ROYAL INSTITUTE OF TECHNOLOGY

Scalable Selective Traffic Congestion Notification

Győző Gidófalvi
Division of Geoinformatics
Deptartment of Urban Planning and Environment
KTH Royal Institution of Technology, Sweden
gyozo@kth.se

Outline

- Introduction
- Related work
- Method
  - Grid-based directional flow and mobility statistics
  - Directional congestion detection
  - Directional congestion notification
  - SQL-based implementation
- Empirical evaluations
  - Accuracy assessment
  - Scalability assessment
- Conclusion and future work

Introduction

- Congestion is a serious problem
  - Economic losses and quality of life degradation that result from increased and unpredictable travel times
  - Increased level of carbon footprint that idling vehicles leave behind
  - Increased number of traffic accidents that are direct results of stress and fatigue of drivers that are stuck in congestion

- Road network expansion is not a sustainable solution
- Instead, utilize increasingly available Floating Car Data (FCD) to: monitor → understand → control movement and congestion

Modern Traffic Prediction and Management System (TPMS)

- Motivated by:
  - Widespread adoption of online GPS-based on-board navigation systems and location-aware mobile devices
  - Movement of an individual contains a high degree of regularity
- Use vehicle movement data as follows:
  - Vehicles periodically send their location (and speed) to TPMS
  - TPMS extracts traffic / mobility patterns from the submitted information
  - TPMS uses traffic / mobility patterns + current / recent historical locations (and speeds) of the vehicles for:
    - Short-term traffic prediction and management:
      - Predict near-future locations of vehicles and near-future traffic conditions
      - Inform the relevant vehicles in case of an (actual / predicted) event
      - Suggest how and which vehicles to re-route in case of an event
    - Long-term traffic and transport planning

Approach, Unique Features, and Contributions

- Use a data-driven approach and a directional grid-based, time-inhomogeneous, Markov jump process model for the detection congestion and the selective dissemination of this congestion information to vehicles
- Unique features
  - Grid-based model: no need to road network information and can be easily scaled to any geographical level of detail
  - Markov jump process: direct estimation of future location, not prone to error propagation
  - Representation flow and movement direction on the grid
  - Novel congestion definition
  - Simple, scalable, portable SQL-based implementation
  - Relevant performance evaluations

Outline

- Introduction
- Related work
- Method
  - Grid-based directional flow and mobility statistics
  - Directional congestion detection
  - Directional congestion notification
  - SQL-based implementation
- Empirical evaluations
  - Accuracy assessment
  - Scalability assessment
- Conclusion and future work

Related Work: Trajectory Prediction

- Prediction model
  - Markov model
  - Sequential rule / trajectory pattern
- Model basis / generality
  - General model for all objects
  - Type-base model for similar (type of) objects
  - Specific model for each individual object
- Definition of Regions Of Interest (ROI) for prediction
  - Application specific ROIs (road segments, network cells, sensors, etc.)
  - Density-based ROIs
  - Grid-based ROIs
- Prediction provision
  - Sequential spatial prediction (location of next ROI)
  - Spatio-temporal prediction
- Additional movement assumptions or models: YES / NO

Outline

- Introduction
- Related work
- Method
  - Grid-based directional flow and mobility statistics
  - Directional congestion detection
  - Directional congestion notification
  - SQL-based implementation
- Empirical evaluations
  - Accuracy assessment
  - Scalability assessment
- Conclusion and future work

Method Outline

1. Map the directional flow / movement of objects to the grid-based framework.
2. Form tumbling windows over the mapped input stream and treat them as ***temporal analysis windows***.
3. Extract ***Current Directional Flow Statistics (CDFS)*** and ***Current Directional Mobility Statistics (CDMS)*** from the ***Recent Trajectories (RT)*** that are within the current tumbling / temporal analysis window.
4. Incrementally summarize the CDFS / CDMS into ***Historical Directional Flow / Mobility Statistics (HDFS / HDMS)*** for different *temporal domain projections*.
5. Detect a grid cell *g* to be congested from a particular direction *dir* if the current mean speed of vehicles that have entered the grid cell *g* from the direction *dir* is significantly and substantially below the normal according to the temporally relevant HDFS.
6. Notify an object *o* about a detected ***directional congestion*** *(g, dir)* if, based on HDMS and the current position and movement direction of *o*, the likelihood that *o* will enter the grid cell *g* from the direction *dir* during the ***prediction horizon*** is greater or equal than a user / system defined ***minimum notification probability threshold*** *min_prob*.

