Intel Xeon E5-2620v4@2.1GHz machine was used.

On TTCTR, we tested several configurations for the query number of user trips from stop X to stop Y, labeled as xy_* in Table 4. The entry for xy with no subindices applies no line nor time restriction. xy_S and xy_E restrict, respectively, the <u>S</u>tarting and <u>E</u>nding stop to a specific line. xy_T denotes a <u>T</u>emporal restriction (at one random day). Therefore, combinations of these subindices stand for combinations of these three restrictions. We randomly generated a set of 10,000 query patterns by choosing trips from all the available user trips. In the densest setup ($t_{\Psi} = 32$, RRR = 32) all the queries are answered in around 6-30 μ s.

The last row also includes the times to solve the query: How many users got on in a stop X on a given line during a given day? (J^kS^1) . We also implemented this query in AcumM to compare the efficiency for these type of queries.

		RRR = 32			RRR = 6	4		RRR = 12	8
	$t_{\Psi} = 32$	$t_{\Psi} = 128 \ t$	$\Psi = 512$	$t_{\Psi} = 32$	$t_{\Psi} = 128$	$t_{\Psi} = 512$	$t_{\Psi} = 32$	$t_{\Psi} = 128$	$t_{\Psi} = 512$
ху	6.94	10.28	22.89	6.90	10.23	22.80	6.87	10.21	22.82
xys	6.89	10.23	22.95	6.86	10.22	22.82	6.90	10.21	22.74
×y _E	29.37	62.63	192.86	29.13	62.60	192.01	29.14	62.38	192.28
XYSE	29.21	61.84	192.11	28.88	61.83	190.46	28.92	62.05	190.47
×ут	31.85	64.66	195.62	31.63	65.52	195.06	31.95	65.47	195.31
ху _{ST}	31.61	63.83	193.41	31.12	65.05	193.14	31.54	64.78	194.23
×у _{ЕТ}	31.73	65.04	195.53	31.66	65.73	195.15	32.01	65.09	195.45
×у _{SET}	31.36	63.96	194.12	31.17	64.99	193.11	31.42	64.42	193.40
J ^k S ¹	5.05	5.85	9.14	5.20	6.01	9.30	5.41	6.31	9.63

Table 4. Performance at query time shown in μ secs/query for TTCTR.

	Accumulated matrix	Differential Matrix
$J^kS^1(column)$	131	211
J^1S^* (rows)	107	221
$J^kS^k\ (\mathrm{window})$	76	182

Table 5. Performance at query time for the variants of AcumM (in *ns* per query).

To test AcumM we considered three types of queries: The query J^kS^1 discussed above; total number of passengers that got on in all the stops of a given journey (1-row), labeled J^1S^* ; and total passengers that got on along a range of consecutive stops from several consecutive journeys (window), labeled J^kS^k . We generated 20,000 query patterns based on the real data (line number, stop number, and journeys), and then run the tests obtaining query times around 0.1μ s when using the accumulated matrix. As expected, the differential accumulated matrix performs around twice slower. Yet, AcumM performs more than one order of magnitude faster than TTCTR on query J^kS^1 . Table 5 shows the results.

8 Conclusions and future work

We have analyzed the problem of representing trips over a public transportation network and presented two data structures designed to efficiently answer two subsets of useful queries. Both approaches use some common data structures defined to represent the transportation network, that is the *offer* (lines, schedule of their journeys and stops) it provides.

The first proposal, TTCTR, represents the whole set of user trips during a period of time. Each user trip is composed of stages performed over specific journeys of different lines. This data structure is useful to analyze user trip patterns, that represent the real *demand* over the transportation network. Yet, TTCTR enables not only counting-based queries for the number of passengers related to any stop of the network, but also queries for stops or stops-lines were users start/end trips or switch lines. It also allows to retrieve individual trips.

The second structure (AcumM) focuses on the usage of lines. For each line, it keeps in an accumulated fashion, the number of passengers that, at each stop, get on/off from the vehicle performing a specific journey of the line. This simplifies solving queries regarding the load of the different journeys and therefore to analyze when specific lines must be reinforced with a more frequent schedule.

We understand more research is needed in this topic. Even we can see AcumM as a data warehouse were we store basically the same information but in an aggregated way, we consider that avoiding the redundancy between both structures would be desirable. As future work, we are interested in developing a unique self-indexed structure providing the functionality included in both TTCTR and acumM.

References

- Brisaboa, N.R., Fariña, A., Galaktionov, D., Rodriguez, M.A.: A compact representation for trips over networks built on self-indexes. Information Systems 78, 1–22 (2018). https://doi.org/10.1016/j.is.2018.06.010
- Claude, F., Navarro, G., Ordóñez, A.: The wavelet matrix: An efficient wavelet tree for large alphabets. Information Systems 47, 15–32 (2015)
- Fariña, A., Brisaboa, N.R., Navarro, G., Claude, F., Places, Á.S., Rodríguez, E.: Word-based Self-Indexes for Natural Language Text. ACM Transactions on Information Systems 30(1), article 1: (2012). https://doi.org/10.1145/2094072.2094073
- Ferragina, P., Manzini, G.: Opportunistic data structures with applications. In: Proc. 41st IEEE Symposium on Foundations of Computer Science (FOCS). pp. 390–398 (2000)
- Funke, S., Schirrmeister, R., Skilevic, S., Storandt, S.: Compass-based navigation in street networks. In: Proc. 14th International Symposium on Web and Wireless Geographical Information Systems (W2GIS). pp. 71–88. LNCS 9080 (2015)
- Kellaris, G., Pelekis, N., Theodoridis, Y.: Map-matched Trajectory Compression. Journal of Systems and Software 86(6), 1566–1579 (2013). https://doi.org/10.1016/j.jss.2013.01.071
- Koide, S., Tadokoro, Y., Yoshimura, T.: SNT-index: Spatio-temporal index for vehicular trajectories on a road network based on substring matching. In: Proc. 1st International ACM SIGSPATIAL Workshop on Smart Cities and Urban Analytics (UrbanGIS@SIGSPATIAL). pp. 1–8 (2015). https://doi.org/10.1145/2835022.2835023
- Krogh, B., Pelekis, N., Theodoridis, Y., Torp, K.: Path-based queries on trajectory data. In: Proc. 22nd ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems (SIGSPATIAL). pp. 341–350 (2014)
- Munizaga, M.A., Palma, C.: Estimation of a disaggregate multimodal public transport Origin–Destination matrix from passive smartcard data from Santiago, Chile. Transportation Research Part C: Emerging Technologies 24, 9–18 (2012). https://doi.org/10.1016/j.trc.2012.01.007
- Okanohara, D., Sadakane, K.: Practical entropy-compressed rank/select dictionary. In: Proc. 9th Workshop on Algorithm Engineering and Experiments (ALENEX). pp. 60–70 (2007). https://doi.org/10.1137/1.9781611972870.6
- Raman, R., Raman, V., Rao, S.S.: Succinct indexable dictionaries with applications to encoding k-ary trees and multisets. In: Proc. 13th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA). pp. 233–242 (2002)
- Richter, K.F., Schmid, F., Laube, P.: Semantic Trajectory Compression: Representing Urban Movement in a Nutshell. Journal of Spatial Information Science 4(1), 3–30 (2012). https://doi.org/10.5311/JOSIS.2012.4.62
- Sadakane, K.: New text indexing functionalities of the compressed suffix arrays. Journal of Algorithms 48(2), 294–313 (2003)