

TITLE OF INVENTION

[Title of the invention]

A Method for Searching or Comparing Sites
Using Travel Between Sites and Places
Within a Transportation System

CROSS-REFERENCE TO RELATED APPLICATIONS

[001] This application is a continuation-in-part of applications:

[Country]	[Application Number]	[Filing Date]
USA	16/274,242	February 13, 2019,
USA	16/864,157	May 1, 2020,
USA	16/894,761	June 6, 2020,
USA	16/940,418	July 28, 2020,

which are incorporated herein by reference as if fully set forth. The rest of the current application
5 contains additional matter, which may, or may not, have been disclosed by the applications. Any
prior art cited in this invention disclosure is considered to be ordinary knowledge in the art; any
person of ordinary skill in the art has that knowledge.

BACKGROUND OF THE INVENTION

[002] The present invention relates to searching or comparing sites. A traditional goal of searching is to find a site, from among a range of possible alternatives, that achieves an optimization objective, such as minimize a route length given specific travel requirements and desired features of the sought site. For example, when searching for real estate properties given required destinations of commutes and real estate property features, a goal may be to enumerate real estate properties with matching features that have the shortest commute durations. Other goal may be to compare any real estate properties using commute durations.

BRIEF SUMMARY OF THE INVENTION

[003] Embodiments include a method for searching or comparing sites, a computer system that implements and executes the method, and a computer service that receives search or compare requests from users and responds with site and route information.

[004] According to an embodiment of the present invention, a method for searching or comparing sites using routes or route lengths is provided. The method uses extensive preprocessing to precompute and store in a database routes or route lengths between each site and representatives within a transportation system. The method introduces a search-or-compare framework for sites. When a request containing a route specification is received, precomputed data is retrieved from the database to rapidly compute a route or a route length for each site. Sites may be searched or compared using routes or route lengths.

[005] According to an embodiment of the present invention, a computer system for searching or comparing sites using routes or route lengths is provided. The system is a combination of hardware and software. It obtains data about a transportation system and sites from one or more data providers. The system builds graphs that model travel between the sites and representatives within the transportation system. The system computes graph paths, and stores graph paths or graph path lengths. This enables to quickly compute routes or route lengths for every site when a request is received, and search or compare sites using routes or route lengths.

[006] According to an embodiment of the present invention, a computer service for searching or comparing sites using routes or route lengths is provided. The service allows the user to specify

a search or compare request through a User Interface on a device, for example a smartphone. The request contains a route specification and a filtering condition. In response, the service presents sites that match the filtering condition along with routes or route lengths for the matched sites, or the service compares sites using routes or route lengths.

5 [007] It is fundamentally necessary to quickly compute a route length for each site. We sketch a mathematical proof. Consider any search-or-compare method $\mathcal{M}$. We can design an adversarial request that will force $\mathcal{M}$ to respond with any given ordered list of sites. The adversary has two mechanisms at its disposal: (1) provide a request with a route specification that creates an order of sites with respect to the route length, the order selected by the adversary, and (2) provide a request
10 with a filtering condition on sites that matches a subset of sites that is selected by the adversary. Thus, the $\mathcal{M}$ must respond with an ordered list of sites arbitrarily selected by the adversary at request time. Details of the proof sketch are outside of the scope of a patent application.

[008] The method, the computer system, and the computer service each jointly performs tasks that are not simply generic nor well-understood by prior art. Prior art includes: US 8,417,409
15 B2; continuation US 8,738,286 B2; division US 8,756,014 B2;; KR 10-1692501 B1; continuation PCT/KR2016/01083; WO 2017/065431 A1; US 2018/0232824 A1;; and KR 10-1905593 B1.

[009] The embodiments of the invention presented here are for illustrative purpose; they are not intended to be exhaustive. Many modifications and variations will be apparent to those of ordinary skill in the art without departing from the scope and spirit of the embodiments.

20 [010] In the presentation, the terms “the first”, “the second”, “the”, and similar, are not used in any limiting sense, but for the purpose of distinguishing, unless otherwise is clear from the context. An expression in a singular form includes the plural form, unless otherwise is clear from the context. The terms “having”, “including”, “comprising”, and similar, indicate an existence of components or features, and do not preclude other components or features from existing or being
25 added.

BRIEF DESCRIPTION OF DRAWINGS

[011] The drawings included in the invention disclosure exemplify various features and advantages of the embodiments of the invention:

- FIG. 1: depicts an example color rendering of travel durations of commute paths according

to an embodiment of the invention, legend: “An example color rendering of travel durations of commute paths for real estate properties with two commute paths: property $\rightarrow$ geo.0.0 $\rightarrow$ geo.0.1 $\rightarrow$ property, and property $\rightarrow$ geo.1.0 $\rightarrow$ property. We used a map tile engine by Google, any other may be used instead.”;

- 5 • FIG. 2: depicts an example process flow of preprocessing and storing of data according to an embodiment of the invention, legend: “An example process flow of preprocessing and storing of data.”;
- FIG. 3: depicts an example process flow of responding to a request using preprocessed stored data according to an embodiment of the invention, legend: “An example process flow of responding to a request using preprocessed stored data.”;
- 10 • FIG. 4: depicts example commute paths according to an embodiment of the invention, legend: “Example commute paths.”;
- FIG. 5: depicts example extensions of a public transportation system graph for precomputing shortest graph paths according to an embodiment of the invention, legend: “Example public transportation system graph extensions for precomputing shortest graph paths.”;
- 15 • FIG. 6: depicts an example table / vector storage of precomputed shortest travel durations according to an embodiment of the invention, legend: “An example table / vector storage of precomputed shortest travel durations.”;
- FIG. 7: depicts an example of a commute path decomposition according to an embodiment of the invention, legend: “An example commute path $H \rightarrow W_1 \rightarrow W_2 \rightarrow W_3 \rightarrow W_4 \rightarrow H$ decomposed into a home-dependent part (green) and a home-independent part (black), illustrated on three homes H_1 , H_2 , and H_3 .”;
- 20 • FIG. 8: depicts an example of computing the shortest travel duration at the start of a commute path according to an embodiment of the invention, legend: “An example of computing the shortest travel duration at the start of a commute path.”;
- 25 • FIG. 9: depicts an example of computing the shortest travel duration at the end of a commute path according to an embodiment of the invention, legend: “An example of computing the shortest travel duration at the end of a commute path.”;

- FIG. 10: depicts an example pseudocode for computing the shortest travel duration at the start of a commute path according to an embodiment of the invention, legend: “An example pseudocode for computing PathFromHomeDurations($\rightarrow W$).”;
- FIG. 11: depicts an example flowchart of a computer system according to an embodiment of the invention, legend: “An example flowchart of a computer system for: precomputing and storing data in a database, and processing requests using data retrieved from the database.”;
- FIG. 12: depicts an example user request and a response by a computer service on a smartphone of the user according to an embodiment of the invention, legend: “An example user request and a response by a computer service on a smartphone of the user.”;
- FIG. 13: depicts an example of smoothing near a start of travel;
- FIG. 14: depicts an example of smoothing near an end of travel;
- FIG. 15: depicts an example of smoothing both near a start of travel and near an end of travel;
- FIG. 16: depicts an example of smoothing with at least one stopstation cluster near each endpoint of travel;
- FIG. 17: depicts an example of smoothing of three travel routes; and
- FIG. 18: depicts an example of smoothing.

[012] The drawings are for illustrative purpose only. Other drawings can exemplify the invention without departing from the principles of the invention, as will be readily recognized by one of ordinary skill in the art.

DETAILED DESCRIPTION OF THE INVENTION

4 Detailed description

[013] The invention concerns a general case of searching or comparing arbitrary sites using an arbitrary optimization objective that uses routes or route lengths to arbitrary places. However,

for the sake of ease of presentation, we first illustrate the invention through specific sites that are real estate properties, specific places that are workplaces, and a specific optimization objective of minimizing travel durations of commutes between real estate properties and workplaces. This illustration is not limiting. In later sections, we explain how the method works in the general case.

5 4.1 Real estate properties and commute paths

[014] Finding a home is a complex endeavor. People spend a significant effort of time and money on a search. Technology has come to help, however. There are several services available online that aggregate real estate property listings, and allow people to search for real estate properties with specific features, for example a price, a geographical location, a number of bedrooms, and so on, in the convenience of a web browser or a smartphone. This search yields a shortlist that the person then typically inspects in flesh.

[015] Location is arguably the most important feature of any real estate property, as attested by the “location, location, location” slogan of real estate agents. Our invention concerns this feature.

[016] Residents need to commute to work, school, or other places. The travel durations of these commutes bear on the value of the real estate property personalized for its specific residents. In addition, there is much global economic value in reducing the travel durations for all residents of a metropolitan area, in terms of time that can be allocated to more productive activities, energy that does not need to be spent on moving people around, etc. Our invention helps deliver these personal and global values.

[017] As a simple illustration, consider a family of two who live in the capital of South Korea. One person is a government employee at the City Hall, and the other person is a librarian at the Main Library of Seoul National University. They currently live in a two-bedroom one-bathroom apartment of size 69 square meters that requires a deposit payment of 350,000,000 KRW. The apartment is near latitude and longitude (37.5333, 127.0746). Their combined daily roundtrip commute lasts 1h 41m (about 25 minutes each way, on average). However, it turns out that they can decrease the travel duration to just 1h 16m (25 minutes savings) after moving to an apartment with the same features, but located near (37.5041, 126.8888). The shortest travel duration for any apartment is 50 minutes. However, that apartment has other price and size. (Illustration as of

February 19, 2018.)

[018] Why is it hard to deliver these improvements for all? A naive approach is to consider every real estate property available on the market that has the size and other features required by the family, and, given the work places, compute the travel durations by querying any existing
 5 online routing engine. This approach does not scale, however. One problem is the high number of real estate properties available on the market in a modern metropolitan area. Other problem is the high number of families/users who may want to seek improvements. Even if we assume that we can query a routing engine cheaply, the quadratic nature of the problem still makes the queries expensive in aggregate.

10 [019] How can we deliver these improvements for all? Our invention explains. It includes the following components:

1. The invention defines a model for a commute path. The model is versatile to cover a wide range of commute paths that occur in practice, for example going to school, then to a piano class, and then returning home. The utility of our model is enhanced by our ability to quickly
 15 find real estate properties that minimize travel durations.
2. The invention teaches an optimization method that rapidly computes travel durations. The invention identifies the parts of any commute path that are dependent on any real estate property. Travel durations for these parts are precomputed and stored. As a result, when travel durations need to be found for a commute path, the invention can rapidly assemble
 20 time parts to produce a travel duration *for every* real estate property.
3. An embodiment of the invention is a live computer service. The service enables the 25 million residents of the Seoul Metropolitan Area to search or compare real estate properties using travel durations.

4.2 Method outline

25 [020] We use the term *travel* in a broad sense that includes moving objects or data. A *description of travel* is anything that a person of ordinary skill in the art would name so. Here are some examples of a description of travel: (1) “hey buddy, you need to go one block north, and then turn slightly left”, and (2) ”5 dollars”. We may use a term *travel path* when we mean a

description of travel. A *length of travel* is a numeric value that a person of ordinary skill in the art can associate with travel, for example a distance, a monetary cost, etc. As other example, we may use a term *travel duration* when we mean a length of travel that represents time. A length of travel is by itself a description of travel. A description of travel: may not include any length of travel, may include only a length of travel, or may also include some other data.

[021] We illustrate some capabilities of the method and introduce terminology subsequently used. The method can compute travel durations of commute paths for every real estate property. FIG. 1 illustrates this. The metropolitan area is colored with squares. The colors represent how long it takes to commute from a real estate property in each specific area using public transportation: green means short commute, yellow longer, and red longest. A commute path will start at a *property*, then visit specific places *geo*, ..., and finally return to the *property*. In a sense the *property* is a free variable, while the *geo*, ... are fixed. In FIG.1, there are two commute paths: an “open-jaw” commute path $property \rightarrow geo.0.0 \rightarrow geo.0.1 \rightarrow property$ twice per week, and a “roundtrip” commute path $property \rightarrow geo.1.0 \rightarrow property$ three times per week. We can see that given these two commute paths and their frequencies, the real estate properties with small weighted sum of shortest travel durations form an irregular patch (dark green), which may not come as a surprise, given possibly intricate public transportation routes.

[022] At a high level, the method is composed of two parts. The first part computes travel durations between real estate properties and representatives that are vehicle stops of a transportation system. These travel durations are stored in a database, so that they can be readily retrieved when a request is received. See FIG. 2 for an illustration. The second part processes requests. A request contains a commute path and desired features of a real estate property. When a request is received, appropriate travel durations are retrieved from the database, and used together with other data to produce a travel duration for the commute path for every real estate property that has the desired features. See FIG. 3 for an illustration. Details and variants of this outline are described in subsequent sections.

4.3 Commute paths

[023] Our method handles a broad range of commutes that people perform within a metropolitan area. A commute starts at a site H that we call a *home*. The H is an arbitrary

location. It can be any real estate property, for example an apartment, a rented room, a house with a garden, a ranch, a hotel, etc. It can also be a site where a person works, a restaurant, a shop, etc. However, as a naming convention, we use the phrase “home” in most of the disclosure; this convention is not limiting. In one embodiment, a commuter travels to various places, and
 5 then returns to the site H .

[024] In one embodiment, a commute lasts within one day, for example a commuter departs from H in the morning and returns back to H in the evening of the same day. In other embodiment, a commute lasts within more than one day, for example if a person works a night shift, or has a work shift of 25 hours. In one embodiment, any travel may start at a specific time, or may end at
 10 a specific time, for example 8:12 AM. In other embodiment, any travel may begin, or end, within a range of time, for example “in the morning”.

[025] In a simplest form, a commuter travels from H to a place W , that we call *work*. The W is an arbitrary location. It includes a school, a grandparent’s home, a weekend golf course, a favorite restaurant, a doctor’s office, a place of worship, etc. It can also be a place where a person
 15 lives. However, as a naming convention, we use the phrase “work” in most of the disclosure; this convention is not limiting. Then the commuter returns from W to H . We call these two travels a *roundtrip commute path*. See FIG. 4A for an illustration.

[026] An *open-jaw commute path* is an example of a more complex commute. See FIG. 4B for an illustration. Here the commuter travels from H to a place W_1 . Then the commuter travels from
 20 the place W_1 to other place W_2 . Finally, the commuter travels from the place W_2 back to the site H . One example is when a person travels to a school and then to a piano class.

