Prediction of duration and impact of non-recurrent events in transportation networks

Sébastien Blandin, Vikneswaran Gopal, Karthik Thirumalai
IBM Research Collaboratory – Singapore,
Email: sblandin@sg.ibm.com

Joanne Cheong
Intellligent Transport Systems Development Division,
Land Transport Authority, Singapore

1 Motivation

The need for efficient monitoring and management of transportation networks, driven by the growth of large cities worldwide, motivates the development of improved forecasting capabilities, required for the shift from a paradigm of *reactive* control methods, to a paradigm of *proactive* control actions. In the context of recurrent congestion, resulting from a structural lack of capacity at peak hours, statistical methods able to leverage large volumes of historical and online data have been shown to provide accurate predictions up to one hour in the future, see Min and Wynter (2011). However, for non-recurrent events such as traffic accidents, occurring outside of peak hours, and known to account for about half of the total delay in the US, see Schrank and Lomax (2011), innovative techniques are required for handling both the complex causal structure of incident duration and impact, as well as the relatively low volume of available data.

In this article, we propose to combine a statistical hierarchical approach suitable for low volume of heterogeneous data, with a causal flow model of impact propagation on a network. The contributions of this work include:

- **Robust incident classification method**: incident are classified based on their observed duration, using a hierarchical method robust to unobserved incident features and allowing for online updates.
- **Adaptive online predictions**: a Bayesian change-point method combined with a statistical smoothing technique updates online the local impact prediction.
- **Causal model-based impact propagation**: a conservation law propagation model allowing for model-based control predicts the incident impact on a sub-network.

- **Large-scale implementation in Singapore**: the algorithms proposed are implemented on the full Singapore road network including more than 60 thousands links, and tested using 1 year of traffic measurements.

The rest of the document is organized as follows. Section 2 presents the classification model and breakpoint method composing the local incident impact model, and section 3 presents the macroscopic model proposed for the network impact prediction.

2 Incident duration and local impact prediction

The incident duration and local impact are predicted using a combination of a statistical tree technique, a Bayesian change-point method, and a linear regression model. In this section we provide an overview of these algorithms and their interaction.

2.1 Model definition

The first component of the local impact prediction is a regression tree model predicting the duration of the incident. It is an extension to the technique introduced in He et al. (2011), which allows for updated predictions as new data becomes available. Specifically, with X the random variable representing the duration of an incident, the conditional distribution $X|X > t$ is used to provide an updated prediction of the duration of the incident as it evolves.

The second component of the algorithm is a Bayesian change-point detection model. The observation model for the change-point time κ reads:

$$Y(t_i) \sim \begin{cases} \mathcal{N}(\beta_0\, t_i + \alpha_0, \sigma^2) & \text{if} \quad t_i \leq \kappa \\ \mathcal{N}(\beta_1\, t_i + \alpha_1, \sigma^2) & \text{if} \quad t_i > \kappa, \end{cases} \tag{1}$$

where the observations $Y(t_i)$ corresponding to different observation times t_i are independent. The priors $\beta_0, \beta_1, \alpha_0, \alpha_1, \sigma^2$ and associated hyper-priors are not explicited here in the interest of space. The first and second equations in (1) respectively describe the evolution of traffic conditions at the location of the incident, before and after the change-point time denoted by κ. If the posterior distribution of κ is such that $\Pr(\kappa = t|Y(t_1), \dots, Y(t_N)) > 0.5$ for some $t \in [t_1, t_N)$, we conclude that a change-point has occurred. For example, by setting $\beta_0 = 0$ and constraining β_1 to be negative, and if $Y(\cdot)$ denotes the speed, then κ allows us to assess if speed has evolved into a phase where the speed drops. In our implementation, a Gibbs sampler is used to compute the desired posterior distributions (see Carlin et al., 1992).

The third component of the local impact prediction is a regression model that is applied whenever the change-point module determines that $\beta_1 < 0$. In such a case, the linear regression model is used to determine the magnitude of the speed drop from the previous phase, using β_1 as the predictor. This regression model is fit offline using historical incident data.