Grid-based Directional Flow and Mobility Statistics

- Directional flow and movement: grid cell and its immediate 8 neighbors
- Directional flow statistics for a grid cell-direction combination (*g*, *dir*):
  - # of objects in (*g*, *dir*)
  - Average speed of objects in (*g*, *dir*)
  - Standard deviation of speeds of objects in (*g*, *dir*)
- Directional mobility statistics for a pair of not necessarily neighboring source and destination grid cell-direction combinations (g_s, dir_s) and (g_d, dir_d):
  - # of objects that move from (g_s, dir_s) to (g_d, dir_d)
  - # of objects that move from (g_s, dir_s) to any directional grid cell

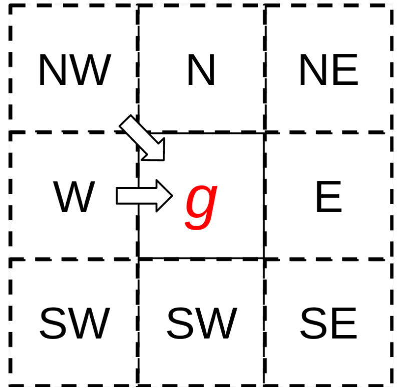

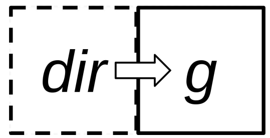

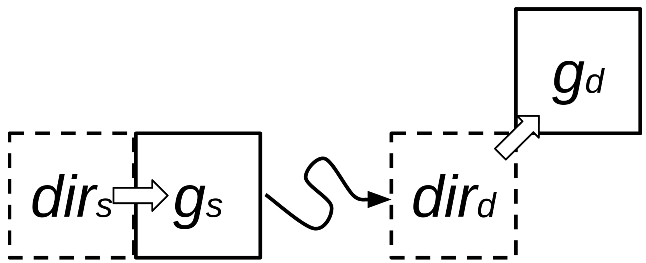

Directional Congestion Detection

- Define a grid cell-direction combination (*g*, *dir*) as a *directional congestion* based on the current $(\dot{n}, \dot{\mu}, \dot{\sigma})$ and historical $(\bar{n}, \bar{\mu}, \bar{\sigma})$ directional flow statistics if the following four criteria are satisfied:

1. Sample size criterion: $\dot{n} >= min\_veh$
2. Sample dispersion criterion: $\dot{\sigma}/\dot{\mu} < max\_cv$
3. Statistical power criterion: $(\dot{\mu} - \bar{\mu})/(\bar{\sigma}/\sqrt{\dot{n}}) < max\_z$
4. Speed difference criterion: $(\dot{\mu} - \bar{\mu})/\bar{\mu} < max\_relspddiff$

Directional Congestion Notification

- Notify an object that has currently entered grid cell g_s from direction dir_s about a direction congestion (g_d, dir_d) according to one of two criteria:

- *Mobility Statistics Criterion* (*MSC*):

$$P((g_d, dir_d)|(g_s, dir_s)) = \frac{n_{(g_s, dir_s) \to (g_d, dir_d)}}{n_{(g_s, dir_s) \to (\cdot, \cdot)}} \geq min\_prob$$