[027] In general, our method allows for arbitrary travels. In FIG. 4C, we see an illustration of a commute path with a missing travel from W_4 to W_2 , and a repeated travel from W_2 to W_3 . In FIG. 4D, we see illustrations of “open” commute paths: a commute path that departs from H ,
 25 but does not return to H ; and a commute path that returns to H , without previously departing from H . In FIG. 4E, we see an illustration of a “disconnected” commute path. A commute path could start at a home, but end at other home. Our method defines a commute path as follows.

Definition 1 A *commute path* is a collection of travels $W_2 \rightarrow W_3$, $W_4 \rightarrow W_5$, $\dots$, $W_{k-2} \rightarrow W_{k-1}$, for $k \geq 2$ that is even, together with $H_{first} \rightarrow W_1$ or $W_k \rightarrow H_{last}$ (so a commute path always
 30 contains at least one home and at least one work), that occur at arbitrarily moments of time.

[028] A commute path can be viewed as a specification of a route within a transportation system. The various W and H in Definition 1 specify where the commuter wishes to travel.

[029] In one embodiment, we consider a simpler commute path $H \rightarrow W_1 \rightarrow W_2 \rightarrow \dots \rightarrow W_k \rightarrow H$, for any $k \geq 1$, that has shared endpoints between travels, and the same home at the start and the end. We use this simpler form in most of the disclosure, because it is common in practice and simplifies our presentation. However, it will be obvious to anyone of ordinary skill in the art that our method applies to our (general) definition of a commute path.

[030] Our method finds a shortest travel duration along any commute path. This by itself has been addressed by various prior art. However, our method finds travel durations for all homes in a dramatically reduced amount of time.

4.4 Preprocessing of transportation system

[031] The method preprocesses data about public transportation system, so as to precompute and store shortest travel for all homes.

4.4.1 Computation of shortest travel

[032] We describe a method for efficiently computing shortest travel between all homes, and all locations of public transportation stops. We call the stops *stopstations* in most of the disclosure, and they include bus stops, subway stations, or both.

[033] The method starts with an arbitrary *public transportation system graph* GT that models a public transportation system, and that can be obtained from prior art. The graph may contain vertexes that represent bus stops, subway stations, or both. Other vertexes may exist in the graph, for example representing a stop or a turn of a vehicle, or a stop or a turn of a walk. Vertexes in the graph that correspond to stops of vehicles are denoted by

$$STOPSTATION_s,$$

indexed by s . There are directed weighted edges, each representing a travel duration along a segment from the source vertex of the edge to the target vertex of the edge. Other edges may exist in the graph. An edge may represent a transfer from one vehicle to other vehicle. The graph may contain data about departure times or arrival times of various public transportation vehicles. A Dijkstra's shortest graph paths algorithm, or an A* (A star) search algorithm, is commonly

used to compute a graph path with a shortest graph path length (sum of weights), representing a shortest travel duration, from a

$$STOPSTATION_s'$$

to other

$$STOPSTATION_s'',$$

for any s' and s'' , also if travel starts at a specific time, or ends at a specific time.

[034] Our method extends the public transportation system graph GT . See FIG. 5 for an illustration.

[035] The first extension introduces clusters of stopstations. The method clusters stopstations using any clustering algorithm; in one embodiment the method puts two stopstations into one cluster when the geographical distance between the two stopstations is at most a threshold, for example 5 meters, or when a travel duration between the two stopstations is at most a threshold. We may refer to the *location* of a cluster, when we mean a geographical location within the cluster, for example the center of the cluster. For each stopstation cluster c , the method adds two vertexes

15

$$STOPSTATION_CLUSTER_SOURCE_c$$

and

$$STOPSTATION_CLUSTER_TARGET_c,$$

and edges that connect the cluster with stopstations of GT

$$STOPSTATION_CLUSTER_SOURCE_c \rightarrow STOPSTATION_s'$$

labeled FirstWaitGetOn, and

$$STOPSTATION_s' \rightarrow STOPSTATION_CLUSTER_TARGET_c$$

labelled Zero, for any s' , whenever s' is in cluster c . The edges have zero weight. The resulting graph is denoted by GC , having vertexes VC and edges EC . A subset of the vertexes

25

$$STOPSTATION_CLUSTER_SOURCE_c,$$

for all c , is denoted by VS .

[036] The second extension introduces clusters of homes. The method clusters homes using any clustering algorithm; in one embodiment the method uses the same algorithm as when clustering stopstations. Similarly, we may refer to the location of a home cluster. For each home cluster s ,
 5 the method adds a vertex

$$HOME_CLUSTER_SOURCE_s$$

and a vertex

$$HOME_CLUSTER_TARGET_s.$$

The method connects each home cluster to stopstation clusters using walks. Specifically, the
 10 method adds an edge

$$HOME_CLUSTER_SOURCE_s \rightarrow STOPSTATION_CLUSTER_SOURCE_c$$

labelled Walk, if there is a walk from the location of the home cluster s to the location of the stopstation cluster c , for any s and c , with edge weight set to the duration of the walk; and a “reversed” edge

$$15 \quad STOPSTATION_CLUSTER_TARGET_c \rightarrow HOME_CLUSTER_TARGET_t$$

labelled Walk when there is a walk in the reverse direction from the location of the stopstation cluster c to the location of the home cluster t , for any c and t , with edge weight set to the duration of the walk. In one embodiment, the method limits walks to shortest walks at a specific speed, for example 4 km/h, lasting at most a threshold, for example 1 hour. The resulting graph is denoted
 20 by G . We denote by VH the set of vertexes

$$HOME_CLUSTER_SOURCE_s,$$

for all s ; and denote by EH the set of edges

$$HOME_CLUSTER_SOURCE_s \rightarrow STOPSTATION_CLUSTER_SOURCE_c,$$

for all s and c .

[037] We are interested in computing a shortest graph path from each stopstation cluster to each home cluster, and a shortest graph path in the reverse direction from each home cluster to each stopstation cluster.

5 [038] We remark that clustering allows our method to significantly improve the performance of a shortest graph paths computation, by “unifying” locations that are essentially the same with respect to shortest graph paths. For example, in a tall apartment complex, there may be hundreds of homes, and our method will “unify” them into just one home cluster. Thus, a shortest graph paths algorithm needs to compute shortest graph paths for just one home cluster, rather than for
10 hundreds of constituent homes.

[039] A standard application of a Dijkstra’s algorithm still yields poor performance, however. That application runs the algorithm from each

$$STOPSTATION_CLUSTER_SOURCE_c$$

in the G without any

$$15 \quad HOME_CLUSTER_SOURCE_s$$

vertexes, and then runs the algorithm from each

$$HOME_CLUSTER_SOURCE_s$$

in the G without any

$$HOME_CLUSTER_TARGET_t$$

20 vertexes. The combined asymptotic time complexity is

$$\begin{aligned} & O\left(|VS| \cdot \left(|EC| + |EH| + (|VC| + |VH|) \log(|VC| + |VH|)\right)\right) \\ & + \\ & O\left(|VH| \cdot \left(|EC| + |EH| + (|VC| + |VH|) \log(|VC| + |VH|)\right)\right). \end{aligned}$$

[040] Our method improves on this standard application. We observe that in a large metropolitan area it is often the case that the number of home clusters is significantly larger than

the number of stopstation clusters $|VH| \gg |VS|$, given the same clustering threshold. Utilizing this observation, our method uses a different algorithm for computing shortest graph paths from home clusters: the method reverses the edges of G , and for each

$STOPSTATION_CLUSTER_TARGET_c$

5 in the reversed graph, computes shortest graph paths to all

$HOME_CLUSTER_SOURCE_s$

vertexes using Dijkstra's algorithm (any

$HOME_CLUSTER_TARGET_t$

vertexes can be removed). This yields the desired effect, because when we reverse the edges of any

10 shortest graph path from

$STOPSTATION_CLUSTER_TARGET_c$

to

$HOME_CLUSTER_SOURCE_s$

in the reversed graph, we obtain a shortest graph path from

15 $HOME_CLUSTER_SOURCE_s$

to

$STOPSTATION_CLUSTER_TARGET_c$

in the original (not reversed) graph. Hence, the asymptotic time complexity of our method is just

$$O\left(|VS| \cdot ((|EC| + |EH|) + (|VC| + |VH|) \log(|VC| + |VH|))\right).$$

In practice, our method offers a significant reduction in the total running time on a graph G of
20 the Seoul Metropolitan Area.

[041] The method uses a symmetric algorithm in the opposite case when there are more stopstation clusters than home clusters $|VH| < |VS|$. In that case, the method reverses the edges when computing shortest graph paths from stopstation clusters to home clusters.

[042] In one embodiment, the method uses departure times. The method computes shortest
25 graph paths from each

STOPSTATION_CLUSTER_SOURCE_c

to every

HOME_CLUSTER_TARGET_t,

for a given departure time from the

5 *STOPSTATION_CLUSTER_SOURCE_c*.

In that case, the method uses an appropriate prior art graph *GT* that allows to specify the departure time from stops of public transportation vehicles; this case is sometimes represented using a vertex that corresponds to a time and a geographical location, and an edge that represents a travel duration starting from the geographical location at the time. Similarly, the method computes in
10 the reversed graph based on a given arrival time at

STOPSTATION_CLUSTER_TARGET_c.

Shortest graph path lengths convert these arrival times to the departure times from each

HOME_CLUSTER_SOURCE_s.

A similar graph is used to compute shortest graph paths given an arrival deadline. An appropriately
15 constructed graph can be used to compute a probability of arrival before a deadline.

[043] In one embodiment, we do not extend *GT* with any vertexes

STOPSTATION_CLUSTER_SOURCE_c

nor any vertexes

STOPSTATION_CLUSTER_TARGET_c.

20 Instead, we connect any vertexes

HOME_CLUSTER_SOURCE_s

and any vertexes

HOME_CLUSTER_TARGET_t

by direct Walk edges with any vertexes

$STOPSTATION_{s'}$.

[044] In one embodiment, we do not extend GT with any vertexes

$HOME_CLUSTER_SOURCE_s$

nor any vertexes

5 $HOME_CLUSTER_TARGET_t$.

Instead, we add vertexes

$HOME_s$,

each representing a home, that we connect to any vertexes

$STOPSTATION_{s'}$

10 directly by Walk edges.

[045] In one embodiment, we do not cluster homes.

[046] In one embodiment, we do not cluster stopstations.

[047] In one embodiment, we compute shortest graph paths given an arrival deadline: from stopstation clusters to home clusters given an arrival deadline at each home cluster, or from home
15 clusters to stopstation clusters given an arrival deadline at each stopstation cluster.

[048] In one embodiment, we use any shortest graph paths algorithm other than a Dijkstra's algorithm, for example an A* (A star) search algorithm. In one embodiment, we use an approximation algorithm for shortest graph paths. We may use an algorithm without any of the performance improvements.

20 [049] In one embodiment, the weights on some graph edges represent monetary costs of travel, instead of travel durations. Then our method searches or compares homes based on the monetary cost of commute paths. Any other semantic of edge weights can be used, for example: a number of transfers among public transportation vehicles, a duration of wait, a monetary cost of wait, or a distance of travel.

25 [050] In one embodiment, we apply a multi-objective optimization search based on a multi-dimensional cost. For example, we search for a shortest graph path whose length represents a

travel duration, such that a monetary cost that represents the graph path is at most a threshold, or when a monetary cost is added as a penalty to the graph path length.

[051] In one embodiment, a travel path is represented by a graph path. We compute various features of a shortest travel path using a shortest graph path, for example: the first (i.e., boarding) or the last (i.e., disembarking) stopstation of the shortest graph path, the main transit stopstation of the shortest graph path, the number of vehicle transfers, the vehicle where most time is spent during travel (for example bus 1234), the total wait time at stopstations, the total walk distance, whether the specific shortest graph path is typically congested during rush hours, or a sequence of geographical locations along the graph path. These features may be used during request processing to filter travel paths whose features match a condition specified in the request.

[052] In one embodiment, we compute two or more graph paths between some graph vertexes. For example, one graph path from a vertex u to a vertex v could be with at most one bus, and other graph path from the vertex u to the vertex v could be without any subway, but no more than 10 minutes longer than a shortest graph path.

[053] In one embodiment, we compute graph paths that match various filtering conditions. For example, graph paths that have at most one transfer, or at most a specific walk duration.

[054] In one embodiment, we use a routing engine (for example prior art) to compute shortest travel between homes and public transportation stops. Thus, we may sometimes not use the graphs described in the current Section 4.4.1.

[055] The method described in Section 4.4.1 computes a shortest travel from each home to each stopstation, and a shortest travel in the reverse direction. Next, we describe a data structure for efficient storage and processing of shortest travel data used by our method.

4.4.2 Storage of shortest travel

[056] In one embodiment, our method stores travel durations in a vector form. The method sequences stopstation clusters as $s_1, \dots, s_n$, and sequences home clusters as $h_1, \dots, h_m$ in some orders, for example random orders. In one embodiment, these sequences become fixed. For each s_i , the method stores a vector of shortest travel durations from s_i to the home clusters using the home cluster sequence

$$v_i = (t_{i,1}, t_{i,2}, \dots, t_{i,m}),$$

so that the value of the vector v_i at a coordinate j , $v_i[j]$, is $t_{i,j}$ equal to the shortest travel duration from the stopstation cluster s_i to the home cluster h_j . See FIG. 6 for an illustration.

[057] The home cluster sequence simplifies the calculation of a shortest travel duration to *every* home cluster. For example, if a commuter needs to travel from a stopstation cluster s_{i_1} and in
 5 addition from a stopstation cluster s_{i_2} , we simply add vectors v_{i_1} and v_{i_2} coordinate-wise, and the resulting sum contains the total travel duration from both stopstation clusters to each home cluster $h_1, \dots, h_m$.

[058] In one embodiment, the method stores durations of reversed travel, from homes to stopstations, using the same home cluster sequence. That is, the method stores a vector

$$v'_i = (t'_{i,1}, t'_{i,2}, \dots, t'_{i,m}),$$

10 so that $v'_i[j] = t'_{i,j}$ is the shortest travel duration from the home cluster h_j to the stopstation cluster s_i (note the transposition; the vector v'_i is for one stopstation cluster, even though travel is in the opposite direction).

[059] In general, the two vectors for the same stopstation cluster s_i are not equal, $v_i \neq v'_i$, because the minimum travel duration to a home may be different than from a home (travel is
 15 not symmetric in general). However, in one embodiment the method computes and stores just one of the two vectors, and uses it instead of the other, which can save time and space. In other embodiment, the method stores a coordinate-wise weighted average of the two vectors, which may decrease a worst-case error. The weights may be set to 0.5, or may favor travel from, or may favor travel to a specific home cluster, for example based on request frequency.