2.2 Sequential online predictions

In this section we describe the online interactions of the models described in the previous section. In order to know the current phase, the change-point algorithm is applied at each discretization time. If a new change-point is detected, a new phase is set to have begun; the previous phase does not enter into subsequent consideration due to the assumed independence (see equation (1)). The prediction of the local impact is performed differently depending on the nature of the phase $\beta_1 < 0$ (congestion increasing), $\beta_1 > 0$ (congestion decreasing) or $\beta_1 = 0$ (stationary phase). The various predictions are described below, where we let α_0 represent the speed before it was adversely affected by the incident.

$\beta_1 < 0$: first, the regression model estimates the speed that the profile is going to drop *to*. Second, the tree model predicts how much longer the incident (speed drop) will last, after which it is assumed that speed will recover to α_0.

$\beta_1 = 0$: if the current speed value is below α_0, the tree model is used to predict how much longer this phase will continue, before recovering to α_0. If the current speed is above α_0, it is assumed that the incident impact has already played out to completion, and that speed will remain constant for the foreseeable future.

$\beta_1 > 0$: this would indicate that the speed is beginning to recover. Hence the forecast consists of predicting that speed will increase at rate β_1 until it reaches α_0, whence it would level off.

The local incident impact prediction, corresponding to perturbations in traffic conditions at the location of the incident, are then propagated on a sub-network using a network conservation law model, described in the following section.

3 Incident impact prediction

The prediction of the impact of an incident on the road network is based on the propagation of the prediction of the local impact of the incident, computed using a statistical model as described in the previous section. The prediction of the impact of the incident on the network accounts for historical relations between speed and flow which characterize the performance of a link, and for calibrated splitting rates at junctions (see Work et al. (2010) for application of a similar model to traffic estimation).

The link model used for incident impact propagation consists of the offline calibration of a quadratic-hyperbolic-linear fundamental diagram, expressed in speed-flow coordinates:

$$q(v) = \begin{cases} c\, v\, (1 - \frac{v}{v_{\max}}) & \text{if} \quad v_{c+} \leq v \\ q_{\max} & \text{if} \quad v_{c-} \leq v \leq v_{c+} \\ \rho_{\max}\, \frac{v}{1+\frac{v}{w}} & \text{if} \quad v \leq v_{c-}, \end{cases} \tag{2}$$

where the parameters $v_{\max}, q_{\max}, \rho_{\max}$ respectively denote the maximal velocity, flow, density, where v_{c-}, v_{c+} are calibrated from historical data, and where c, w are defined by continuity of the flux function. Based on the incident features, the fundamental diagram at the incident location is scaled in order to account for the potential capacity drop and speed reduction at the location. The model computes speed and flow values in a spatio-temporal cone around the incident, using a first-order scalar conservation law model expressed in velocity variables:

$$\frac{\partial\, \frac{q(v(t,x))}{v(t,x)}}{\partial t} + \frac{\partial\, q(v(t,x))}{\partial x} = 0. \tag{3}$$

The flow across each junction is obtained by solving a linear program, which computes the maximal allowable flow across the junction, based on link traffic conditions, and calibrated splitting rates. At each time discretization, the linear program provides link boundary conditions for the link model (2)-(3), at the junctions located in the interior of the area impacted by the incident. Network boundary conditions for this propagation model correspond on the cone axis to the local incident impact prediction computed from the statistical model described in the previous section, and to the prediction of recurrent traffic conditions on the cone boundary.

References

B.P. Carlin, A.E. Gelfand, and A.F.M. Smith, Hierarchical Bayesian analysis of change-point problems, *Applied statistics*, 389–405, 1992.

Q. He, Y. Kamarianakis, K. Jintanakul, and L. Wynter. Incident duration prediction with hybrid tree-based quantile regression. IBM Research Report, 2011.

W. Min and L. Wynter, Real-time road traffic prediction with spatio-temporal correlations, *Transportation Research Part C: Emerging Technologies*, 19(4), 606–616, 2011.

D. Schrank and T. Lomax, 2011 Urban Mobility Report, *Texas Transportation Institute*, 2011.

D.B. Work, S. Blandin, O.-P. Tossavainen, B. Piccoli and A.M. Bayen, A traffic model for velocity data assimilation, *Applied Mathematics Research eXpress*, 1, 1–35, 2010.