



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Probabilistic Modelling of the Impact on Bus Punctuality of a Speed Limit Proposal in Edinburgh

Citation for published version:

Reijsbergen, D & Ratan, R 2016, Probabilistic Modelling of the Impact on Bus Punctuality of a Speed Limit Proposal in Edinburgh. in *9th EAI International Conference on Performance Evaluation Methodologies and Tools*. ACM, 9th EAI International Conference on Performance Evaluation Methodologies and Tools, Berlin, Berlin, Germany, 14/12/15. <https://doi.org/10.4108/eai.14-12-2015.2262703>

Digital Object Identifier (DOI):

[10.4108/eai.14-12-2015.2262703](https://doi.org/10.4108/eai.14-12-2015.2262703)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

9th EAI International Conference on Performance Evaluation Methodologies and Tools

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



Probabilistic Modelling of the Impact on Bus Punctuality of a Speed Limit Proposal in Edinburgh (Extended Version)

Daniël Reijsbergen

Rajeev Ratan

November 18, 2015

Abstract

We propose a data-driven methodology for evaluating the impact of the introduction of a speed limit on the punctuality of bus services. In particular, we use high-frequency Automatic Vehicle Location data to parameterise a model that represents the movement of a bus along predefined patches of the route. We fit the probability distributions of the time spent in each patch to two classes of probability distributions: hyper-Erlang distributions, for which we use the tool HyperStar, and a variation of the three-parameter gamma distributions recommended by the Traffic Engineering Handbook. In both cases we obtain models that can be expressed using the framework of Probabilistic Timed Automata, allowing us to evaluate bus punctuality using the model checking tool UPPAAL. We conduct a case study involving a proposed speed limit in Edinburgh. This is an extended version of a paper presented at ValueTools 2015 [15].

1 Introduction

Anticipating the possible impact of changes in the operating environment on mass transit timetables has been a long-standing challenge for planners and operators. A present manifestation of this situation is seen in the City of Edinburgh, where the city council has agreed to implement a new speed limit along the majority of the city's roads after years of campaigning by road safety groups. The concerned roads are heavily utilised by Lothian Buses, the main urban bus company that operates in the City of Edinburgh, who pride themselves on maintaining a high level of service quality and reputation. Consequently, the delays emanating from the introduction of these speed limits have the possibility to adversely affect timetable adherence, resulting in missed targets set by regulating bodies, and an inconvenienced travelling public. The proposed speed limits (see Figure 1) affect the majority of their routes currently oper-

ated. As such, quantifying the possible impacts of these speed limits is of major concern.

In simulating the effects of changes to real world systems, there are two basic approaches. The first involves direct experimentation on the existing system under the new conditions (or a close approximation thereof). This, if done correctly, produces a realistic measure of the effect. However, many real world systems do not lend themselves to such direct experimentation as it is either costly or logistically infeasible. As such, the second approach seeks to create a model to mimic the relevant behaviour of the system. This model can then be used to simulate a variety of conditions or constraints. However, this poses a challenge as real world bus systems are often subject to stochastic behaviour, thus necessitating a suitable model to reflect this.

In this paper we perform a case study on two bus routes in Edinburgh and provide a quantitative answer to the question of how the proposed speed limit impacts bus punctuality. Our methodology for developing bus models involves an expanded and modified implementation of the work done in [16] and [12]. Using this approach, we build a model that represents the continuous-time movement of buses as they move along predefined segments of a route. The time needed to move from one of the segments (referred to as 'patches' throughout this paper) to the next is modelled as a random variable whose distribution is chosen to best fit high-frequency (approximately every 5 seconds) GPS data provided to us by Lothian Buses. In this paper, we compare two classes of distributions used to model the time in each of the patches. The first of these is the class of hyper-Erlang distributions, a modelling class with good performance in a broad range of settings and which is easily parameterised using the user-friendly tool HyperStar [18]. The second is a variation of the three-parameter gamma distribution, which is an industry standard recommended by the Traffic Engineering Handbook [13]; its parameters are determined using a purpose-built optimisation tool written in the language of the

statistical package R [14]. The resulting distributions are then used to build probabilistic models in the tool UPPAAL [4], which facilitates the modelling formalism of Probabilistic Timed Automata and the simulation-based verification framework of Statistical Model Checking (SMC) [6].

We apply this modelling approach to a case study involving two bus services. The first is the Airlink, a popular bus route that connects Edinburgh’s city centre to the airport. Big delays may cause disaffection among passengers travelling to the airport. However, we find that the proposed speed limit has a small impact, namely between 7.61 and 24.3 seconds depending on the time and the direction. The second is the N11 night service, which is particularly affected by the new speed limit seeing as 50% of its route is affected by the speed limit, and because buses are particularly able to speed up to over 20mph at night. In this case, average delays are still only about 32 seconds. Finally, we fully exploit the stochastic model and the SMC engine of UPPAAL to compute ‘on-time’ probabilities for the final stops of the Airlink, namely the probability of the bus arriving at the stop between at most one minute early and five minutes late.