20 [060] In one embodiment, the method stores travel durations rounded to the closest minute, using one byte of computer memory, represented as an unsigned integer from 0 to 254, with 255 denoting an unknown or too large travel duration. This storage enables efficient vector addition using modern computer hardware support for Vector Operations and Saturation Arithmetic (for example the AVX-512 instruction set, or GPU intrinsics), while keeping the error at a practically
 25 acceptable level of at most half a minute, and covering common travel durations up to over 4 hours. Any other rounding can be used, for example a duration in seconds can be divided by 120, and rounded to an integer, which will represent durations at a 2-minute granularity.

[061] In one embodiment, the method stores a vector not for clusters. For example, some of the $s_1, \dots, s_n$ represent stopstations (not stopstation clusters), or some of the $h_1, \dots, h_m$ represent

homes (not home clusters).

[062] In one embodiment, the method uses other form of a vector to store travel durations. For example, a hash map, or a (coordinate, value) list. These sparse forms may be advantageous when there are many unknown or too large travel durations.

5 [063] In one embodiment, the method stores travel paths using vectors that follow the home cluster sequence $h_1, \dots, h_m$, or using other form of a vector, for example a hash map or a (coordinate, value) list. These travel paths may be used during request processing to filter travel paths whose features match a condition specified in the request.

4.5 Travel duration of one commute path

10 [064] We describe how our method efficiently computes the shortest travel duration of any commute path. Consider any commute path from a home cluster H , through $k \geq 1$ works, back to the home cluster: $H \rightarrow W_1 \rightarrow W_2 \rightarrow \dots \rightarrow W_k \rightarrow H$. We want to find the shortest travel duration of this commute path for every H . This may be expensive, because the number of home clusters may be rather large, for example 500,000 given a clustering radius of a few meters. However, our
15 method introduces a technique that dramatically accelerates the search: the method decomposes any commute path into a home-independent part, and a home-dependent part, as illustrated in FIG. 7.

4.5.1 Middle part: travel $W_1 \rightarrow W_2 \rightarrow \dots \rightarrow W_k$

[065] The method finds the shortest travel duration for travel $W_1 \rightarrow W_2 \rightarrow \dots \rightarrow W_k$ that
20 excludes the home cluster. For each segment $W_i \rightarrow W_{i+1}$, the method queries a routing engine (for example prior art) that computes the shortest travel duration from W_i to W_{i+1} using the transportation system (for example, involving walks, subway and bus rides, transfers, including a direct walk from W_i to W_{i+1}). Then the method adds up the shortest travel durations across the different i . The resulting sum is denoted by

$$\begin{aligned} \text{PathNonHomeDuration}(W_1 \rightarrow W_2 \rightarrow \dots \rightarrow W_k) = \\ \sum_{1 \leq i \leq k-1} \text{RoutingEngineShortestDuration}(W_i \rightarrow W_{i+1}). \end{aligned}$$

We note that only $k - 1$ queries to the routing engine are needed. This number is independent of the number of homes.

[066] In one embodiment, the method uses a departure time, an arrival deadline, or other parts of a request, when querying the routing engine.

5 [067] Next, the method computes the shortest travel duration for the open-jaw part of the commute path: the two segments $H \rightarrow W_1$ and $W_k \rightarrow H$ that involve the home cluster. This computation needs to be especially quick, because its result is needed for every home cluster.

4.5.2 Start part: travel $H \rightarrow W_1$

[068] The method considers two ways of getting from H to W_1 . See FIG. 8 for an illustration.
 10 The first way is a direct walk. When H and W_1 are nearby, a true shortest travel will often be a direct walk. The method thus queries a walk engine (for example prior art) to compute the shortest walk duration $walk(H \rightarrow W_1)$. The second way uses the transportation system. The method finds the stopstation clusters that are within a threshold distance, for example 2000 meters, from W_1 , denoted by a set A . The set A is a subset of $\{s_1, \dots, s_n\}$ of all stopstation clusters (in FIG.
 15 8 the $A = \{s_1, s_2, s_3\}$). The method queries a walk engine to retrieve a shortest walk duration $walk(s_i \rightarrow W_1)$, for each s_i in A . When H and W_1 are not nearby, a true shortest travel will often pass through a stopstation in the set A , and then continue along a corresponding walk. Due to the two-way approach we apply, the resulting travel duration often is a shortest travel duration.

[069] The two-way approach can be applied to each home cluster. We recall that a vector
 20 v'_i contains shortest travel durations to stopstation cluster s_i from consecutive home clusters. Therefore, the method computes a shortest travel duration from a home cluster h_j by retrieving vectors v'_i from the database, and using the following formula

$$a_j = \min \left\{ walk(H_j \rightarrow W_1), \min_{s_i \in A} \{v'_i[j] + walk(s_i \rightarrow W_1)\} \right\}. \quad (\text{Equation 1})$$

[070] In one embodiment, the method computes $walk(H_j \rightarrow W_1)$ only when the distance from home cluster H_j to W_1 is below a threshold, for example 2000 meters, or when a travel duration
 25 from home cluster H_j to W_1 is below a threshold. Let us denote the set of such home clusters by J .

[071] In one embodiment, because of the vector representation of the v' , the method jointly

computes the travel duration from every home cluster, using vector operations on the v' as follows:

$$\begin{aligned}
 &\text{for } j = 1 \text{ to } m \\
 &\quad a_j = \infty \\
 &\text{for all } s_i \in A \\
 &\quad w = \text{walk}(s_i \rightarrow W_1) \\
 &\quad \begin{pmatrix} a_1 \\ a_2 \\ \vdots \\ a_m \end{pmatrix} = \begin{pmatrix} \min\{a_1, v'_i[1] + w\} \\ \min\{a_2, v'_i[2] + w\} \\ \vdots \\ \min\{a_m, v'_i[m] + w\} \end{pmatrix} \\
 &\text{for all } j \in J \\
 &\quad w = \text{walk}(H_j \rightarrow W_1) \\
 &\quad a_j = \min\{a_j, w\}.
 \end{aligned}$$

[072] In one embodiment, the method uses a mathematical formula

$$\min_{s_i \in A} \{ \text{walk}(s_i \rightarrow W_1) + v'_i \},$$

where the “+” operation adds a number to a value at each coordinate of a vector, and the “min” operation computes a minimum value at each coordinate across several vectors.

5 [073] The vector $(a_1, \dots, a_m)$ of travel durations is denoted by

$$\text{PathFromHomeDurations}(\rightarrow W_1) = (a_1, \dots, a_m).$$

[074] An example pseudocode for computing PathFromHomeDurations is illustrated in FIG. 10.

[075] In one embodiment, a uint8 number format is used to store a_j , $v'_i[j]$, or w . In one embodiment, the unit of a number stored is minutes. In one embodiment, a sum $v'_i[j] + w$, for some
 10 j , is performed using Saturation Arithmetic in the range from 0 to a threshold, for example 255. In one embodiment, a summand $v'_i[j]$ is first converted into a wider number format, for example into an fp16 format or a uint16 format, and only then added to w , to avoid an arithmetic overflow

during an addition. In one embodiment, some calculations included in the mathematical formula are performed using at least one instruction for Vector Operations (including tensor operations), for example an AVX-512 instruction or a GPU intrinsic supported by hardware. In one embodiment, some calculations included in the mathematical formula are partitioned, and the partitions are
 5 executed in parallel.

[076] In one embodiment, the method uses a nearest-neighbor data structure, for example a KD-tree, on stopstation cluster locations to quickly compute the set A during request processing.

[077] In one embodiment, the method limits the set A to at most a certain number of stopstation clusters that are nearest W_1 , for example at most 100.

10 [078] In one embodiment, the method limits the set A to stopstation clusters within a walk of at most a certain length to W_1 , for example 2000 meters.

[079] In one embodiment, the method precomputes shortest durations of walks between points within a threshold distance from each home cluster location and the home cluster location, or precomputes shortest durations of walks between points within a threshold distance from
 15 stopstation cluster locations and the stopstation cluster location. Then, during request processing, the method does not query any walk engine, but instead uses the precomputed walk durations.

[080] In one embodiment, a shortest duration of a walk is estimated using a geodesic line that ignores any obstacles. This may speed up the computation of a $walk(s_i \rightarrow W_1)$ at the expense of accuracy.

20 [081] In one embodiment, travel paths are used in Equation 1. For example, if a user request specifies a condition that limits the number of vehicle transfers, we filter s_i and j in the equation: we look up the number of transfers stored in a travel path, and when that number exceeds the limit, we ignore the specific s_i and j in the equation.

[082] In one embodiment, grouping travel durations by the stopstation s_i , just like we did in
 25 v'_i , improves data access performance. Indeed, although the set A depends on a user request that is not known in advance, for each s_i in the set A , predictably a large number of travel durations need to be accessed.

[083] In one embodiment, the vectors v'_i , for a collection of i , are distributed randomly across different processing units. This may decrease the latency of computing the vector $(a_1, \dots, a_m)$,
 30 because the stopstations included in the set A will often be evenly divided across the processing

units.

[084] In one embodiment, the vectors v'_i , for a collection of i , are grouped geographically inside a processing unit. This may increase the throughput of computing the vector $(a_1, \dots, a_m)$, because of a reduced need for data transfer, due to the fact that the set A often consists of stopstations
 5 located near one another.

4.5.3 End part: travel $W_k \rightarrow H$

[085] The computation is analogous, but uses vectors v , not v' . See FIG. 9 for an illustration. The method computes a set B of stopstation clusters within a threshold distance from W_k (in FIG. 9 the $B = \{s_4, s_5\}$). Then the method computes a shortest travel duration to the home cluster h_j

10 as

$$b_j = \min \left\{ \begin{aligned} &walk(W_k \rightarrow H_j), \\ &\min_{s_i \in B} \{walk(W_k \rightarrow s_i) + v_i[j]\} \end{aligned} \right\}. \quad (\text{Equation 2})$$

[086] Similar as before, in one embodiment, because of the vector representation of the v , the method jointly computes a travel duration to every home cluster using vector operations on the v according to Equation 2. In one embodiment, the method uses a mathematical formula

$$\min_{s_i \in B} \{walk(W_k \rightarrow s_i) + v_i\}.$$

[087] The vector $(b_1, \dots, b_m)$ of travel durations is denoted by

$$\text{PathToHomeDurations}(W_k \rightarrow) = (b_1, \dots, b_m).$$

15 [088] The method uses embodiments similar as in Section 4.5.2. For example, in one embodiment, the method uses a vector operation that together computes a part of the vector $(a_1, \dots, a_m)$ and a part of the vector $(b_1, \dots, b_m)$.

4.5.4 Combining start, middle, and end parts

[089] Finally, the method computes the shortest travel duration of the commute path for every
 20 home cluster. The method simply adds the two vectors and shifts values at the coordinates. We denote the resulting vector by

$$\begin{aligned}
&\text{PathDurations}(\rightarrow W_1 \rightarrow W_2 \rightarrow \dots \rightarrow W_k \rightarrow) = \\
&\quad \text{PathFromHomeDurations}(\rightarrow W_1) \\
&\quad + \text{PathToHomeDurations}(W_k \rightarrow) \\
&\quad + \text{PathNonHomeDuration}(W_1 \rightarrow W_2 \rightarrow \dots \rightarrow W_k), \quad (\text{Equation 3})
\end{aligned}$$

where the first “+” is a coordinate-wise addition of vectors, and the second “+” is an addition of a number to a value at every coordinate of a vector.

[090] In one embodiment, the method computes travel durations given specific departure times from, or arrival deadlines to, any geographical location along any commute path.

5 [091] In one embodiment, the method precomputes PathDurations for specific commute paths. For example, the method computes and stores

$$\text{PathFromHomeDurations}(\rightarrow W)$$

and

$$\text{PathToHomeDurations}(W \rightarrow),$$

for every W that is a school. During request processing, the method retrieves precomputed vectors from the database, instead of computing them with Equation 1, Equation 2, and Equation 3.

10 [092] In one embodiment, the value of PathDurations may be changed. Consider a school W that has a zoning requirement that specifies that only certain homes may send children to the school. This can be simply achieved by setting the values at coordinates of a vector $\text{PathFromHomeDurations}(\rightarrow W)$ to infinity for these home clusters that are outside of the zone of the school, and similarly setting values at coordinates of a vector $\text{PathToHomeDurations}(W \rightarrow)$. A
 15 similar change can be made when W is a restaurant with a limited delivery area, or a government office with a limited jurisdiction. Thus changed values may be stored, and retrieved during request processing.

[093] In one embodiment, our method applies an arbitrary “manipulation” function to the values at coordinates of vectors v_i , v'_i , or PathDurations. For example, this can help implement a
 20 policy of a government that reduces the transportation charge for commuting from or to certain parts of a metropolitan area in a special economic zone affected by a disaster. An application can

be performed before request processing (for example a global policy), or during request processing (for example a user-personalized policy).

[094] In one embodiment, the method computes PathDurations for a subset of home clusters. For example, the subset is determined from a condition on home features specified by a user request. In one embodiment, we select the sequence $h_1, \dots, h_m$ of home clusters so that these subsets often reside in a short segment of the sequence, which reduces data transfer.

4.5.5 Smoothing

[095] In one embodiment, the method performs *smoothing* of travel that starts at a work location. This is illustrated in Figure 13 on a case of an end part of travel $W_k \rightarrow H$ from a work location W_k to a home cluster H .

[096] In one embodiment, the method precomputes and stores in a database at least one part of travel (sometimes referred as “at least one partial description of travel”), each including a part of a description of travel from a stopstation cluster to a home cluster (sometimes referred as “first partial description of travel”). We use the term “part” in a broad sense, consistent with an interpretation of the term by a person of ordinary skill in the art. In this invention disclosure, any entity is considered to be a part of itself, so the whole entity is an instance of a part of the entity. The part of travel can be interpreted more generally as: a part of a description of travel starting from any endpoint location that can reach W_k and ending at H . In one embodiment, we precompute and store a part of a description of travel between each stopstation cluster and each home cluster. In one embodiment, we precompute and store a part of a description of shortest travel. In one embodiment, we precompute and store a description of travel (without removing any part). In one embodiment, we decrease the storage space needed to keep a precomputed description of travel, for example by storing its part at a reduced granularity, such as every at least 100 meters, or every at least 10 seconds. In one embodiment, we compress overlapping travel routes, for example when two travel routes pass along the same road segment. In one embodiment, we store a prefix of a description of travel, for example within a distance of 2000 meters, or a travel duration of 2 minutes, from a start of travel, or simply store a start of travel itself. In one embodiment, we do not store the part of a description of travel beyond a prefix, that is, we do not store a tail of travel, such as the tail 1304. In one embodiment, we store a summary of a tail of travel, for example we store a duration of a tail of travel. This allows our method to compute a duration of smoothed

travel from a work location to a home cluster, without storing any other information about a tail of travel. In one embodiment, we combine two or more approaches for determining a part of a description of travel. We use other embodiments that concern precomputing and storing, and that are described in the invention disclosure. Many other ways of precomputing and storing at least
 5 one part of travel will be apparent to those of ordinary skill in the art.