- *Linear Movement Criterion* (*LMC*):

$$\cos(\alpha) \geq min\_cos \wedge dist(g_s, g_d) \leq max\_r$$

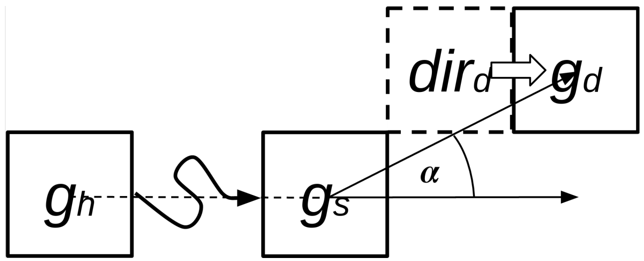

SQL: Schema

- Five database tables:

```
RT = <oid, seqnr, dgid, spd>
CDFS = <dgid, nr, mu, sig, nr_suc>
CDMS = <dst_dgid, src_dgid, nr_src2dst>
HDFS = <dgid, nr, mu, sig, nr_suc>
HDMS = <dst_dgid, src_dgid, nr_src2dst>
```

- Directional grid ID `dgid` columns contain an integer concatenation of grid coordinates and direction (`gx`, `gy`, `dir`)
- `seqnr` column in `RT` records the position of the `dgid` in a partial trajectory of an object; `seqnr = 1` denotes most recent / current
- Underline denotes DB indexes

SQL: Calculation of CDFS and CDMS

SQL 1 FUNCTION calc_CDFS()

```
SELECT dgid, count(*) AS nr, avg(spd) AS mu,
       COALESCE(stddev(spd),0) AS sig
FROM RT
GROUP BY dgid;
```

SQL 2 FUNCTION calc_CDMS()

```
SELECT dst.dst_dgid, src.src_dgid,
       count(*) AS nr_srs2dst
FROM (SELECT oid, seqnr, dgid AS dst_dgid
      FROM RT) AS dst,
     (SELECT oid, seqnr, dgid AS src_dgid
      FROM RT) AS src
WHERE dst.oid = src.oid
  AND dst.seqnr < src.seqnr
GROUP BY dst.dst_dgid, src.src_dgid;
```

Same object trajectory (WHERE dst.oid = src.oid)

Destination after source (AND dst.seqnr < src.seqnr)

SQL: Incremental Calculation of HDFS and HDMS

SQL 3 FUNCTION ud_HDFS()

```
UPDATE HDFS AS gh
SET nr = (c.nr+gh.nr),
    mu = (c.nr*c.mu+gh.nr*gh.mu)/(c.nr + gh.nr),
    sig = sqrt((gh.nr * gh.sig^2 + c.nr * c.sig^2) /
               (gh.nr + c.nr) +
               (gh.nr * c.nr * (gh.sig - c.sig)^2) /
               (gh.nr + c.nr)^2),
    nr_suc = (c.nr_suc+gh.nr_suc)
FROM CDFS AS c
WHERE gh.dgid = c.dgid;

INSERT INTO HDFS (dgid, nr, mu, sig, nr_suc)
SELECT c.gid, c.dir, c.nr, c.mu, c.sig
FROM CDFS AS c
LEFT JOIN HDFS AS gh
ON (gh.dgid = c.dgid)
WHERE gh.dgid IS NULL;
```

No previous HDFS

- Incrementally update previously observed HDFS based on non-overlapping subset / tumbling window statistics
- Insert new / not-yet-observed statistics
- Analogous calculations for HDMS

SQL: Calculation of Directionally Congested Cells

SQL 4 FUNCTION CongCell(min_veh, max_cv, max_z, max_relspddiff)

```
1   SELECT c.dgid AS dgid
2   FROM HDFS AS gh, CDFS AS c
3   WHERE gh.dgid = c.dgid
4     AND c.nr >= min_veh
5     AND c.sig / c.mu < max_cv
6     AND (c.mu - gh.mu) / (gh.sig / sqrt(c.nr)) < max_z
7     AND (c.mu - gh.mu) / gh.mu < max_relspddiff;
```

Directional congestion criteria (4-7)