The structure of this paper is as follows. In Section 2 we discuss the context of our work, describe the dataset, and provide an overview of related research. In Section 3 we give a formal description of the modelling and parameter fitting techniques used in this paper, and discuss the use of UPPAAL. In Section 4, we use the presented methodology to quantify the impact of the speed limit plan on the Airlink service. Section 5 concludes the paper.

2 Model & Preliminaries

2.1 Motivation & Background

The speed limit plan discussed in this paper was approved by Edinburgh’s City Council on 13 January 2015, a decision which followed three years of research and three months of public consultation on the part of the council [7]. The speed limit will be introduced gradually starting in April 2016 [2]. The move has been hailed by campaign groups such as Living Streets Edinburgh [11], 20’s Plenty For Us and Friends of the Earth Scotland. The arguments in favour of the speed limit include public support for the speed limit witnessed across several opinion polls, improved road safety, cycling and walking being encouraged leading to lower carbon emissions and health benefits, and a nicer and more attractive city overall.

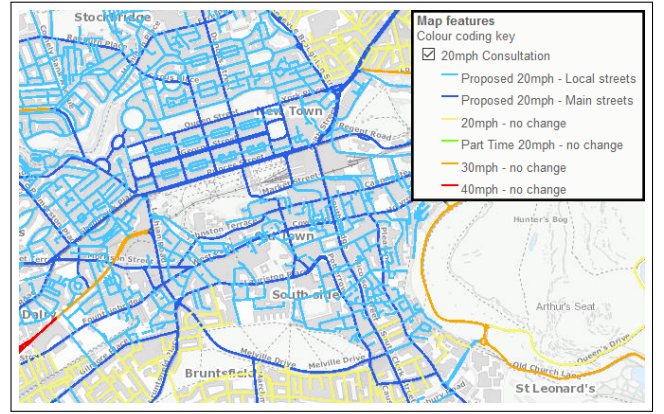


Figure 1: Map of proposed 20mph streets, zoomed in slightly to focus on Edinburgh’s city centre (blue and dark-blue roads will be changed into having 20mph speed limits). An interactive version is available at <http://edinburghcouncilmaps.info/LocalViewExt/Sites/20mphConsultation/>

The move has not been entirely uncontroversial [8], with a Facebook campaign called ‘Say no to 20mph’ being started and a related petition drawing 2700 signatories [20]. Opponents of the speed limit plan are generally skeptical of the suggested benefits and argue that the speed limit will encourage antisocial behaviour (such as aggressive tailgating), that it will be ineffective if it is not properly enforced, and that journey times will be adversely affected. Longer journey times also have an impact on bus timetable adherence, which is of specific interest to Lothian Buses due to its effect on passenger satisfaction and the fact that they have been fined £10,500 in the past for non-compliance with punctuality targets [21]. The cause of the fine was that the impact of road works was overestimated, leading to buses running early with high probability. Accordingly, there is a demand for predictions of the impact of the speed limit to be as accurate as possible.

As we demonstrate in this paper, the accuracy of these predictions — particularly for timetable adherence probabilities — can be affected by the modelling choice. In previous work [16], the hyper-Erlang distribution was chosen to model the time spent in patches for several reasons, including its general applicability and the fact that the resulting models are Continuous-Time Markov Chains, for which many efficient analysis techniques exist. However, as we discuss later on, an important alternative, namely the probability distribution that is recommended for traffic engineers, may not yield a Markov chain, but the resulting model can still be expressed using the framework of Probabilistic Timed Automata, allowing us to use the tool UPPAAL [4] and its powerful

stochastic simulation engine. We compare the two in this paper, not just with the intent to make the predictions of the speed limit impact as accurate as possible, but also to inform the construction of more general urban traffic models as part of the EU project QUANTICOL.

2.2 Patch Structure

Although the Automatic Vehicle Location (AVL) measurements supplied to us by Lothian Buses are roughly done in continuous space, we discretise our model by dividing the full route into ordered patches and consider the travel time through each patch. Since our case study does not require the same road-level detail as in [16], the patches were chosen to be relatively large, namely between 500m and 2.5km.

Our patches were designed to encapsulate areas between major bus stops as shown in Figure 2. Another important criterion was to ensure that the road segments covered by a patch had to either remain unaffected by the speed limit or be changed to 20mph, but not both. This was done so that we could implement the speed limit effects at a per-patch level. Finally, specific patches were constructed for the end points of the route, because buses tend to spend a large amount of time there; these are the black patches in Figure 2. We did not include the time spent in these patches in the total journey times.