[097] Our method uses precomputation and storage in a broad sense. In one embodiment, the method precomputes fragments of descriptions of travel, and stores them in a database. This can be done in many ways. In one embodiment, the method determines a coarsened graph, which is obtained from a transportation graph using any graph coarsening method known in the
 10 art. In one example, an edge of a coarsened graph is a shortcut graph edge determined using a method of Geisberger, Sanders, Schultes and Delling: Contraction Hierarchies: Faster and Simpler Hierarchical Routing in Road Networks, WEA 2008. In other example, a coarsened graph is obtained using any graph summarization method known in the art, for example, described in a survey paper by Liu, Safavi, Dighe and Koutra: Graph Summarization Methods and Applications:
 15 A Survey, ACM Computing Surveys, Vol. 51(3) 2018. Later, a step of retrieving from a database a precomputed description of travel or a precomputed part of a description of travel between two endpoints, comprises computing a path within the coarsened graph between the two endpoints, and combining at least one fragment of a description of travel along the path, such as combining two shortcuts, or combining two edges within a graph summary. Thus, precomputation and storage is
 20 simply an activity performed before request processing, so as to make the request processing more efficient.

[098] After a request has arrived, comprising a work location W_k , the method retrieves from a database at least one precomputed part P of travel (sometimes referred as “second partial description of travel”), each including a part of a description of travel $s \rightarrow H$ from a stopstation
 25 cluster s (sometimes referred as “second representative”), to a home cluster H (examples: 1302, 1303 and 1304 illustrate a description of travel from the stopstation cluster s to the home cluster H ; and 1301 illustrates a description of travel from the work location W_k to the stopstation cluster s). This can be done in many ways. In one embodiment, the retrieval uses a constraint based on a predetermined threshold, for example a distance of 2000 meters, or a travel duration of 2 minutes.
 30 In one embodiment, the method determines a stopstation cluster s within the predetermined threshold of W_k , and the part P is the stopstation cluster s . In one embodiment, the method

allows the stopstation cluster s to be arbitrarily far from W_k , but: the part P resides within the predetermined threshold of W_k , or the part P includes a location within the predetermined threshold of W_k . In one embodiment, the method retrieves all parts P that satisfy the constraint. In one embodiment, the retrieval uses a constraint based on a predetermined limit, for example
 5 100. In one embodiment, the method retrieves no more parts P than the predetermined limit. In one embodiment, the retrieved parts P are arbitrarily far from W_k , for example with respect to distance or travel duration. In one embodiment, the retrieved parts P are nearest W_k . In one embodiment, the method retrieves maximally many parts P that are nearest W_k . In one embodiment, both constraints are used. In one embodiment, a constraint based on a limit
 10 includes a constraint based on a threshold, for example when the stopstation clusters are sparsely located. In one embodiment, the method determines a distance or a travel duration using: a transportation graph, an approximate shortest path, a path in a coarsened graph, or a geodesic. In one embodiment, the method uses a nearest neighbor data structure to determine parts P whose distances or travel durations are within a threshold from W_k . Many other ways of retrieving the
 15 at least one precomputed part P of travel will be apparent to these of ordinary skill in the art

[099] The method determines a location x included in the part P . This can be done in many ways. In one embodiment, the location x is arbitrarily selected to be within a predetermined threshold of the work location W_k , for example within a distance of 2000 meters, or within a travel duration of 2 minutes. In one embodiment, the location x is a nearest location of the part P ,
 20 with respect to distance or travel duration from the work location W_k . In other embodiment, the location x is a most downstream (following the sequence of travel $s \rightarrow H$) location of the part P . Many other ways of determining the location x will be apparent to these of ordinary skill in the art.

[100] Then the method determines a description of travel $W_k \rightarrow x$ from the work location
 25 W_k to the location x (an example: 1305, then 1303). This can be done in many ways. In one embodiment, the method determines a connector c that belongs to the part P , and that can be reached from the work location W_k . In one embodiment, the connector c is the stopstation cluster s . In other embodiment, the connector c is the location x . In other embodiment, we select the connector c to be nearest the work location W_k , for example by running a shortest path algorithm
 30 (for example, the Dijkstra's algorithm) starting at the work location W_k , and stopping the shortest path algorithm the first time the algorithm visits a location that belongs to the part P . In other

embodiment, we select the connector c to be a furthest from W_k location that is included in the part P , such that the connector c is no further from W_k than a predetermined threshold, for example a distance or a travel duration from W_k to P , plus a predetermined margin, such as a distance of 1000 meters, or a travel duration of 1 minute. The method determines a description of travel $W_k \rightarrow c$ from the work location W_k to the connector c (an example: 1305). Then the method determines a description of travel $c \rightarrow x$ from the connector c to the location x (an example: 1303). In one embodiment, any of the two descriptions of travel is determined using a shortest path algorithm. In one embodiment, we determine a description of travel $W_k \rightarrow x$ using: a transportation graph, an approximate shortest path, a path in a coarsened graph, or a geodesic. Although we may give up an ability to determine a true shortest path, but we gain an ability to efficiently determine a description of travel leading arbitrarily far from W_k . In one embodiment, we precompute and store in a database at least one description of travel $W_k \rightarrow x$, and retrieve from a database a precomputed description of travel $W_k \rightarrow x$ at a time of request processing. In one embodiment, we determine a description of travel $W_k \rightarrow x$ using the road congestion information that is current at a time of request processing. Many other ways of determining a description of travel $W_k \rightarrow x$ will be apparent to these of ordinary skill in the art.

[101] In addition, the method determines a description of travel $x \rightarrow H$ from the location x to the home cluster H (an example: 1304). This can be done in many ways. In one embodiment, the method determines a description of travel $x \rightarrow H$ using the precomputed part of a description of travel $s \rightarrow H$, that gets retrieved from a database at a time of request processing. In one embodiment, the description of travel $s \rightarrow H$ uses an average road congestion information, for example an average based on measurements during a time window within any weekday. In other embodiment, the method determines a description of travel $x \rightarrow H$ by running a shortest path algorithm at a time of request processing. An example advantageous effect is that the road congestion information that is current at a time of request processing may be used when running the shortest path algorithm. In other embodiment, a description of travel $x \rightarrow H$ is determined using: a transportation graph, an approximate shortest path, a path in a coarsened graph, or a geodesic. An example advantageous effect is that less space may suffice for storing at least one part of travel in a database. Many other ways of determining a description of travel $x \rightarrow H$ will be apparent to these of ordinary skill in the art.

[102] Then the method finishes the smoothing of travel. In one embodiment, the method finds

a part P included in the at least one retrieved part P that minimizes a duration of “smoothed” travel $W_k \rightarrow x \rightarrow H$ (an example: 1305, then 1303, then 1304). In other embodiment, instead of using a minimum, the method uses any criterium for computing a description of travel described in the invention disclosure, for example: a filtering condition, a multi-dimensional cost, a departure
 5 time, or an arrival deadline. Thus, the resulting description of “smoothed” travel may not pass via any stopstation cluster s that is a start of a precomputed description of travel (see Section 4.4.2). It is even possible that it is too far to travel from the work location W_k to the stopstation cluster s (think: a river and a bridge far away), in which case 1301 practically does not exist.

[103] Smoothing typically produces a description of travel that has a more natural shape
 10 compared to prior art methods. This benefit can be seen on an example where a stopstation cluster is a few meters North off the work location W_k , but a shortest travel from the work location W_k to the home cluster H actually leaves from the work location W_k going South.

[104] Smoothing typically does not reduce the length of travel much. In one embodiment, we do not perform smoothing. One example advantageous effect occurs when we are interested merely
 15 in a length of travel, instead of a more general description of travel.

[105] The part P used by our smoothing method can be interpreted as a one-dimensional generalization of a zero-dimensional technical feature used by prior art, and called by any of these names: landmarks, portals, hubs, beacons, seeds, transit nodes, intermediate waypoints, vertexes of a coarsened graph, and so on, or representatives — see, for example, a survey paper by Sommer:
 20 Shortest-Path Queries in Static Networks, ACM Computing Surveys, Vol. 46(4) 2014. However, our generalization has a distinctly different technical character. In some sense, the one-dimensional part P allows to spread in space the benefit of a good entry point into the transportation system. Thus, two work locations that may be significantly far from each other, can use the same part P as an entry “point” (actually, a one-dimensional section) into the transportation system, on a
 25 way towards the home cluster H . In one embodiment, the part P is a part of a shortest-path tree rooted at a representative; the part is determined using any embodiment described in this invention disclosure. Thus, the part P may include branches of the shortest-path tree. Our generalization has an effect that is similar to the effect achieved by using an increased number of representatives, without explicitly using that increased number of representatives, and without any additional cost
 30 of preprocessing that would have to be expended, if we were to use a traditional method with that

increased number of representatives.

[106] The sequence of the steps presented above is not limiting. Above, we presented the steps in one sequence: retrieving a part P , determining a location x , determining a description of travel $W_k \rightarrow x$, and determining a description of travel $x \rightarrow H$. This sequence was selected to improve the clarity of presentation. However, in other embodiment, the method executes some of the steps in other sequence, or in parallel. For example, in one embodiment, the method executes a shortest path algorithm starting at W_k , and as locations get visited, the method retrieves the encountered parts P . The method executes the shortest path algorithm until a stop condition, such as: a number of parts P that have been retrieved reaches a predetermined limit, for example 100, or a predetermined threshold distance has been reached, for example a distance of 2000 meters, or a travel duration of 2 minutes. The method selects a location x for each of the retrieved parts P , for example a furthest location reached by the shortest path algorithm within each part P . The shortest-path tree constructed by the shortest path algorithm automatically yields descriptions of travel from W_k to each location x . Then the method determines a description of travel from each location x to the home cluster H . In one embodiment, a step is precomputed, and the precomputed result is retrieved during request processing. For example, we precompute at least one part P for at least one work location W_k . Many other ways of executing the steps will be apparent to those of ordinary skill in the art.

[107] In one embodiment, we select which travel to smooth. Given a home cluster H and a work location W_k , we determine at least one part P . In one embodiment, we perform the selection based on a preset condition that uses a predetermined threshold. For example, we smooth a part P that is within a predetermined threshold, for example 1000 meters, or 1 minute, from the work location W_k . One example advantageous effect of the selection is that a part P that is near W_k and that is not smoothed, is more likely to yield travel with an unnatural shape. In one embodiment, we perform the selection based on a preset condition that compares travel distances or durations. For example, we perform smoothing for a part P that minimizes a duration of travel $W_k \rightarrow x \rightarrow H_j$; for example, we achieve this by: determining stopstation clusters $s_1, \dots, s_n$ following Section 4.5.2, and selecting the part P by identifying a stopstation cluster s_i that minimizes a travel duration $walk(W_k \rightarrow s_i) + v_i[j]$. One example advantageous effect of the selection is that we can focus the effort of smoothing on a part P that is more likely to yield a shortest smoothed travel compared to any other part P . In one embodiment, we perform the selection based on a preset condition

that uses a predetermined angle. For example, we determine an angle between the part P and a geodesic that arrives from the W_k , for example an angle at a stopstation clusters s that is upstream along the part P (illustrated by an angle between 1301 and 1302). When the angle sufficiently deviates from 0 degrees, for example the angle is outside of an interval of $[-40, 40]$ degrees, then we perform smoothing using the part P . In other example, we run a shortest path algorithm starting at W_k , and determine an angle between the shortest-path tree and the part P , at a vertex of the part P that is first visited by the shortest path algorithm. In one embodiment, we continue running the shortest path algorithm, until the angle of a visited vertex falls inside an interval. In one embodiment, we combine two or more selection approaches. One example advantageous effect of the selection is that we can focus our effort on smoothing that will likely yield travel with a more natural shape. Many other ways of selecting which travel to smooth will be apparent to those of ordinary skill in the art. In one embodiment, we finish the smoothing by finding a shortest travel among: at least one “smoothed” travel, and at least one travel computed without smoothing.

[108] In one embodiment, given home clusters $H_1, \dots, H_m$, we perform smoothing for an arbitrary subset of the home clusters. For example, we select a home cluster that is sufficiently close to the work location W_k , such as within 5 kilometers, or 5 minutes. In other embodiment, we perform smoothing for all home clusters $H_1, \dots, H_m$.

[109] In one embodiment, we improve the efficiency of smoothing at least one travel to at least one home cluster. We run a shortest path algorithm starting at the work location W_k . We keep running the shortest path algorithm until a stop condition. One example of the stop condition is: having visited at least a predetermined number of parts P , each including a part of a description of travel leading to one of the home clusters. In other embodiment, we stop after having visited 100 parts P . In other embodiment, we stop after having visited at least one part P , for each of the one or more home clusters. In other embodiment, we stop based on a preset condition that uses a predetermined threshold, or a preset condition that uses a predetermined angle. In one embodiment, we combine two or more stop conditions. Many other stop conditions will be apparent to those of ordinary skill in the art. Then, we determine descriptions of travel from the work location W_k to the locations x , and from the locations x to the home clusters.

[110] In one embodiment, we use any embodiment for computing a description of travel described in the invention disclosure, when we determine: a precomputed part of travel, a description of travel from the work location W_k to a location x , or a description of travel from a

location x to the home cluster H .

[111] A case of smoothing an end part of travel $H \rightarrow W_1$ is symmetric to the case of smoothing a start part of travel described above. Figure 14 illustrates an example embodiment. 1404, 1403 and 1402 illustrate a description of travel $H \rightarrow s$ from a home cluster H to a stopstation cluster s .
 5 1403 and 1402 illustrate a precomputed part Q of travel $H \rightarrow s$. The part Q includes a location x .
 1401 illustrates a description of travel from the stopstation cluster s to the work location W_1 . The method determines a description of travel $x \rightarrow W_1$ from the location x to the work location W_1 (an example: 1403, then 1405). In addition, our method determines a description of travel $H \rightarrow x$ from the home cluster H to the location x (an example: 1404). Finally, our method produces the
 10 following “smoothed” travel: $H \rightarrow x \rightarrow W_1$

[112] In one embodiment, our method constructs a description of travel by performing one smoothing at a work location W_1 , and by performing other smoothing at a work location W_k .