SQL: Calculation of Directional Congestion Notifications

SQL 5 FUNCTION CongNotif(min_veh, max_cv, max_z, max_relspddiff, min_notif_prob)

```
1  WITH cond_prob AS
2       (SELECT m.src_dgid, m.dst_dgid,
3               m.nr_src2dst::float / f.nr_suc AS cond_p
4        FROM HDMS m, HDFS f
5        WHERE m.src_dgid = f.dgid)
6  SELECT t.oid, c.dgid AS con_dgid
7  FROM cond_prob AS gcp, RT AS t,
8       CongCell(min_veh, max_cv,
9                max_z,max_relspddiff) AS c
10 WHERE t.seqnr = 1
11   AND gcp.src_dgid = t.dgid
12   AND gcp.dst_dgid = c.dgid
13   AND gcp.cond_p >= min_prob;
```

CTE for $P((g_d, dir_d)|(g_s, dir_s))$

Object is currently located (10) in a source dgid (11) from which the conditional probability (13) of a directionally congested (12) dgid is larger or equal than the threshold

Temporal Domain Projections

- To capture temporal regularities in flows and movements the proposed method extracts HDFS and HDMS for different values of day-of-week and hour-of-day temporal domain projections
- Clients calculate `dow` and `how` projections of their status reports
- The `HDFS` and `HDMS` tables store the domain projected aggregates using the value `-1` to denote the "any" value
- Detection and notification queries combine a disjunction of conditions using the relevant domain projected information in their decision criteria
  - Detection if the statistical power criterion and the speed difference criterion are satisfied either based on the `dow`-projected, the `hod`-projected or the global statistics
  - Notification if the likelihood of encountering a congestion is above the threshold using either the `dow`-projected, the `hod`-projected or the global statistics

Generality of the Model and the Method

- Directional grid IDs can be replaced with adjacent road segments
- Gapless / spatially contiguous trajectories are not required but provide more robust statistics
- Congestion model can be replaced as needed
- Selective dissemination system has other applications, e.g., LBA

Outline

- Introduction
- Related work
- Method
  - Grid-based directional flow and mobility statistics
  - Directional congestion detection
  - Directional congestion notification
  - SQL-based implementation
- Empirical evaluations
  - Accuracy assessment
  - Scalability assessment
- Conclusion and future work

Empirical Evaluations: Environment + Data

- Environment: 64bit Ubuntu 14.04 LTS with PostgreSQL 9.3.9 on a PC with Intel Core i7-5600U @ 2.60GHz × 4 CPU, 16GB main memory and 512GB SSD
- Data set: 6 day sample of 11K taxis in Wuhan, China (85M records)
  - Outlier removal
  - 18km x 18km city center
  - Sampling gaps of more the 120 seconds delimit trips
  - Linear interpolation of trips between samples
  - Eliminate short trips (less than 300 seconds / 10 100m-grids)
  - → 2 million trips that have an average length of 1268 seconds and 82 grid cells; ~185M status reports