2.3 Description of Data

The dataset used in this paper was gathered using the WiFi capabilities of some of the vehicles used by Lothian Buses. These WiFi services are provided via Long Term Evolution (LTE) 4G cellular modems which allow for reliable high-frequency — once every 5 seconds — AVL measurements. It should be noted that non-WiFi equipped buses also relay AVL data as part of the MyBusTracker service that provides information for arrival time prediction at stops and web applications.¹ However this is at a much lower frequency, namely between 30 to 40 seconds. For analysis of bus speeds, speed averages of over 40 seconds are typically too imprecise for accurate model parameterisation.

The data set used for the case study contains information for all buses operating the Airlink route from Monday 20 July 2015 to Wednesday 22 July 2015. We also have access to data involving Route 31 and N11 night route vehicles from the same period, although we will only briefly consider the latter

in Section 4. The data was provided in the form of files (in CSV format) containing all AVL measurements logged per bus per day, starting at 3AM in the morning and ending at 3AM the next day. Speed measurements were included in the dataset, although these could otherwise be calculated from the distance travelled.

Since traffic conditions vary throughout the day, we focus on two non-rush hour periods, namely the period between 9AM and 3PM and the period between midnight and 5AM. The reasons are twofold: even between individual hours in the rush hour period there are still considerable differences, and the impact of the speed limit is expected to be the largest outside the rush hours. We also categorise measurements based on the direction that is being travelled, as this may have an impact; in fact, the distance crossed in Patch 10 of Figure 2 is not the same in both directions.

Routes are typically operated by several individual buses per day. In order to comply with the timetable, buses first leave the depot and then head towards the starting points of the route. Buses then leave these starting points in staggered intervals as per their timetable. This allows buses to service stops along the route in frequencies ranging from once per 5 to 15 minutes (60 minutes for the night bus). Within a single day (or night), buses may change route. This can be problematic, as patch crossing time measurements are inaccurate if only portions of a patch were completed. For the Airlink service this is not the case, because the vehicles of this service are painted in a distinct livery and hence never service other routes. For Route 31 it did occur frequently and for that reason, in combination with the fact that there were roadworks on the A772 during the observation period, this route will not be further considered in this paper.

In all datasets, some measurements appeared to be erroneous or the result of extreme conditions. For example, gaps have been observed in the dataset, possibly caused by temporary connection losses. This could lead to extremely small measurements if the connection was restored right before a patch was exited. Hence, a data cleaning operation was performed in which measurements were discarded if they were over the median plus three times the standard deviation in that patch, or if they were below 30 seconds (the time needed to clear the smallest patch — Patch 8 of Figure 2 — at 36mph). The impact was limited: for example, 18 out of 742 measurements across all patches were discarded in the east-west direction between 9AM and 3PM.

¹An API for the low-frequency data is publicly available via <http://www.mybustracker.co.uk/>.

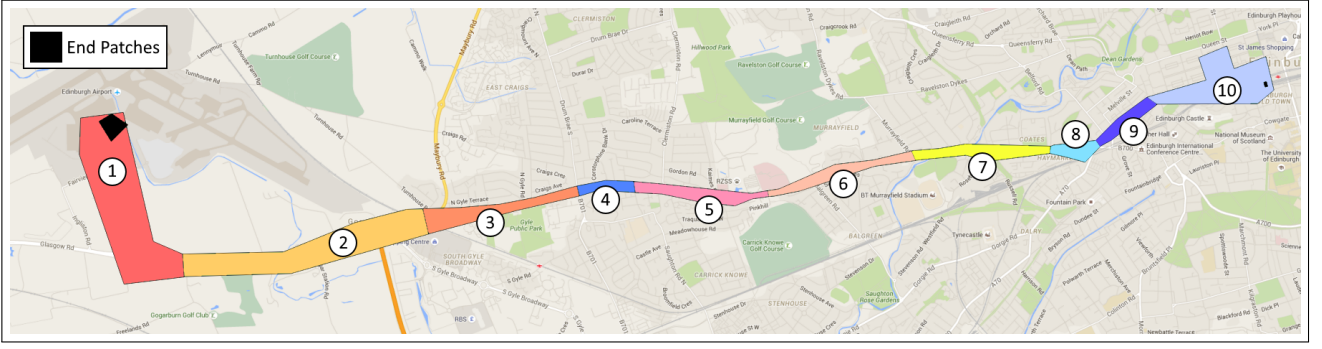


Figure 2: Illustration of the patches used to partition the Airlink route. Patches 4, 8, 9, and 10 will be affected by the new speed limit.