[113] In one embodiment, the method performs smoothing within a middle part of travel $W_1 \rightarrow W_2 \rightarrow \dots \rightarrow W_k$. Figure 15 illustrates an example embodiment of smoothing of both ends
 15 of travel from a work location W_i to a work location W_{i+1} , for any $1 \leq i < k$. Using the earlier embodiments, the method precomputes and stores in a database at least one part of travel, each including a part of a description of travel from a stopstation cluster s to a stopstation cluster s' (sometimes referred as “at least one partial description of travel”). The part of travel can be interpreted more generally as: a part of travel from any endpoint location that can reach s to
 20 any endpoint location that can be reached from s' . After a request has arrived, comprising a work location W_i and a work location W_{i+1} , the method retrieves from a database at least one precomputed part P and at least one precomputed part Q , wherein the part P and the part Q include a part of a description of travel $s_1 \rightarrow s_2$ from a stopstation cluster s_1 to a stopstation cluster s_2 . The at least one precomputed part P is determined using the earlier embodiments for
 25 smoothing travel $W_k \rightarrow H$, whose W_k is interpreted as the work location W_i , and whose H is interpreted a stopstation cluster within the predetermined threshold of the work location W_{i+1} . The at least one precomputed part Q is determined symmetrically: using the earlier embodiments for smoothing travel $H \rightarrow W_1$, whose W_1 is interpreted as the work location W_{i+1} , and whose H is interpreted a stopstation cluster within the predetermined threshold of the work location W_i .
 30 For example, 1501, 1502, 1503, 1504, and 1505 illustrate a description of travel $s_1 \rightarrow s_2$ from the stopstation cluster s_1 to the stopstation cluster s_2 . 1501 and 1502 illustrate the part P . 1504 and

1505 illustrate the part Q . Using the earlier embodiments, the method determines: a location x_1 (sometimes referred as “location x_P ”) included in the part P , a location x_2 (sometimes referred as “location x_Q ”) included in the part Q , a description of travel $W_i \rightarrow x_1$, and a description of travel $x_2 \rightarrow W_{i+1}$. 1511 illustrates a description of travel from the work location W_i to the stopstation cluster s_1 . 1512 illustrates a description of travel from the work location W_i to a connector c_1 . 1521 illustrates a description of travel from a stopstation cluster s_2 to the work location W_{i+1} . 1522 illustrates a description of travel from a connector c_2 to the work location W_{i+1} . The method determines a description of travel $x_1 \rightarrow x_2$, using the earlier embodiments for determining a description of travel $x \rightarrow H$ or a description of travel $H \rightarrow x$. For example, in one embodiment, a precomputed description of travel $x_1 \rightarrow x_2$ is retrieved from a database at a time of request processing. Then the method finishes the smoothing and responds to the request: in one embodiment, among the retrieved parts, the method finds a part P and a part Q that minimize a duration of “smoothed” travel $W_i \rightarrow x_1 \rightarrow x_2 \rightarrow W_{i+1}$ (an example: 1512, 1502, 1503, 1504, and 1522). A response to the request is illustrated in Figure 16 by 1601. In theory, it is possible for a part P to overlap a part Q . However, in practice, the work location W_i is typically far enough from the work location W_{i+1} , and so an overlap can be avoided. We use several ways to handle an overlap. In one embodiment, we select the connector c_1 to appear upstream from the connector c_2 within a description of travel from the stopstation cluster s_1 to the stopstation cluster s_2 . In other embodiment, we compute a description of travel from the work location W_i to the work location W_{i+1} using a shortest path algorithm, when the two work locations are within a predetermined threshold, for example within a distance of 2000 meters, or within a travel duration of 2 minutes. Other ways of handling an overlap will be apparent to one of ordinary skill in the art. In one embodiment, we perform smoothing at W_i or at W_{i+1} , but not at both. We use other embodiments described for the above cases of smoothing of $H \rightarrow W_1$ and smoothing of $W_k \rightarrow H$. Our method of smoothing of $W_i \rightarrow W_{i+1}$ significantly improves over the method described in claim 1 of US 8,417,409 B2.

[114] In one embodiment, our method determines a description of travel that includes a smoothing of two travel routes. Figure 17 illustrates example embodiments. The method uses a predetermined set of representatives, such as the intermediate waypoints of US 2011/0251789 A1. Using the earlier embodiments, the method precomputes and stores in a database at least one part of travel, each including a part of a description of travel between a representative and a

location within the transportation system (sometimes referred as “at least one partial description of travel”). The location is included in an arbitrary preset set of endpoint locations. Embodiments of the location include: a road junction, a turn, a vehicle stop, a source location or a target location included in a request, a representative, a location of a vertex of a transportation graph, and so on. After a request has arrived, comprising a source location *source* and a target location *target*, the method determines at least one representative w included in travel $source \rightarrow w \rightarrow target$. This determination can be achieved, for example, using prior art, such as US 2011/0251789 A1. The method retrieves from a database at least one precomputed part Q of travel, each including a part of a description of travel from a representative to the target location *target*. The retrieval is performed using the earlier embodiments for smoothing travel $W_k \rightarrow H$, whose W_k is interpreted as the representative w , whose H is interpreted as the target location *target*, and whose stopstation clusters are interpreted as the predetermined set of representatives. In a symmetric manner, the method retrieves from a database at least one precomputed part P of travel, each including a part of a description of travel from the source location *source* to a representative. Using the earlier embodiments, the method determines: a location x_P included in the part P , and a location x_Q included in the part Q . Next, the method determines: a description of travel $source \rightarrow x_P$ and a description of travel $x_Q \rightarrow target$, which can be interpreted as the earlier embodiments for determining a description of travel $H \rightarrow x$ and a description of travel $x \rightarrow H$. For example, in one embodiment, a precomputed description of travel $source \rightarrow x_P$ is determined by finding a precomputed description of travel that starts within a predetermined threshold of the *source* and ends at x_P . Similarly, in one embodiment, a precomputed description of travel $x_Q \rightarrow target$ is determined. Then the method determines which parts to smooth. A location x_P and a location x_Q may be arbitrarily far away from each other, for example because the part P and the part Q were determined for two different representatives w . In one embodiment, the method smooths the part P and the part Q when they were determined for the same representative w . In one embodiment, the method uses a constraint bases on a predetermined threshold, as before. For example, provided that the location x_P and the location x_Q are within a predetermined threshold, for example a distance of 2000 meters, or a travel duration of 2 minutes, the method determines a description of travel $x_P \rightarrow x_Q$ using the earlier embodiments for determining a description of travel $W_k \rightarrow x$ or a description of travel $x \rightarrow W_1$. In one embodiment, the method uses a constraint bases on a predetermined limit, as before. For example, the method selects for smoothing at most

100 pairs of x_P and x_Q that are nearest each other. In other embodiment, the method determines a description of travel $x_P \rightarrow x_Q$ using: a transportation graph, an approximate shortest path, a path in a coarsened graph, or a geodesic. Then the method finishes the smoothing and responds to the request: in one embodiment, among the retrieved parts, the method finds a part P and a part Q that minimize a duration of “smoothed” travel $source \rightarrow x_P \rightarrow x_Q \rightarrow target$. This embodiment is illustrated in Figure 17 with: s_3 denoting the source location, s_5 denoting the target location, s_4 denoting the representative w , x_4 denoting the location x_P , and x_5 denoting the location x_Q . We use other embodiments described above. Our method of smoothing significantly improves over the method described in claim 1 of US 2011/0251789 A1.

[115] In one embodiment, our method determines a description of travel that includes at least one smoothing of two travel routes. In one embodiment, the method can be characterized by: (1) determining a sequence of $k \geq 2$ travel routes $s_1 \rightarrow s_2, s_3 \rightarrow s_4, \dots, s_{2k-1} \rightarrow s_{2k}$; (2) selecting at least one pair of the travel routes; and (3) connecting the pair via a smoothed travel $x_P \rightarrow x_Q$ using one of the earlier embodiments. In one embodiment, the method computes a coarsened graph, so that each edge of the coarsened graph corresponds to a graph path in the transportation graph. Using the earlier embodiments, the method precomputes and stores in a database at least one part of travel, each including a part of a description of travel within the transportation system between two vertexes of the coarsened graph (sometimes referred as “at least one partial description of travel”). After a request has arrived, comprising a source location and a target location, the method computes a shortest path C , in the coarsened graph, from a vertex corresponding to a location within a predetermined threshold of the source location to a vertex corresponding to a location within a predetermined threshold of the target location. A vertex of the coarsened graph is an embodiment of an endpoint location. Next, the method performs smoothing, by joining at least one pair of travel routes within the transportation system, that correspond to two edges of the shortest path C , for example by joining travel routes for: each consecutive pair of edges, two edges that reside within the predetermined threshold from each other, or no more pairs of edges than the predetermined limit. For this purpose, using the earlier embodiments, the method retrieves from a database at least one precomputed part of travel. For example, Figure 17 illustrates a travel route $s_3 \rightarrow s_4$ and a travel route $s_4 \rightarrow s_5$ that correspond to two consecutive edges of a coarsened graph; these two travel routes get smoothed by $x_4 \rightarrow c_4 \rightarrow c_5 \rightarrow x_5$. Figure 17 also illustrates smoothing of a travel route $s_1 \rightarrow s_2$ and a travel route $s_3 \rightarrow s_4$ that reside within the

predetermined threshold from each other. In one embodiment, the method performs smoothing at s_1 , in order to smooth travel from the source location. In one embodiment, the method performs smoothing at s_{2k} , in order to smooth travel to the target location. Then the method finishes the smoothing and responds to the request: in one embodiment, the method finds a shortest among
 5 the thus computed descriptions of smoothed travel. In one embodiment, the join operations for two or more pairs are performed in parallel. We use other embodiments described above. Other embodiments of determining a description of travel that includes at least one smoothing of two travel routes will be apparent to one of ordinary skill in the art.

[116] In one embodiment, the smoothing described in the current Section 4.5.5 is applied
 10 to a more general case described in Section 4.9.2. In one embodiment, we modify the method described in the current Section 4.5.5. We use *REPRESENTATIVE_SOURCE* vertexes and *REPRESENTATIVE_TARGET* vertexes, instead of stopstation clusters, and perform other modifications following Section 4.9.2, as will be apparent to one of ordinary skill in the art.

[117] In one embodiment, we use representatives that are locations that frequently occur in
 15 shortest travels; the locations can be obtained, for example, using prior art. There are many known methods for obtaining a desired set of representatives for any transportation system. Some of the methods appear in patent literature. An arguably simplest method proposed by CN104240163A selects an equally spaced geographical grid as representatives. Other method, proposed by US 8,756,014 B2, selects transit stations of a transportation system as representatives. Another
 20 method selects a random subset of locations within a transportation system as representatives. Another method, also proposed by US 8,756,014 B2, first computes the average weight of a shortest-path subtree rooted at a vertex, by randomly sampling a root of a tree; and then selects a small set of vertexes with the highest average weight, for example as small as one percent of the number of vertexes. Other methods appear in non-patent literature. The literature describes
 25 the following relevant technical features: (a) the notion of Reach discussed in Section 3 of Gupta, Kosowski and Viennot: Exploiting Hopsets: Improved Distance Oracles for Graphs of Constant Highway Dimension and Beyond, ICALP 2019; (b) the construction of Hub Sets discussed in Section 3.1 of Kosowski and Viennot: Beyond Highway Dimension: Small Distance Labels Using Tree Skeletons, SODA 2017; (c) the notion of Access Nodes discussed in Section 5.2 of Abraham,
 30 Dellinger, Fiat, Goldberg and Werneck: Highway Dimension and Provably Efficient Shortest Path Algorithms, Journal of the ACM, Vol. 63(5) 2016; (d) the notion of in-labels and out-labels stated

in Definition 3.3 of Cohen, Halperin, Kaplan and Zwick: Reachability and Distance Queries via 2-Hop Labels, SIAM Journal on Computing, Vol. 32(5) 2012; and (e) the notion of Shortest Path Covers discussed in Section 4.1 and 5.4 of Abraham, Delling, Goldberg and Werneck: A Hub-Based Labeling Algorithm for Shortest Paths in Road Networks, Microsoft Research Technical Report / SODA 2010. The above methods provide examples of sets of representatives. The examples are not intended to be exhaustive. Many other examples exist. The reader is referred to a survey paper by Sommer: Shortest-Path Queries in Static Networks, ACM Computing Surveys, Vol. 46(4) 2014, which is an excellent introductory reading for obtaining basic / ordinary knowledge in the technical field of the invention disclosure. In one embodiment, a set of representatives can be characterized as: a subset of the transportation system, such that a number of representatives is at most a size of the transportation system multiplied by a predetermined ratio, and such that there is a representative near an endpoint of a route between two endpoints. In one embodiment, a size of a transportation system corresponds to the number of vertexes of a graph that models the transportation system. In one embodiment, a representative corresponds to a vertex of such a graph. The characterization is found in US 8,756,014 B2, and also in the survey paper by Sommer. The predetermined ratio has a value at most one. See Section 4.4 for example tradeoffs concerning the value. Our method uses any set of representatives, including any set of representatives that is unknown publicly at the time of making this invention disclosure.

[118] In one embodiment, various embodiments are combined.

[119] The steps described in the current Section 4.5.5 are optional. In one embodiment, the method does not perform smoothing. Our method of smoothing can be used for computing routes in general.

[120] We include a glossary of selected phrases that occur in the claims, and example references to the specification. These references are not intended to be exhaustive; other references exist. A reference may be semantic (not verbatim).

phrase	example reference
information	paragraph [195]
part	paragraph [096]
at least one partial description of travel	paragraph [096] , [113]. [114], [115]
database	paragraph [123]

first partial description of travel	paragraph [096]
second partial description of travel	paragraph [098]
second representative	paragraph [098]
location x	paragraph [099]
endpoint location	paragraph [096], [113], [114], [115]
part P	paragraph [113], [114]
part Q	paragraph [113], [114]
location x_P	paragraph [113], [114]
location x_Q	paragraph [113], [114]

[121] Embodiments of the invention include the following methods. Each claimed method is automated. Each claimed method includes no “mental process” (no step of any claimed method is performed in the human mind). Antecedent basis is sometimes tracked with boxes: a term in a claim can later be used as the term.