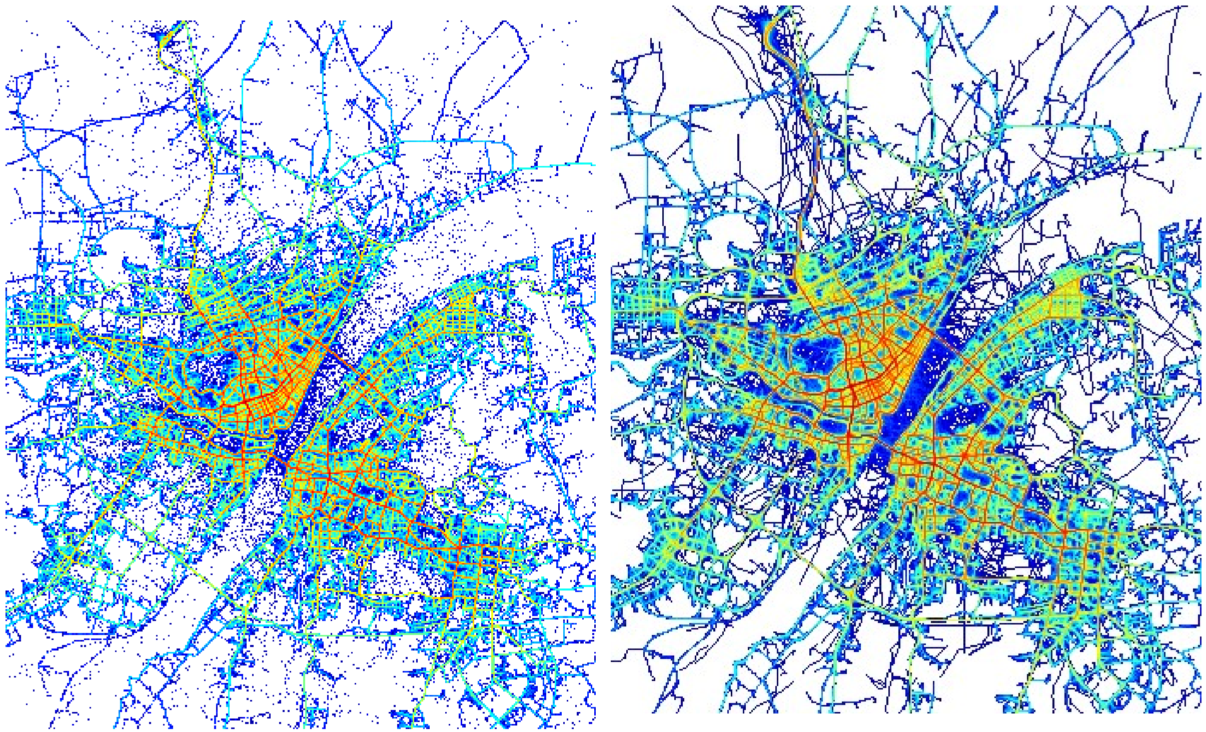

Raw sample vs. interpolated trips

Empirical Evaluations: Setup

- Accuracy + scalability assessments
- Temporal data alignments: hod (for accuracy) vs fixed (for scalability)
- Robustness of results: n-fold cross-validation
- Four notification systems:
  1. a system using hod-projected HDMS (*DMSChod*)
  2. a system using global HDMS (*DMSCglobal*)
  3. a system using non-directional, hod-projected HDMS (*NDMSChod*)
  4. a system using the LMC for notifications (*LMC*)
- Default parameters:
  - prediction horizon: Δt_{pred} = 60 seconds
  - minimum number of current status reports: *min_veh* = 2
  - maximum sample dispersion: *max_cv* = 0.5
  - maximum negative z-score: *max_z* = -1.65 (significance level of α = 0.05)
  - maximum negative relative speed difference: *max_relspddiff* = -0.5
  - minimum notification probability threshold: *min_prob* = 0.06

Empirical Evaluations: Framework + Measures

- Detected congestions are treated as ground truth: their spatio-temporal distribution and clustering are reasonable [Gid15]
- Modified binary assessment framework:
  - Baseline *B* for notifications: objects that can reach a congestion within the prediction horizon ==> *TN* = *B* – *TP* – *FP* – *FN*
  - Detections and notifications at different prediction times are unique
- Accuracy measures:
  - *TPR* = *TP*/(*TP* + *FN*) [sensitivity]
  - *FPR* = *FP*/(*FP* + *TN*) [1-specificity]
  - Cohen's kappa coefficient: discount for classification agreement due to chance
  - *AUC* (Area Under the [ROC] Curve): probability the a classifier assigns a higher positive-class probability to a randomly chosen positive case than a randomly chosen negative case
- Scalability measures: time and storage (# of DB rows) that the computation phases use