2.4 Related Work

This paper broadly follows the same approach as in [16], in which a patch-based analysis was carried out to evaluate the impact of the introduction of trams to Edinburgh’s city centre, and [12], in which Erlang distributions were used to model the time spent by buses between major stops, and to suggest improvements to bus timetables. The contributions of this paper are the use of a probability distribution recommended in the traffic engineering literature [13], the use of the statistical model checking engine of UP-PAAL [10] and the application to a new case study that has received considerable media attention in Edinburgh.

The literature on the impact of speed limits does not appear to be very extensive. One approach would be to use of a traffic simulation tool such as the microscopic tool SUMO [3] or the mesoscopic formalism of [5]. Another approach would be the use of statistical analysis of the impact of speed limits in other cities; see [1] for a literature overview.

3 Probabilistic Modelling

In this section, we present the probabilistic models used in the analysis section of this paper. The most important probabilistic concept that we use is that of the exponentially distributed *phase*. In particular, the time spent in each patch will be modelled as a function of several such phases. We begin with a formal discussion of the above, then discuss parameter fitting and the expression of these models using the language of the model checking tool UPPAAL.

3.1 The Basic Distributions

Recall that the time T spent in a phase is *exponentially distributed* with rate $\lambda > 0$ if its *probability*

density function (pdf) is given by

$$f_T(t; \lambda) = \lambda e^{-\lambda t}$$

for $t \geq 0$ and by 0 otherwise. The expected amount of time spent in a phase is given by $\frac{1}{\lambda}$; hence, large values of λ mean that the time spent in a phase is small on average. The pdf $f_T(t)$ of a random variable T determines the probability of observing values from a small interval around t , and will be used later in this paper to measure how well a fitted probability distribution corresponds to the data.

If T_1, \dots, T_k are exponentially distributed random variables with rate λ , then $S = T_1 + \dots + T_k$ is *Erlang-distributed* with rate λ and shape k . Its pdf equals

$$f_S(s; k, \lambda) = \frac{\lambda^k s^{k-1} e^{-\lambda s}}{(k-1)!} \quad (1)$$

for $s \geq 0$ and 0 otherwise.

As mentioned in Section 2.1, we will compare the modelling accuracy of two probability distributions used to model the time spent in a patch; both are generalisations of the Erlang distribution. The first is the *hyper-Erlang distribution*, which is constructed in the following way. Given some $m \in \mathbb{N}$, let $\alpha_1, \dots, \alpha_m$ be a probability distribution over $\{1, \dots, m\}$, i.e., $\forall i, \alpha_i \in [0, 1]$ and $\sum_{i=1}^m \alpha_i = 1$. Furthermore, let $S_i, i \in \{1, \dots, m\}$, be Erlang-distributed with rate λ_i and shape k_i — these random variables are called the *branches* of the hyper-Erlang distribution. We then say that a random variable Y is hyper-Erlang-distributed with shapes $\vec{k} = k_1, \dots, k_m$, rates $\vec{\lambda} = \lambda_1, \dots, \lambda_m$ and branch probabilities $\vec{\alpha} = \alpha_1, \dots, \alpha_m$ if

$$f_Y(y; \vec{k}, \vec{\lambda}, \vec{\alpha}) = \sum_{i=1}^m \alpha_i \frac{\lambda_i^{k_i} y^{k_i-1} e^{-\lambda_i y}}{(k_i-1)!} \quad (2)$$

for $y \geq 0$, where f_Y denotes the pdf of Y .

The second generalisation of the Erlang distribution is one where a real constant $c > 0$ is added. This

gives rise to a random variable Z with pdf

$$f_Z(z; k, \lambda, c) = \frac{\lambda^k (z - c)^{k-1} e^{-\lambda(z-c)}}{(k-1)!} \quad (3)$$

for $z \geq c$. This is a variation of the probability distribution recommended in the Traffic Engineering Handbook [13]; in their version, the factorial in (3) is replaced by a gamma function, thereby allowing k to be any positive real value instead of a positive integer. We require k to be an integer in order to be able to express the distribution in UPPAAL. In the following, we will call this distribution ‘Erlang+ c ’.