- 5 1. A method for searching or comparing at least one site using at least one description of travel between the at least one site and at least one place within a transportation system, the method comprising:
 - (a) receiving at least one representative included in the transportation system;
 - (b) receiving a request comprising the at least one place; and
 - 10 (c) responding to the request with information obtained using the at least one description of travel;

the method characterized by:

- (d) determining at least one partial description of travel and storing it in a database, wherein each first partial description of travel comprises a part of a description of travel within the transportation system between a first site included in the at least one site and a first representative included in the at least one representative; and
- 15 (e) computing the at least one description of travel that comprises a description of travel within the transportation system between a start/end site included in the at least one

site and a start/end place included in the at least one place,

wherein the computing comprises:

- i. retrieving from the at least one partial description of travel at least one second partial description of travel, wherein each second partial description of travel comprises a part of a description of travel within the transportation system between the start/end site and a second representative included in the at least one representative;
- ii. determining a description of travel within the transportation system between the start/end place and a location x included in the second partial description of travel; and
- iii. determining a description of travel within the transportation system between the location x and the start/end site.

2. A method for determining a description of travel from a source location to a target location within a transportation system, the method comprising:

- (a) receiving at least one representative included in the transportation system;
- (b) receiving a request comprising the source location and the target location; and
- (c) responding to the request with information obtained using the description of travel;

the method characterized by:

- (d) determining at least one partial description of travel and storing it in a database, wherein each first partial description of travel comprises a part of a description of travel within the transportation system between an endpoint location included in the transportation system and a first representative included in the at least one representative; and
- (e) computing the description of travel using steps comprising:
 - i. determining at least one part P and at least one part Q , wherein the part P and the part Q are included in the transportation system, and wherein:
 - A. the part Q is the target location, and the part P is retrieved from the at least one partial description of travel wherein the endpoint location is the source location,

- B. the part P is the source location, and the part Q is retrieved from the at least one partial description of travel wherein the endpoint location is the target location, or
 - C. the part P is retrieved from the at least one partial description of travel wherein the endpoint location is the source location, and the part Q is retrieved from the at least one partial description of travel wherein the endpoint location is the target location,
- ii. determining a description of travel from a location x_P included in the part P to a location x_Q included in the part Q ,
 - iii. determining a description of travel from the source location to the location x_P , and
 - iv. determining a description of travel from the location x_Q to the target location.

[122] An embodiment is illustrated in Figure 18: 1801 illustrates a source location, 1802 illustrates a target location, 1803 illustrates a part P , 1804 illustrates a part Q , 1805 illustrates a location x_P , 1806 illustrates a location x_Q , 1807 illustrates a description of travel from a location x_P to a location x_Q , 1808 illustrates a description of travel from a source location to a part P , and 1809 illustrates a description of travel from a part Q to a target location.

[123] Our methods use any form of storage medium. For example, a database for storing precomputed descriptions of travel is, in one embodiment, a non-transitory computer-readable storage medium, for example a volatile memory or a non-volatile memory. Embodiments of a non-transitory computer-readable storage medium are well-known in the art, and so need not be described herein. In one embodiment, a method uses a non-transitory computer-readable storage medium, so as to store one or more programs for execution by at least one processor. Examples of a processor include a CPU and a GPU. Embodiments of a processor are well-known in the art, and so need not be described herein. The one or more programs comprise instructions to be executed by the at least one processor, so as to perform at least one step of a method described in the invention disclosure. In one embodiment, the instruction are expressed in any programming language, for example C++, Java or JavaScript. Embodiments of expressing instructions are well-known in the art, and so need not be described herein.

4.6 Search for one home

[124] Our method teaches how to efficiently search for a home using commutes. We present a few example search requests first, before we introduce a general search request for one home.

4.6.1 Weighted sum request

5 [125] Consider a family with one parent going to work located at *geo1* five times per week, and the other parent going to work at *geo2* three times per week. This family wants to find a home with a short weekly travel duration. We can find the weekly travel duration for the family for every home cluster as a weighted sum of two vectors, as follows:

$$\begin{aligned} & 5 \cdot \text{PathDurations}(\rightarrow \text{geo1} \rightarrow) \\ & + 3 \cdot \text{PathDurations}(\rightarrow \text{geo2} \rightarrow). \end{aligned}$$

4.6.2 Minimum request

10 [126] Consider a single mother who works from home, and is sending a child to school. The mother wants to find a home that will be near any school from among schools E . We can find the daily travel duration for the child for every home cluster as a coordinate-wise minimum of vectors, as follows:

$$\min_{e \in E} \text{PathDurations}(\rightarrow e \rightarrow).$$

4.6.3 General search for one home

15 [127] In one embodiment, our method defines a request as any sequence of commute paths $path_1, \dots, path_q$, for any q , and a function $\text{Deriver} : \mathbb{R}^q \rightarrow \mathbb{R}$ that maps any vector of q numbers into a number. The method computes the travel duration vectors (from scratch, by retrieving precomputed vectors, or both)

$$\begin{aligned} \text{PathDurations}(path_1) &= (d_{1,1}, \dots, d_{1,m}), \\ &\dots, \\ \text{PathDurations}(path_q) &= (d_{q,1}, \dots, d_{q,m}), \end{aligned}$$

and applies the function Deriver coordinate-wise to the vectors, to produce a vector

$$\text{RequestDurations}(\text{path}_1, \dots, \text{path}_q, \text{Deriver}) =$$

$$\left(\begin{array}{l} \text{Deriver}(d_{1,1}, \dots, d_{q,1}), \\ \text{Deriver}(d_{1,2}, \dots, d_{q,2}), \\ \dots, \\ \text{Deriver}(d_{1,m}, \dots, d_{q,m}) \end{array} \right)$$

with a “derived” travel duration for each home cluster.

[128] In one embodiment, the function Deriver is a weighted sum of numbers

$$\text{Deriver}(x_1, \dots, x_q) = w_1 \cdot x_1 + \dots + w_q \cdot x_q,$$

for weights $w_1, \dots, w_q$. Weights can be positive or negative. If a travel duration is interpreted as a monetary cost of travel, then a negative weight can be interpreted as a monetary gain, for example if the commuter is expected to benefit from performing the specific commute (e.g., a delivery of a parcel). Weights can represent relative importance of commute paths, for example a commute path of a Chief Executive Officer may have a higher weight than a commute path of a first-line manager.

[129] In one embodiment, the function Deriver is a minimum of numbers

$$\text{Deriver}(x_1, \dots, x_q) = \min_{1 \leq i \leq q} x_i.$$

[130] In one embodiment, the function Deriver is a weighted sum and a minimum

$$\text{Deriver}(x_1, \dots, x_r, x_{r+1}, \dots, x_q) = w_1 \cdot x_1 + \dots + w_r \cdot x_r + \min_{r+1 \leq i \leq q} w_i \cdot x_i.$$

[131] In one embodiment, the function Deriver is a conditional, for example

$$\text{Deriver}(x_1, x_2) = \left[\begin{array}{l} \text{if } x_1 < 30 \text{ then} \\ \quad \text{return } x_1 + x_2 \\ \text{else} \\ \quad \text{return } \infty \end{array} \right],$$

or any algorithm.

4.7 Compare two or more homes

[132] Our method teaches how to efficiently compare multiple homes using commutes. We illustrate the method on a few examples, before we introduce a general comparison request.

4.7.1 Current home

5 [133] Consider a family that currently lives in a home. The family members have specific commute paths to work places, schools, and other places. The family considers a move to other home, and wants to compare the total travel duration for their current home with the total travel duration for other prospective homes. Our method makes such a comparison rather simple.

[134] In one embodiment, the method computes the travel duration for every home cluster,
10 including the home cluster h_j of the current home S of the family

$$\text{RequestDurations} = (q_1, \dots, q_m).$$

The method then responds with a “difference” vector: the method subtracts the value at the j -th coordinate from the value at each coordinate

$$(q_1 - q_j, \dots, q_{j-1} - q_j, 0, q_{j+1} - q_j, \dots, q_m - q_j).$$

A negative value at a coordinate of the “difference” vector indicates that the home corresponding to the coordinate has a shorter travel duration compared to the current home S .

15 [135] In one embodiment, the method uses some commute paths for the current home, but other commute paths for other homes. Thus, the method can help a family evaluate a what-if scenario: “Suppose we change workplaces, and move to some other home. How will the new commute duration compare to our current commute duration?”

4.7.2 General comparison for two or more homes

20 [136] Consider a family where the mother and the father must continue commuting to the same work places (no job change), but the children can change schools. Consider other example of two families: parents and their maternal grandparents. They want to find two homes within around 30 minutes from each other, so that one home is near any hospital, and the other home is near

specific schools and work places. Our method makes comparisons of such homes with other homes rather simple.

[137] In one embodiment, the method computes travel durations PathDurations for home clusters for a range of commute paths

$$\begin{aligned} \text{PathDurations}(\text{path}_1) &= (d_{1,1}, \dots, d_{1,m}), \\ &\dots, \\ \text{PathDurations}(\text{path}_q) &= (d_{q,1}, \dots, d_{q,m}). \end{aligned}$$

5 Then the method applies a “generalized” function Deriver that does not operate coordinate-wise as in Section 4.6.3, but instead operates jointly on all $q \cdot m$ travel durations, and produces a vector of one or more numbers. In other words, the generalized function is

$$\text{Deriver} : \mathbb{R}^{q \cdot m} \rightarrow \mathbb{R}^y,$$

for some y .

4.8 Search-or-compare

10 [138] In one embodiment, the function Deriver is any algorithm, for example a randomized algorithm, that takes any input (for example: the commute paths, shortest travel paths, shortest travel durations, the homes, or any condition specified by a user request), and produces any output (for example: a “top list” containing homes that satisfy the user condition, and that have shortest travels that also satisfy the user condition, sorted by the shortest travel durations).

15 4.9 Variants

[139] Many modifications and variations will be apparent to those of ordinary skill in the art without departing from the scope and spirit of the embodiments. We present of few variants for illustration.

4.9.1 Extending travels

20 [140] In one embodiment, the method first computes incomplete shortest travels, and then extends them to selected home clusters. This trades off additional processing and storage for a

decrease in storage and processing due to fewer home clusters in vectors v_i and v'_i .

[141] The method determines *home connectors* which are some elements of the transportation system. In one embodiment, the home connectors are clusters of homes or clusters of stopstations, which may differ from the clusters discussed in Section 4.4.1. Then the method precomputes and
 5 stores in a database shortest travels between the stopstation clusters and the home connectors, using embodiments similar to these of Section 4.4.

[142] When a request is received, the method computes shortest travels between a work W contained in the request, and the selected home clusters. For this purpose, the method determines stopstation clusters nearby the work W , and retrieves precomputed shortest travels between the
 10 nearby stopstation clusters and the home connectors, to compute shortest travels between the work W and the home connectors, similar to Section 4.5. In one embodiment, this computation uses vector operations similar to these of Sections 4.5.2 and 4.5.3. Then these shortest travels are extended beyond the home connectors, to form shortest travels between the work W and the selected home clusters. This can be achieved by simply finding, for each home cluster, a minimum
 15 extended travel (which may be not just a walk) that passes through any home connector near the home cluster, for example within 2000 meters, and sometimes finding a direct shortest travel between the work W and the home cluster (similar to Sections 4.5.2 and 4.5.3). In one embodiment, in an overlap area the method uses both the home connectors and the vectors v_i and v'_i .

[143] The performance of extending travels is proportional to the total number of home
 20 connectors, and also proportional to the number of home connectors near any selected home cluster. In one embodiment, the extending is used in sparse parts of a metropolitan area, where the two quantities are likely to be low. In one embodiment, the method uses a performance cost function to: (1) determine home connectors, (2) select home clusters, and (3) determine a subset of nearby home connectors for each of the selected home clusters.

25 4.9.2 Commute by car

[144] In one embodiment, we compute travel durations of commute paths by car, rather than by public transportation. This can be achieved by simple modifications to earlier sections.

[145] When preprocessing a transportation system for homes, Section 4.4, instead of taking a GT to be a graph of a public transportation system, we take a GT to be a graph of a car drive
 30 system. This car drive graph can be obtained, for example, from prior art. That graph may

have vertexes representing geographical locations on roads, and edges representing driving a car or turning a car; edges may contain data about drive durations at various times of the day, for example during rush hours.

[146] We extend GT . For each home cluster s , we add vertexes

5 $HOME_CLUSTER_SOURCE_s$

and

$HOME_CLUSTER_TARGET_s$.

We connect each home cluster with some roads that are within a threshold distance, for example 100 meters, by adding a vertex

10 $CONNECTOR_r$,

for at least one r . In one embodiment, the vertex denotes a shortest-distance projection of a home cluster location onto a road that is within a threshold distance; the road may be inside a car parking assigned to the home cluster. The vertex is connected to the two home cluster vertexes by edges labeled Zero with weight 0. The vertex is also connected to vertexes representing endpoints
15 of a segment of the road. Instead of adding clusters of stopstations, which we should not because now GT does not model any subway station nor bus stop, we add vertexes

$REPRESENTATIVE_SOURCE_s'$

and

$REPRESENTATIVE_TARGET_t'$

20 for a collection of s' and t' ; these vertexes represent locations that frequently occur in shortest travels; the locations can be obtained, for example, using prior art. These vertexes may represent clusters. These vertexes get appropriately connected with other vertexes using edges.

[147] We use the extended GT to compute shortest durations of graph paths from each

$REPRESENTATIVE_SOURCE_s'$

25 to every

$HOME_CLUSTER_TARGET_s$,

and from each

$$HOME_CLUSTER_SOURCE_s$$

to every

$$REPRESENTATIVE_TARGET_t',$$

5 possibly using embodiments similar to these described in Section 4.4.1. These travel durations will be stored in a vector form v_i and v'_i as described in Section 4.4.2.

[148] When computing a travel duration of a commute path, we use car rides, instead of using walks as in Section 4.5.

[149] For a start part of a travel, we make the following modifications to Section 4.5.2. Instead
10 of $walk(H_j \rightarrow W_1)$, we use a duration of a car ride from a home cluster H_j to W_1 , for example obtained from prior art; in one embodiment we instead set this duration to infinity when the distance between H_j and W_1 is above a threshold. The set A is a set of vertexes

$$REPRESENTATIVE_TARGET_t'$$

that are within a threshold distance from W_1 . For each vertex $s_i \in A$, we compute a car drive
15 duration from s_i to W_1 , for example using prior art. Then we apply these modifications to Equation 1.