Accuracy Assessment

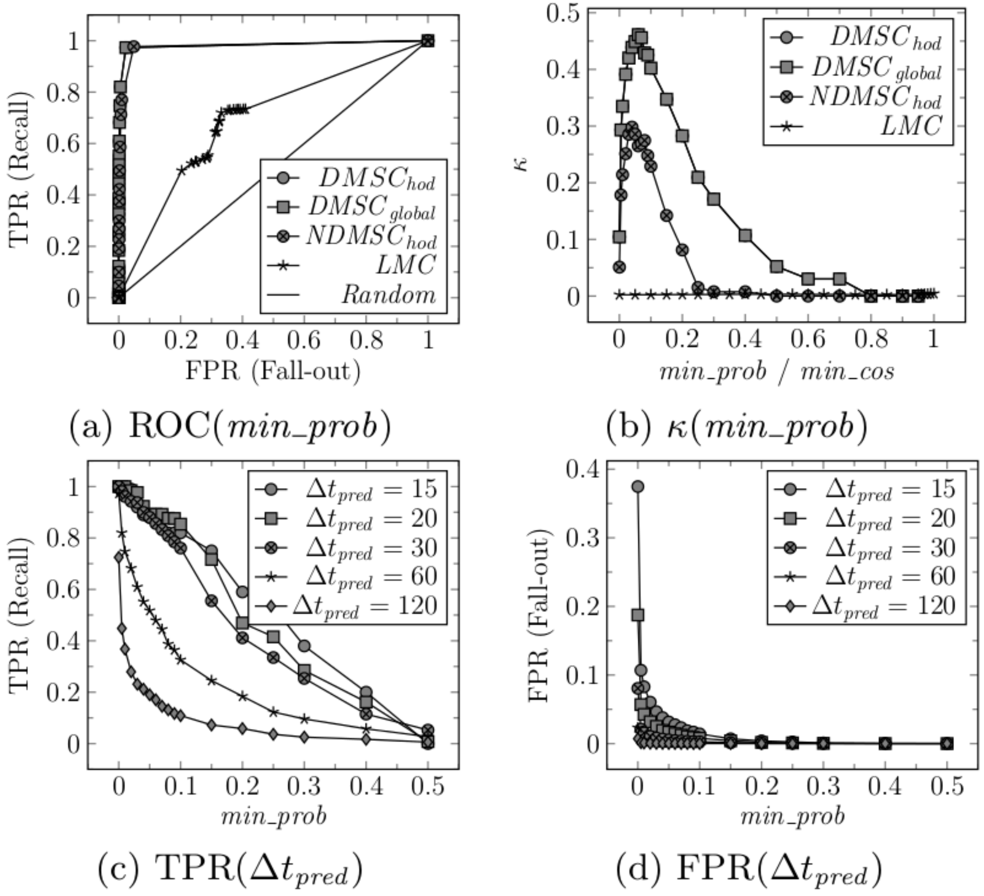

Figure 1: ROC (Fig. 1(a)) and Cohen's kappa coefficient (Fig. 1(b)) for varying min_prob / min_cos values for four notification systems. TPR (Fig. 1(c)) and FPR (Fig. 1(d)) for varying min_prob values for five $DMSC_{global}$-systems with different prediction horizon and temporal analysis window.

- Sensitivity to notification criteria thresholds
  - AUC of HDMS-based models (0.9799-0.9831) >> AUC of *LMC* (0.6907)
  - *LMC*: 41K *FP* to 55K objects about approx. 20 congestions
  - *AUC*(*DMSC*$_{hod}$) = *AUC*(*DMSC*$_{global}$)
- Sensitivity to prediction horizon length
  - While "shorter" mobility patterns cover most positive cases, "longer" patterns are spatially more specific
  - Similar results for sensitivity to spatio-temporal resolution: length of trajectories linearly increases with the resolution (1/*glen*)