3.2 Parameter Fitting

Assuming that we have chosen either the hyper-Erlang or the Erlang+ c distribution for the time spent in a patch, the next question is how to choose the distribution parameters such that they correspond to the data in the best possible way. There are several methods for this; examples include the method of moments and maximum likelihood estimation. We use the latter, for which the idea is to maximise the product of the densities in each individual sample — or rather the logarithm thereof. To make this more concrete, assume we have a sample $\vec{x} = (x_1, \dots, x_n)$ of observations of the time spent in a patch. The total *log-likelihood* of observing \vec{x} as realisations of Y with parameters \vec{k} , $\vec{\lambda}$, and $\vec{\alpha}$ is then given by

$$l_Y(\vec{x}; \vec{k}, \vec{\lambda}, \vec{\alpha}) = \sum_{i=1}^n \log(f_Y(x_i; \vec{k}, \vec{\lambda}, \vec{\alpha}))$$

with f_Y as in (2). The log-likelihood l_Z of observing \vec{x} as realisations of Z is defined analogously for parameters k, λ , and c , and f_Z as defined in (3). The functions l_Y and l_Z are then maximised over the available parameters to obtain an optimal fit to the data. For the hyper-Erlang distribution, this is done using the tool HyperStar [18], which uses the expectation maximisation algorithm described in [22] and [17]. For the Erlang+ c distribution, we use a purpose-built numerical optimisation routine based on gradient/steepest ascent (see, for example, [9]) with golden section search. We discuss this in more detail below.

We first assume that k is given and that we need to find the values of λ and c on \mathbb{R}^2 that maximise l_Z . The (partial) derivatives $\frac{\partial}{\partial \lambda} l_Z$ and $\frac{\partial}{\partial c} l_Z$ are known in closed form, and we use this knowledge to construct an iterative procedure in which we move closer to the nearest optimum in each step. First, we initialise λ and c ; with \bar{x} denoting the average of \vec{x} and $\min(\vec{x})$

its minimum, we have found that $c = \frac{9}{10} \min(\vec{x})$ and $\lambda = \frac{k}{\bar{x}-c}$ are generally good choices for λ and c respectively (the idea behind the factor $\frac{9}{10}$ is that we want c to start close to the smallest observation but not exactly at it). Then, given the values of λ and c in the current iteration, we define

$$h(u) = l_Z(\vec{x}; k, \lambda + u \frac{\partial}{\partial \lambda} l_Z(\vec{x}; k, \lambda, c), c + u \frac{\partial}{\partial c} l_Z(\vec{x}; k, \lambda, c)). \quad (4)$$

This reduces each iteration to single-dimensional optimisation of h , as h is simply l_Z defined on a line in the direction of the steepest ascent. Note that $u = 0$ corresponds to λ and c remaining the same as in the previous iteration, so we search for a non-zero value of u that maximises (or at least improves) $h(u)$. When this value has been found, the next iteration’s values of λ and c can be computed as $\lambda + u \frac{\partial}{\partial \lambda} l_Z(\vec{x}; k, \lambda, c)$ and $c + u \frac{\partial}{\partial c} l_Z(\vec{x}; k, \lambda, c)$ respectively.

The function h equals $-\infty$ in parts of \mathbb{R} ; this happens when the corresponding value of λ is negative or when c becomes either negative or larger than the smallest observed value. This proves problematic for all of the general-purpose optimisation procedures implemented in R through the function `optim`, including Brent’s method which is specifically intended for single-dimensional optimisation. Accordingly, we have implemented our own algorithm which follows the principles of golden section search. We search for the local optimum of h on $[-1, 1]$, which means that in each iteration we move away from the old values of λ and c by at most the gradient (to see this, substitute -1 or 1 into (4)). Then, we fix an integer $l \in \{2, 3, \dots\}$ to be used throughout the entire procedure, and compute the values of h in the points $-1, -1 + \frac{1}{l}, -1 + \frac{2}{l}, \dots, 1$. We then determine the smallest value $i \in \{2, 3, \dots, 2l\}$ for which $h(\frac{i-1}{l} - 1) > h(\frac{i}{l} - 1)$. The reason is that we know that the value of u that maximises $h(u)$ must then be on $[\frac{i-2}{l} - 1, \frac{i}{l} - 1]$ — after all, the derivative of h can only change sign once if we assume that a global maximum exists. We found that $l = 2$ does not necessarily give the best run times; we used $l = 7$ throughout. We run the single-dimensional optimisation procedure several times (in our implementation, we took ten steps), then choose u to be the centre of the remaining interval and complete the iteration. We continue until the changes in terms of λ and c have dropped below a fixed threshold or if a fixed number of iterations has been reached.

Since we now have a procedure for finding the optimal values of λ and c as a function of k and \vec{x} , the final challenge is to find the optimal k . We use the

fact that k is discrete by starting with $k = 1$ and then increasing k until l_Z starts to decrease. In our experiments, this appeared to return the global optimum. As can be seen in the ‘Erlang+ c ’ column of Table 3, the optimal values of k were typically not very large, so it was not necessary from a computational point of view to find a more efficient way of initialising this procedure.

3.3 Probabilistic Timed Automata and UPPAAL

UPPAAL [10] is a toolbox for the analysis of real-time systems modelled as *timed automata*. Of particular interest to this paper is the UPPAAL Statistical Model Checker (SMC) [6], which allows for the verification of Probabilistic Timed Automata (PTAs), the stochastic extension of timed automata. We do not give a full description of PTAs but discuss the key language features used in the paper through a discussion of the PTA models for the patch-based model.