[150] For an end part of a travel, we make similar modifications, but now to Section 4.5.3. Instead of $walk(W_k \rightarrow H_j)$, we use a duration of a car ride from W_k to a home cluster H_j , for example obtained from prior art; in one embodiment we instead set this duration to infinity when
20 the distance between W_k and H_j is above a threshold. The set B is a set of vertexes

$$REPRESENTATIVE_SOURCE_s'$$

that are within a threshold distance from W_k . For each vertex $s_i \in B$, we compute a car drive
duration from W_k to s_i , for example using prior art. Then we apply these modifications to Equation 2.

25 [151] We use other embodiments stated in Section 4.5. For example, we may estimate a travel duration of a car ride using a geodesic.

[152] The search framework of Section 4.6 can be simply extended to allow to restrict what kind of vehicles should be used for each commute path. As a result, for example, we can determine

travel durations of commute paths for every home for a family with one parent going to work located at *geo1 by car* five times per week, and the other parent going to work at *geo2 by public transportation* three times per week.

[153] The comparison framework of Section 4.7 can similarly be extended. As a result, for example, we can search for a new home from which a person will commute by car, and compare with the current travel duration by public transportation. A result of our method may be used as an incentive to purchase a car due to a move.

4.9.3 Commute by other means

[154] In one embodiment, the method uses commute paths by other means, for example: walk only; bicycle only; express buses and walks only; subway and walks only; express buses, subway and walks only; shared vans and walks only; boats; airplanes; and so on. We simply use modifications similar to these described in Section 4.9.2. In one embodiment, the method identifies other home and other means of travel, and recommends these to the user, explaining gains compared with the current home and the current means of travel.

[155] Shortest graph paths in a given graph can be computed without knowing geographical locations of the various vertexes of the graph. Hence, in one embodiment, the method uses a transportation system whose various elements lack geographical location.

[156] The transportation system need not physically move objects. The method merely needs to be able to determine route or route lengths between the elements of the transportation system. Thus, a computer network, that moves data, is an example of a transportation system, comprising these transportation elements: wires/lines (analogous to roads), and hubs/switches (analogous to stops/turns). Many other examples of a transportation system will be apparent to anyone of ordinary skill in the art.

4.9.4 Conditions on travel paths

[157] We can simply realize various filtering conditions on travel paths. For example, we can build a graph G that has no transfer between any subway lines but has walks between home clusters and station clusters. Along with modifications similar to these described in Section 4.9.2, the method will then search-or-compare homes when a commuter may sit all the way during a ride

between home and work (both the home and the work will need to be within a walking distance from a subway stop of one subway line). Similarly, we can build a G and modify shortest graph path algorithms so as to restrict graph paths to at most one transfer, or a subway-bus transfer, or a bus-subway transfer, or a transfer that occurs within a time window, or a type of travel path
 5 commonly used by commuters. Other variations for filtering travel paths fall within the scope of our method, as will be apparent to anyone of ordinary skill in the art.

4.9.5 Conditions on homes

[158] In one embodiment, the method receives a filtering condition on various features of homes, and searches or compares homes whose features match the condition. The features may be: a type
 10 (e.g., a detached home or a high-rise apartment), a transaction type (e.g., sale or rent), a price, a real estate agent commission fee, taxes, a maximum bank loan amount, a number of bedrooms or bathrooms, an area / size of a home, geographical directions of windows, a floor number of a home, the number of floors in the building, or a typical monthly management fee. Simply the method maintains a feature list for each home, and, given a condition, determines the homes whose
 15 features match the condition. The travel durations for these matching homes can be obtained from travel durations for home clusters. Other variations for filtering homes fall within the scope of our method, as will be apparent to anyone of ordinary skill in the art.

4.9.6 Meta search-or-compare

[159] In one embodiment, the method searches or compares homes using an ensemble of prior
 20 search-or-compare requests and responses. This can be viewed as a meta method (a method that uses itself). It is useful, for example, for estimating the value of a proposed new real estate development.

[160] We describe an embodiment of the meta method. We receive some number u of commute paths $path_1, \dots, path_u$. In one embodiment, the commute paths come from a usage log of a
 25 computer service for searching or comparing homes, each commute path may be provided by a different user of the service. For each commute path $path_k$, we compute the travel durations for all home clusters $PathDurations(path_k)$ according to Section 4.5. Then, we apply an aggregator that processes the travel durations. In one embodiment, for each commute path $path_k$, we receive a

vector of weights w_k . In one embodiment, each weight is a probability of a user clicking on a home in the corresponding home cluster, which may be affected by a condition for a real estate property or for a travel path specified in the user request. Then we compute aggregates coordinate-wise, as in the following formula

$$\frac{1}{u} \sum_{k=1}^u w_k \cdot \text{PathDurations}(\text{path}_k).$$

- 5 The j -th aggregate is an average weighted travel duration for home cluster H_j . In one embodiment, the j -th aggregate denotes a normalized contribution of the home cluster to the total travel duration of the entire metropolitan area; in a sense, it is a centrality of the home cluster within the metropolitan area. Because our method rapidly computes $\text{PathDurations}(\text{path}_k)$, we can compute the aggregates quickly. This enables a rapid computation of a desirability of each real estate
10 property with respect to commuting.

[161] Many other embodiments of the meta method will be apparent to those of ordinary skill in the art. In other embodiment, the commute paths are generated from geographical locations of homes and works. In other embodiment, the weights are set to non-zero for some number of home clusters with the lowest travel durations, others are set to zero. In other embodiment, a data
15 scientist evaluates what-if scenarios; because of advantages of our method, these scenarios can be evaluated rapidly. In other embodiment, the aggregator is an arbitrary algorithm, for example one that computes a variance, quantiles, a cumulative distribution function, or a probability of exceeding a threshold.

4.9.7 Commute path involving two or more homes

- 20 [162] In one embodiment, a commute path involves any two or more homes. This is useful, for example, when searching or comparing homes jointly for two or more families.

[163] In one embodiment, H_1 and H_2 are two home clusters. We can compute a travel duration of any commute path $H_1 \rightarrow \text{path} \rightarrow H_2$ that starts at H_1 , but ends at H_2 that may be different from H_1 . We simply use a j corresponding to H_1 in Equation 1, but use a j corresponding to H_2
25 in Equation 2.

[164] In one embodiment, we precompute both vectors $(a_1, \dots, a_m)$ and $(b_1, \dots, b_m)$, and then we find a travel duration for arbitrary two home clusters H_i and H_j by simply adding a value at a coordinate of $(a_1, \dots, a_m)$ corresponding to H_i to a value at a coordinate of $(b_1, \dots, b_m)$

corresponding to H_j , plus the travel duration of the middle part of the *path* (which is independent of H_i and H_j). We remark that a travel duration for any pair of home clusters, from among the m^2 pairs, can be found in constant $O(1)$ time using only linear $O(m)$ space. This is because of an additive structure of the travel duration and the fact that the *path* acts as a separator between
 5 any two home clusters.

[165] In one embodiment, the set of home clusters allowed for H_1 may be different from the set of home clusters allowed for H_2 . For example, a user condition can restrict H_1 to be in the east part of a city, while H_2 to be in the west part of the city. In this case we may use a vector $(b_1, \dots, b_{m'})$ of length m' that is different from m .

10 [166] More generally, the method computes a travel duration for any sequence of $k \geq 2$ home clusters $H_1 \rightarrow path_1 \rightarrow H_2 \rightarrow path_2 \rightarrow \dots \rightarrow path_{k-1} \rightarrow H_k$. Using precomputing, this can be done in $O(k)$ time and $O(m \cdot k)$ space.

[167] In one embodiment, a commute path includes a direct travel between home clusters, $H_x \rightarrow H_y$, as in this example commute path: $H_1 \rightarrow path_1 \rightarrow H_x \rightarrow H_y \rightarrow path_2 \rightarrow H_2$. In this example,
 15 our method partitions the computation into three parts: (1) compute a travel duration of $H_1 \rightarrow path_1 \rightarrow H_x$, (2) compute a travel duration of $H_x \rightarrow H_y$, and (3) compute a travel duration of $H_y \rightarrow path_2 \rightarrow H_2$. The parts (1) and (3) can be performed using the method described earlier in the current Section 4.9.7. The part (2) can be performed similarly, conceptually treating either H_x or H_y as a work. For example, using the vectors v'_i at the coordinate corresponding to H_x , we obtain
 20 travel durations from the home cluster H_x to every stopstation cluster; these can be retrieved as a column of the precomputed vectors v'_i , $1 \leq i \leq n$. Then we extend from each stopstation cluster within a threshold distance from H_y to form a complete travel to H_y , similar to Section 4.5.2, see for example FIG. 8 and Equation 1. Because of precomputing due to our method, the (3) can be performed in time proportional to the number of the stopstation clusters within the threshold
 25 distance. In other embodiment, because the home clusters are known in advance, we precompute a travel duration between every pair of home clusters, and then simply retrieve a travel duration in constant $O(1)$ time. Naturally, a similar partitioning applies to any commute path that contains one or more direct travel between home clusters.

4.9.8 Space exploration

[168] Our method introduces high-performance space exploration algorithms. Consider a case when a commuter travels along a specific commute path $H_1 \rightarrow path \rightarrow H_2$ which has a specific cost; the cost function depends on the two homes H_1 and H_2 . The commuter seeks homes H_1 and H_2 that minimize the cost. This is useful, for example, when a family is willing to change work locations (H_1 and H_2) of the two parents *at the same time*, while minimizing the total travel duration from the current home location (inside *path*) of the family, aka the two-body problem.

[169] One embodiment of space exploration is a gradient descent algorithm. The cost function can be any differentiable function; for example, given a fixed *path*, the function takes two homes H_1 and H_2 as an input, and returns as an output a distance between the two homes H_1 and H_2 multiplied by a travel duration of a commute path $H_1 \rightarrow path \rightarrow H_2$ (this function can be appropriately extended outside of the discrete domain of home pairs, for example via geographical location of homes and extrapolation). At a given step of the gradient descent algorithm, the algorithm computes a gradient of the cost function given a pair of homes, and then picks a pair of homes along the direction of the gradient for the next step of the algorithm. The gradient can be computed using values of the cost function, for example using the two-point formula that needs two values. In one embodiment, the output of the gradient descent algorithm is a pair of homes where the gradient has a small enough norm.

[170] In one embodiment, we precompute the vectors $(a_1, \dots, a_m)$ and $(b_1, \dots, b_m)$ consuming linear $O(m)$ space, and at each step our method computes the gradient in constant $O(1)$ time. Thus, the gradient descent algorithm can often make rapid progress.

[171] In one embodiment, we do not precompute $(a_1, \dots, a_m)$ and $(b_1, \dots, b_m)$, but instead produce the needed values per Equation 1 and Equation 2 on-demand. Such approach may be used when the number of homes is so large that precomputing $(a_1, \dots, a_m)$ or $(b_1, \dots, b_m)$ is infeasible.

[172] In one embodiment, a space exploration algorithm may have constraints on homes. For example, if we require a round-trip commute to one home, a constraint will be $H_1 = H_2$. We may restrict a geographical region allowed for H_1 , or allowed for H_2 .

[173] In one embodiment, a space exploration algorithm uses a commute path involving only one home: $H_1 \rightarrow path$ or $path \rightarrow H_1$. In other embodiment, a space exploration algorithm uses a

commute path involving two or more homes, as in Section 4.9.7.

[174] In one embodiment, a cost function depends on homes and works. For example, such cost function can be used when a person seeks to minimize a blend of expended time and money. In one embodiment, the cost function is: the travel duration between H and W , plus the monetary cost of renting H , minus the salary paid by W , possibly with weights that represent relative importance of the constituents. In one embodiment, because of advantages of our method, a gradient descent can efficiently be used to search for homes and works under such cost function.

4.10 General case

[175] In one embodiment, the words “home” and “work” have an arbitrary semantic meaning. For example, consider a case where a person is looking for work that is located near the current home of the person. The person does not want to move to a different house, but merely wants to find a job closer to the current home. In this case, our method can be simply applied. Given a range of sites where people work throughout a metropolitan area (for example various offices, factories, etc.), the method computes travel durations between every work site cluster, and each stopstation cluster. Thus, the method can be viewed as a “work search-or-compare using commute” method. In one embodiment, the method searches or compares work sites based on a user-specified job type, or a salary range, and based on a travel duration from the current home of the user. As other example, consider a corporation that wants to move its headquarters to a new location. Our method can be used to compute the total “corporate travel duration” for every new location of the headquarters throughout the metropolitan area. Thus, the corporation may determine how each new location will affect commutes of the employees. A new location may be selected, for example, so that: (1) the worst-case commute duration is limited, and (2) the average commute duration is low; thus satisfying individual and social objectives.

[176] Our description so far mainly talked about a travel duration as the search or compare objective. However, the method can use any other objective, for example: a monetary cost of travel; a metric distance; specific features or attributes of travel paths, for example: the number of transfers, or a walking distance; or features of homes, for example: a price, a size, or a type. In one embodiment, this can be simply achieved by building graphs and setting edge weights appropriately. Various objectives may be combined into a multi-objective optimization search

based on a multi-dimensional cost, for example to search for a home that minimizes a travel duration that is penalized by the monetary cost of travel.

[177] In general, the method uses arbitrary *sites* $S_1, \dots, S_m$ (a site was called a home in earlier sections) and arbitrary *places* $P_1, \dots, P_k$ (a place was called a work in earlier sections),
 5 and the method searches or compares the sites $S_1, \dots, S_m$ using routes or route lengths (in earlier sections: a route was called a description of travel, and a route length was called a length of travel) that originate or end at some of the sites and visit some of the places, as specified by arbitrary route specifications (a route specification was called a commute path in earlier sections) involving the sites and the places. The searching or comparing may use any variant described in
 10 Section 4.9. Thus, the method may respond with routes or route lengths, or their representation. Information computed by one embodiment of the method can be used, recursively, as an input to any embodiment of the method.

4.11 Computer system

[178] One of the embodiments of the invention is a computer system that searches or compares
 15 real estate properties using commutes. We illustrate an embodiment of the computer system in FIG. 11.

[179] We use the term “module” in our description. It is known in the art that the term means a computer (sub)system that provides some specific functionality. Our choice of partitioning the computer system into the specific modules is exemplary, not mandatory. Those of ordinary skill
 20 in the art will notice that the system can be organized into modules in other manner without departing from the scope of the invention.

[180] In one embodiment, each travel along any commute path has a specific departure time.