Variability of HDMS

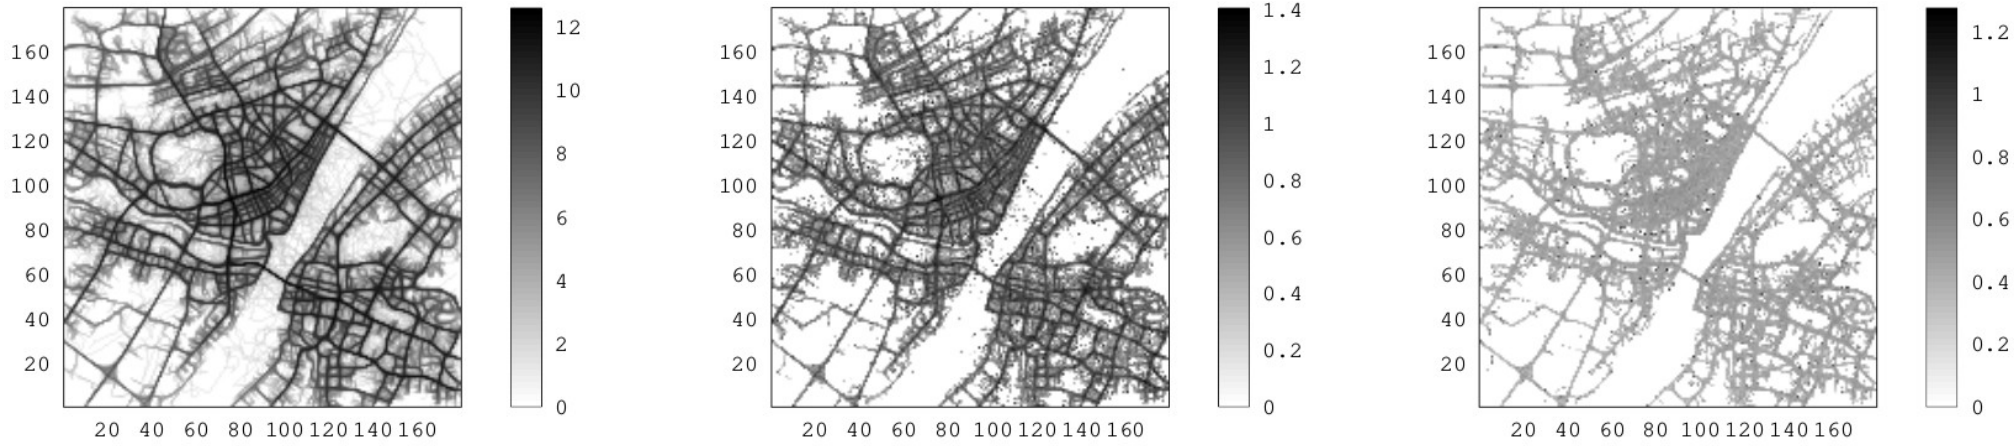

(a) log(#reports) (b) $CV_{dir\_hod}(g)$ (c) $CV_{dircond\_hod}(g)$

Figure 2: Spatial distribution of status reports and directional and temporal variability of HDMS.

- Coefficient of variation $CV = \sigma / \mu$ of
  - Directional hod-projected mobility statistics: $CV_{dir\_hod}(g)$
  - *Directionally-conditioned hod-projected mobility statistics:* $CV_{dircond\_hod}(g)$
- $CV_{dir\_hod}(g) > CV_{dircond\_hod}(g)$: directional aspects of the patterns capture most of the variability