Figure 3: Illustration of a PTA representing a system with two patches with two-branch Hyper-Erlang distributed crossing times.

In Figure 3, we depict a sequence of two patches for which the journey time is modelled using a two-branch hyper-Erlang distribution. States (called ‘locations’ in PTA terminology) are represented using circles, and edges (representing transitions between states) using arrows. The system starts in the left-most location (the double circle), which is exited with a rate of 1 second (this transition would ideally be immediate, but the 1-second delay is negligible). We need this state because the system requires the specification of a unique initial state. When the initial state is left, a probabilistic choice is made between the two branches of the next patch. The edge weights, which determine the relative likelihood of choosing an edge, are required to be integers. Hence, the values of α_1 and α_2 — found to be 0.4938 and 0.5062 by the parameter-fitting algorithm — had to

be converted to the smallest pair of integers with the same ratio, which in this case were 2469 and 2531.

If branch i has been chosen, the system then needs to complete k_i locations that each have rate λ_i . When this has been completed, the system jumps to a so-called branch point (the small circle) from which a probabilistic choice is made between the branches of the next patch. When the final patch has been completed, the system jumps to the final state labelled ‘end’.




Figure 4: Illustration of a PTA representing a system with two patches with Erlang+ c distributed crossing times.

In Figure 4, we depict a similar model for the Erlang+ c distribution. In addition to the discrete states, we now also use a real-valued clock, called x . This clock is reset immediately after a patch is left. We then impose a guard on the first (or only) edge of the patch requiring that this edge cannot be completed before the clock x has reached a value greater than c ; in practice, this adds a constant delay of c to the global clock. Otherwise, the behaviour in each patch is similar to that of a single Erlang branch in Figure 3.

Whereas Figures 3 and 4 depict graphical representations of the PTA models used in this paper, the same models can also be expressed in XML format. These XML files can be generated by R, taking as input either a file with the hyper-Erlang parameters obtained via HyperStar, or a matrix of the Erlang+ c distribution parameters obtained by another function in R. Furthermore, the system property of interest is also embedded in the XML file. We are interested in two probabilities: the property of a bus reaching the final location over one minute early or less than five minutes late.² For example, if the timetabled journey duration to get from Waverley station in Edinburgh’s city centre to the airport is 28 minutes (i.e., 1680 seconds) then the property

$$Pr[i=1620] \ (i \text{ Process.end})$$

²The notion of a bus being ‘on-time’ if it is at most N_1 minutes early and at most N_2 minutes late is a common notion, and the particular choice of $N_1 = 1$ and $N_2 = 5$ is from the Scottish Government’s Bus Punctuality Improvement Partnerships guiding document [19]. This notion of punctuality is less relevant for the Airlink as it is a ‘frequent’ service during the day, meaning that different standards apply; however, it is infrequent at night.

Patch	Erlang+c	Erlang	hyper-Erlang	
			2-branch	3-branch
1	-338	-343	-340	-337
2	-346	-386	-296	-286
3	-336	-336	-336	-334
4	-357	-358	-357	-356
5	-396	-401	-396	-394
6	-333	-339	-336	-334
7	-358	-361	-356	-356
8	-409	-409	-409	-409
9	-388	-388	-387	-384
10	-447	-449	-442	-440

Table 1: Total log-likelihoods for the distributions fitted to the patch journey times in each of the patches for the Airlink, from Waverley Bridge to the airport during midday, rounded to integers.

specifies the property of being over a minute early (‘Process’ is the name of the system and ‘end’ the name of the final state). The SMC engine of UPPAAL is then invoked to estimate the probability of this property being satisfied. In particular, system executions are repeatedly executed and the results used to make statistically justified statements about the relevant probabilities.

4 Analysis & Results

In this section, we present the results of applying the methodology described above to the case study of the proposed speed limit in Edinburgh. The section is divided into two parts. In Section 4.1, we discuss the results of the parameterisation and model construction steps, and discuss the impact of the modelling choices. In Section 4.3 we discuss the actual impact of the speed limit using the model built in the previous step.

4.1 Modelling Results

In order to best illustrate the impact of the modelling choices, we focus on a specific dataset of GPS measurements, namely the east-west direction of the Airlink and the time window between 9AM and 3PM, without the speed limit. This dataset consists of 724 measurements with 68-76 measurements per patch. The differences between the numbers of observations was due to the outlier removal; outliers occurred more often in the patches near the two ends of the route and less so in the centre.