[181] One module (1101) of the system reads data about the transportation system from a data source (1102), and constructs the graph G. During the construction, the module retrieves
 25 data about homes from a real estate data source (1103), and retrieves from other data source (1104) shortest walks between and near home clusters and stopstation clusters. The graph G contains timing data about vehicles. The module outputs the graph without any

HOME_CLUSTER_SOURCE_s

vertexes (1105), and also outputs the graph without any

HOME_CLUSTER_TARGET_s

vertexes but with reversed edges (1106). The module also builds a nearest neighbor data structure (1107) that can find stopstations within a threshold distance from any given geographical location, and precomputes shortest walks near home clusters and stopstation clusters (1108).

5 [182] In the meantime, other module (1109) of the system reads the two graphs, and computes shortest graph paths. The module considers a range of times during the day, in one embodiment every 5 minutes. For each departure time, the module generates one table (1110) with shortest travel durations from stopstation clusters to home clusters using (1105), and the other table (1111) with shortest travel durations from home clusters to stopstation clusters using (1106). In one
10 embodiment, each travel duration, rounded to the nearest minute, is stored as the uint8_t type of the C++ programming language, with the maximum value of 255 reserved to denote an unknown or too large travel duration. In one embodiment, the tables are laid out on HDD disks in the row-major order. In one embodiment, the system uses a cache hierarchy involving HDD disks, SSD disks, and the main memory. In one embodiment, the tables or their parts are compressed
15 using any compression algorithm, for example delta compression. We observe that for any pair comprising a home cluster and a stopstation cluster, travel durations are often similar during a period of time. This similarity often also holds in a neighborhood of the pair. In one embodiment, we select the sequence $h_1, \dots, h_m$ of home clusters so that any h_i and h_{i+1} that are adjacent in the sequence often are nearby home clusters, or select the sequence $s_1, \dots, s_n$ of stopstation clusters
20 so that any s_i and s_{i+1} that are adjacent in the sequence often are nearby stopstation clusters.

[183] The modules (1101) and (1109) operate continuously. As a result, the system maintains fresh data about travel durations given departure times.

[184] Concurrently, a path durations module (1112) computes PathDurations. Given a commute path with departure times, the module queries (1113) for any relevant PathDurations
25 that has already been precomputed. Any missing one is computed from scratch: the module queries a navigation data source (1114) to compute the PathNonHomeDuration of the part of the commute path that does not involve any home. The module also computes the travel durations PathFromHomeDurations and PathToHomeDurations that involve homes, by querying the nearest stopstations (1107), walks (1108), and home travel duration vectors (1110 and 1111) at the

departure times.

[185] Concurrently, the request processing module (1115) searches or compares real estate properties. Any request (1116) contains commute paths including geographical locations along the commute paths and departure times, and a Deriver. When a request is received from a user, the
5 module retrieves PathDurations from the path durations module (1112), applies the Deriver, and responds to the user with information representing an output of the Deriver (1117).

[186] Aspects of the invention may take form of a hardware embodiment, a software embodiment, or a combination of the two. Steps of the invention, for example blocks of any flowchart, may be executed out of order, partially concurrently or served from a cache,
10 depending on functionality or optimization. Aspects may take form of a sequential system, or parallel/distributed system, where each component embodies some aspect, possibly redundantly with other components, and components may communicate, for example using a network of any kind. The invention is not described with reference to any specific programming language. A computer program carrying out operations for aspects of the invention may be written in any
15 programming language, for example C++, Java, or JavaScript. Any program may execute on an arbitrary hardware platform, for example a Central Processing Unit (CPU) or a Graphics Processing Unit (GPU), and associated memory or storage devices. A program may execute aspects of the invention on one or more software platforms, including, but not limited to: a smartphone running Android or iOS operating systems, or a web browser, for example Firefox,
20 Chrome, Internet Explorer, or Safari.

4.12 Computer service

[187] One of the embodiments of the invention is a computer service for searching or comparing real estate properties using commutes. The service is available to users through a user-accessed device, for example a smartphone application or a webpage. It will be obvious to anyone of ordinary
25 skill in the art that the invention is not limited to these devices. It will also be obvious that the presentation of the service in our drawings can be modified (for example by rearranging, resizing, changing colors, shape, adding or removing components) without departing from the scope of the invention.

[188] In one embodiment, the service is accessed through a smartphone application. See FIG.

12 for an illustration. The user inputs a request. In one embodiment, the request includes:

- the desired features of a real estate property (1201), for example “3 bedroom, tall building, high floor”;
- the commute paths (1202, 1203, “schools Townsend or Jericho High” 1204), departure times
5 (1205, 1206), and the frequency of each commute path (1207, 1208); and
- the geographical location of the current home (1209) of the user.

[189] In response, the service returns information representing travel durations. For example, the service renders geographical locations of real estate properties that match the user request (1210). The service renders how the travel durations from these real estate properties compare to
10 the travel duration from the current home. The service renders a summary about each matching real estate property, for example its price. The real estate properties may be stacked on a 2D map, so that a real estate property with a lower travel duration appears above these real estate properties with higher travel durations. When there is clutter on a map, the service may instead render a cluster of real estate properties, the size of which is correlated with the number of real
15 estate properties in the cluster. The cluster may display a summary, for example the number, or typical features of the real estate properties in the cluster.

[190] In one embodiment, the service renders real estate properties with the shortest travel durations. Summaries of the real estate properties are also rendered (1211). The real estate properties may be sorted by the travel duration.

20 [191] In one embodiment, the service renders a “heatmap” that uses a color to indicate a travel duration from each region of a metropolitan area, for example a minimum for the region, see FIG. 1 for an illustration. A heatmap may depict a difference between a travel duration for any home and the travel duration for the current home of the user. In one embodiment, a heatmap renders only the real estate properties that match the user-desired features of a real estate property.

25 [192] In one embodiment, the service renders a histogram of travel durations (1212). The histogram has the travel duration on one axis, and on the other axis the fraction of real estate properties that yield this travel duration. In one embodiment, a histogram renders only the real estate properties that match the user-desired features of a real estate property. In one embodiment, the user can scroll (1213) through a histogram to any section of the histogram, and the service

renders results for travel durations of this specific section. In one embodiment, the method uses other form of a histogram, for example a pie chart.

[193] In one embodiment, the user may restrict the travel paths, for example by limiting a total walk duration, a number of transfers, etc., for example using range sliders (1214).

5 [194] In one embodiment, the service renders a summary of a travel path for a real estate property.

[195] In one embodiment, the service responds to the user with at least one of:

- (a) a geographical location of: a place, or a site, rendered on a map;;
- (b) a departure time or an arrival time for: a place, or a site;;
- 10 (c) a summary of a site; a summary may include at least one of: a name or an address of the site, a price of the site, or a size of the site;;
- (d) a summary of sites that form a cluster of nearby sites;;
- (e) a rendering of the cluster of nearby sites, a size of the rendering being correlated with a number of sites in the cluster of nearby sites;;
- 15 (f) a stacking of sites in a z-index order of a route length, shorter route length higher up the stacking, rendered on a map;;
- (g) information about a route or route length; information may include at least one of: (i) a route: length, duration, monetary cost, speed, or wait duration; (ii) a name of: a vehicle, a vehicle road, a walking path, a stop, a turn, or a transit station of a public transportation vehicle;
- 20 or (iii) a geographical location of: a monetary cost, a speed, a wait duration, a vehicle, a vehicle road, a walking path, a stop, a turn, or a transit station of a public transportation vehicle;;
- (h) a histogram of a route length across the sites; or a heatmap of a route length across the sites rendered on a map;;
- 25 (i) a minimum route length; a histogram of a minimum route length across the sites; or a heatmap of a minimum route length across the sites rendered on a map;;

- (j) a weighted route length; a histogram of a weighted route length across the sites; or a heatmap of a weighted route length across the sites rendered on a map;;
- (k) a difference route length; a histogram of a difference route length across the sites; or a heatmap of a difference route length across the sites rendered on a map;;
- 5 (l) a rendering of an output returned by the Deriver;;
- (m) a site that minimizes a route length across the sites, under a limit on a monetary cost of travel;;
- (n) one of the above items (a) to (m) constrained by a condition specified by a user;; or
- (o) a top list of sites, from among these sites, routes, and route lengths that satisfy the condition,
10 sorted by a route length.

4.13 Note about claims

[196] Those of ordinary skill in the art shall notice that various modifications may be made, and substitutions may be made with essentially equivalents, without departing from the scope of the present invention. Besides, a specific situation may be adapted to the teachings of the
15 invention without departing from its scope. Therefore, despite the fact that the invention has been described with reference to the disclosed embodiments, the invention shall not be restricted to these embodiments. Rather, the invention will include all embodiments that fall within the scope of the appended claims.

4.14 Glossary

20 [197] We include a glossary of selected phrases that occur in the claims, and example references to the specification. These references are not intended to be exhaustive; other references exist.

phrase	example reference
method	paragraph [003]
searching or comparing	paragraph [002]
site	paragraph [177]

route or route length	paragraphs [013] , [156]
place	paragraph [177]
transportation system	paragraph [022]
representative	paragraphs [022] , [146]
storing	paragraph [022]
database	paragraph [022]
precomputed	Section 4.4
description of travel	paragraph [020]
receiving	paragraph [003]
request	paragraph [022]
start place	Section 4.5.2
end place	Section 4.5.3
computing	paragraph [089]
nearby representative	paragraphs [068] , [149] , [150]
length of travel	paragraph [020]
threshold	paragraph [035]
retrieving	paragraph [022]
not nearby site	paragraph [068]
responding	paragraph [003]
information	paragraph [003]
representation	paragraph [177]
real estate property / commute destination	paragraph [002]
departure time / arrival deadline	paragraphs [066] , [090]
place P_i , start place P_1 , end place P_k	Definition 1
first site, last site	Section 4.9.7
involves	Section 4.9.7
site-place route or route length	paragraph [091]
cluster site / cluster representative	paragraphs [035] , [036]
site connector	paragraph [141]

compression	paragraph [182]
manipulation function	paragraph [093]
zone	paragraph [092]
transportation elements	paragraph [156]
segment length	paragraph [033]
graph / graph vertexes	paragraph [033]
transportation element vertexes	Section 4.9.3
site vertex / representative vertex	paragraph [146]
graph edge	paragraph [033]
graph edge source vertex / graph edge target vertex / graph edge weight	paragraph [033]
graph path or graph path length	paragraph [033]
time	paragraph [042]
algorithm	paragraph [033]
reversed	paragraph [040]
coordinate	paragraph [056]
vector v , value $v[i]$	paragraph [056]
vector v' , value $v'[i]$	paragraph [058]
laying out	paragraph [182]
list / hash map	paragraphs [062] , [063]
weight w_i	paragraph [059]
lower bound lb_i , upper bound ub_i , scaling factor sf_i	paragraph [060]
+ operation, min operation	paragraph [072]
first mathematical formula / second mathematical formula	paragraphs [072] , [086]
partitions	paragraph [075]
route specification L_i	paragraphs [028] , [177]
deriver	paragraph [138]
minimum route length	Section 4.6.2
weights $w_1, \dots, w_q$	paragraph [128]

weighted route length	Section 4.6.1
site S	paragraph [134]
difference route length	Section 4.7.1
condition	Sections 4.9.4 , 4.9.5
filtering	paragraph [006] , Sections 4.9.4 , 4.9.5
aggregator	paragraph [160]
centrality	paragraph [160]
cost function	paragraph [169]
exploration algorithm	Section 4.9.8
differentiable function	paragraph [169]
gradient descent	paragraph [169]
multi-objective optimization / multi-dimensional cost	paragraphs [050] , [176]
computer system	Section 4.11
apparatus	Section 4.12
summary / rendering / stacking	paragraph [195]
heatmap	paragraph [191]
histogram	paragraph [192]
top list	paragraph [138]

CLAIMS

1. A [method] for searching or comparing [at least one site] using [at least one description of travel] between the [at least one site] and [at least one place] within a [transportation system], the [method] comprising:
 - 5 (a) receiving [at least one representative] included in the [transportation system];
 - (b) receiving a [request] comprising the [at least one place]; and
 - (c) responding to the [request] with information obtained using the [at least one description of travel];
 the [method] characterized by:
 - 10 (d) determining [at least one partial description of travel] and storing it in a database, wherein each first partial description of travel comprises a part of a description of travel within the [transportation system] between a first site included in the [at least one site] and a first representative included in the [at least one representative]; and
 - (e) [computing] the [at least one description of travel] that comprises a description of travel
 - 15 within the [transportation system] between a [start/end site] included in the [at least one site] and a [start/end place] included in the [at least one place], wherein the [computing] comprises:
 - i. retrieving from the [at least one partial description of travel] at least one [second partial description of travel], wherein each [second partial description of travel]
 - 20 comprises a part of a description of travel within the [transportation system] between the [start/end site] and a second representative included in the [at least one representative];
 - ii. determining a description of travel within the [transportation system] between the [start/end place] and a [location x] included in the [second partial description of travel] and
 - 25 iii. determining a description of travel within the [transportation system] between the [location x] and the [start/end site].
2. A [method] for determining a [description of travel] from a [source location] to a [target location] within a [transportation system], the [method] comprising:
 - (a) receiving [at least one representative] included in the [transportation system];
 - 30 (b) receiving a [request] comprising the [source location] and the [target location]; and
 - (c) responding to the [request] with information obtained using the [description of travel];

the [method] characterized by:

- (d) determining [at least one partial description of travel] and storing it in a database, wherein each first partial description of travel comprises a part of a description of travel within the [transportation system] between an [endpoint location] included in the [transportation system] and a first representative included in the [at least one representative]; and
- (e) computing the [description of travel] using steps comprising:
 - i. determining at least one [part P] and at least one [part Q], wherein the [part P] and the [part Q] are included in the [transportation system], and wherein:
 - A. the [part Q] is the [target location], and the [part P] is retrieved from the [at least one partial description of travel] wherein the [endpoint location] is the [source location];
 - B. the [part P] is the [source location], and the [part Q] is retrieved from the [at least one partial description of travel] wherein the [endpoint location] is the [target location], or
 - C. the [part P] is retrieved from the [at least one partial description of travel] wherein the [endpoint location] is the [source location], and the [part Q] is retrieved from the [at least one partial description of travel] wherein the [endpoint location] is the [target location];
 - ii. determining a description of travel from a [location x_P] included in the [part P] to a [location x_Q] included in the [part Q];
 - iii. determining a description of travel from the [source location] to the [location x_P], and
 - iv. determining a description of travel from the [location x_Q] to the [target location].