Scalability Assessment

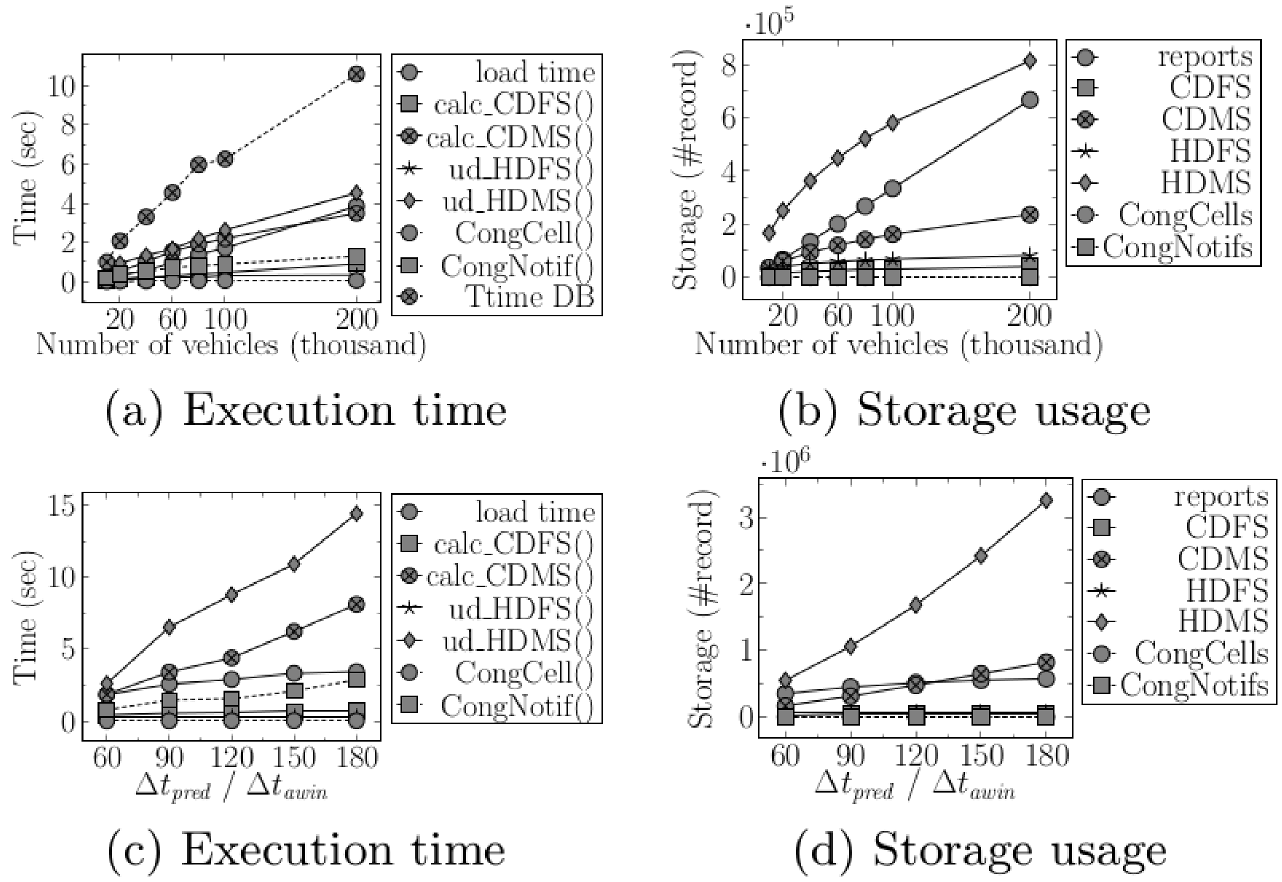

Figure 3: Execution time and space usage of different phases of the congestion detection and notification tasks in the $DMSC_{global}$ system for varying number of vehicles and values of prediction horizon / temporal analysis window size.

- All stages of processing an average temporal analysis window scale in the worst case linearly with the load
- Discounting the dominating load time, given a 60-second real-time processing limit, the system can mange approximately 60/10.5* 0.2K = 1.14M objects
- Linear behavior with prediction horizon length and resolution

Outline

- Introduction
- Related work
- Method
  - Grid-based directional flow and mobility statistics
  - Directional congestion detection
  - Directional congestion notification
  - SQL-based implementation
- Empirical evaluations
  - Accuracy assessment
  - Scalability assessment
- Conclusion and future work

Conclusions and Future Work

- Conclusions
  - Data-driven approach and a directional grid-based, time-inhomogeneous, Markov jump process model for the detection of and selective dissemination of traffic congestion information
  - Superior prediction accuracy
  - Model captures the topology of the road network and the movement on it
  - Highly scalable, simple, portable, SQL-based implementation
- Future work
  - Performance evaluations of road network based adaptation
  - Implementation and evaluation using main-memory and stream based Big Data processing frameworks

Thank you for your attention!

Q/A?