In Table 1, we display for each patch and for four choices of distributions (Erlang, 2-branch hyper-

Erlang, 3-branch hyper-Erlang, and Erlang+c) the log-likelihood of observing the data under the distribution with the parameters that were found to be optimal. As can be seen from the table, the numbers tend not to differ starkly.

In nearly all cases the 3-branch hyper-Erlang has a higher log-likelihood than the Erlang+c. However, this comes at a price: the number of phases tends to be much smaller for the Erlang+c than for the other distributions. In Table 3, we display the expected number of completed phases: this is just k for the Erlang and Erlang+c, and $\sum_i \alpha_i k_i$ for the hyper-Erlangs. The differences are considerable, sometimes even between the Erlang and Erlang+c (see, for example, Patch 1). The consequences are twofold. Firstly, since the time to draw a sample using the SMC engine of UPPAAL scales roughly linearly with the expected number of completed phases, it takes less time to obtain accurate estimates of the on-time probabilities using the Erlang+c. Secondly, high phase counts can be a symptom of overfitting: it could indicate that a sharp peak around a cluster of observations has been created through an Erlang branch with many phases and a rate such that the branch mean is near the centre of the cluster.

Patch	Erlang+c	Erlang	hyper-Erlang	
			2-branch	3-branch
1	4	80	149.5	591.4
2	1	3	186.7	310.8
3	13	32	59.0	569.9
4	3	10	36.4	39.9
5	3	13	47.5	89.0
6	2	38	125.5	161.6
7	3	23	79.6	91.4
8	5	7	13.2	21.8
9	9	9	22.0	272.9
10	6	12	13.5	27.8

Table 3: The expected number of exponential phases for each of the distributions.

The choice of distribution does have a big impact on the tail of the distribution of the time needed to complete the entire journey. This is displayed in Figure 5, in which it can be seen that the hyper-Erlang distributions have a noticeably heavier tail on the right. As mentioned earlier, the time needed to generate samples for the Erlang+c model is lower than for the hyper-Erlang model; whereas Figure 5(a) is based on 25,103 samples generated in 6 seconds, Figure 5(b) is based on 22,116 samples generated in 39 seconds. We found that the distribution of total journey times does seem to have a relatively heavy tail,

Route	Period	Direction	Speed Limit		New	Difference	Perc. Change
			Timetable	Old			
Airlink	9AM - 3PM	Waverley-Airport	1680	1886.34	1896.82	10.48	0.56%
		Airport-Waverley	1800	2038.05	2045.66	7.61	0.37%
	12AM - 5AM	Waverley-Airport	1260	1413.83	1429.01	15.18	1.07%
		Airport-Waverley	1260	1208.13	1232.43	24.30	2.01%
N11	12AM - 5AM	Ocean Terminal-Hyvots Bank	2460	2256.93	2289.27	32.34	1.43%

Table 2: We depict the average total journey times in seconds both with and without incorporation of the the speed limit impact as per (5). This is done for two routes, the Airlink in both directions and during two different time periods, and for the N11 from Ocean Terminal to Hyvots Bank (the N11 only goes in one direction). Timetable predictions were obtained using the journey planner on <http://lothianbuses.com/getting-around/journey-planner>.

which would support the use of the hyper-Erlang distribution.

Finally, we consider the impact of the modelling choice on the computation of the on-time probabilities, which we will discuss in more detail below. As can be seen in Table 4, the choice of distribution can have a big impact, in our case even more so than the speed limit itself.

4.2 Patch Parameter Fitting Results

To better illustrate the results of fitting the hyper-Erlang distribution to the individual patch sojourn time datasets, we provide screen shots taken of the HyperStar tool in Figures 6 and 7. These screen shots depict the datasets of the time Airlink buses spent in each of the ten patches of the Airlink route (see Figure 2), both in terms of an observation histogram and the empirical Cumulative Distribution Function (CDF). We only consider buses heading from Waverley Bridge to the airport between 9AM and 3PM. In the same figure as the histograms and empirical CDFs, we depict the Probability Density Functions (PDFs) and CDFs of the hyper-Erlang distributions fitted to the data. For each dataset, we consider three types of hyper-Erlang distributions: single-branch (i.e., Erlang distributions), 2-branch, and 3-branch. The coloured vertical lines in the graphs depict the manually selected peaks which determine the maximum number of branches and inform the choice of initial parameters for the implemented parameter optimisation procedure (the EM algorithm).

As can be seen in the figures, the dataset for which the hyper-Erlang does best when compared the the Erlang+ c distribution is Patch 2, which appears to contain quite a few outliers. We mention again here that we do apply some form of outlier removal, but the used criterion for removal — namely a large distance from the mean relative to the standard deviation — does not lead to good results when the num-

ber of outliers is so high that the standard deviation becomes too large. Something similar applies to Patch 10, which has outliers in the tail on the left. The improvement of Erlang+ c over standard Erlang is the biggest for Patches 5 and 6. In general, the hyper-Erlang distributions have a tendency to overfit peaks in the data, which is particularly noticeable for Patch 9 and 3 Erlang branches.

4.3 The Impact of the Speed Limit

To evaluate the impact of the speed limit, we create a second dataset in the following way: we test for each GPS measurement (corresponding to a 5-second period) whether the recorded speed of the vehicle s was over 20mph, and if so, we add the value

$$\frac{s - 20}{20} \cdot 5 \quad (5)$$

to the total time spent in the corresponding patch. This quantity represents the additional amount of time needed to complete the same distance under the speed limit. It allows us to parameterise the model using both datasets and compare expected journey times and on-time probabilities. Note that (5) of course represents a rough approximation: if, for example, the speed limit also changes the density of vehicles on the road, then this could have an additional impact on bus travel times. One way of investigating this effect would be to use a microscopic road simulation tool such as SUMO, which incorporates both the buses and the other road traffic as individual vehicles in the simulation. However, since we do not have access to reliable non-bus traffic data, we are currently unable to realistically parameterise such a model, so this is left as a direction for future research.

As can be seen in Table 2, the impact of the speed limit varies across time periods and routes. For the Airlink, which is only impacted by the speed limit on the segment between Waverley Station and Haymarket (Patches 8, 9 and 10 of Figure 2), the im-

Route	Distribution	Speed Limit	Too Early	Too Late
Airlink	Hyper-Erlang	old	[0.091799, 0.092799]	[0.193120, 0.194120]
		new	[0.088044, 0.089044]	[0.210440, 0.211440]
	Erlang+c	old	[0.073240, 0.074240]	[0.335775, 0.336775]
		new	[0.067363, 0.068363]	[0.345214, 0.346214]

Table 4: 90% confidence intervals for the probability that a bus completes the Waverley-Airport journey during the 9AM-3PM period over one minute too early (so before 1620 seconds) or over five minutes too late (so after 1980 minutes). The hyper-Erlang distributions had two branches. All distributions were fully parameterised again for the all patches, even those which would be unaffected by the speed limit.

impact on the timetable seems to be negligible, as buses only rarely manage to go over 20mph in the city centre. As expected, the impact at night is larger than during the day. The N11, which is impacted by the speed limit on over half its route, is slowed down by about 58.91 seconds, although the total impact is still limited relative to the total route duration (namely about 2.27%). We do note that we observed that some N11 buses switched routes during the night, and although this did not appear to have a substantial impact the results are subject to this caveat.

The impact of the speed limit also does not seem to have much of an effect on the probabilities of not being on-time for the Airlink, as can be seen in Table 4. However, for the N11, the full minute being added on to the average route completion did have a noteworthy impact. It should be noted that buses of all routes typically spend a large amount of time stationary, particularly to allow passenger to board or alight, which mitigates the effect of the speed limit.

5 Conclusions

In this paper we have presented a probabilistic model that was able to quantify the impacts on bus timetables of the newly proposed 20mph speed limit plan in Edinburgh. We considered two ways of incorporating the stochastic nature of bus movement within route segments (‘patches’), namely the use of hyper-Erlang distributions and an Erlang distribution plus a constant that is a variation of the distribution recommended by the Traffic Engineering Handbook. The tail behaviour of the former is better able to represent total journey durations, whereas the latter is more efficient to simulate and its parametrisation is more stable across repeated experiments. We combined the patches into a single representative model in UPPAAL which allowed us to compute interesting punctuality metrics. The impact of the speed limit was found to be limited.

Regarding future work, patch selection and dimensioning is a time-consuming non-trivial task that easily becomes cumbersome when the number of patches

and routes is increased. Automated patch selection would greatly enhance the efficiency of our methodology. We did try to find patch centres using the k -means clustering algorithm, however one of the disadvantages is that there is no easy way to align patch boundaries to the roads affected by the speed limit. As such, we would need to find a way to automatically incorporate the map data of the proposed speed limit.


Also, the speed limit adjustment as given in Equation (5) is of course an approximation, and assumes that traffic conditions remain the same. A micro-scale simulation model could be constructed to validate this assumption. However, since we do not have access to full traffic data it would be a challenge to realistically parameterise such a model.

6 Acknowledgments


This work has been supported by the EU project QUANTICOL, 600708. The authors thank Bill Johnston and Philip Lock of Lothian Buses for providing access to the data and for their helpful feedback on parts of this research project, and Jane Hillston for her helpful comments.

References

- [1] J Archer, N Fotheringham, M Symmons, and B Corben. The impact of lowered speed limits in urban and metropolitan areas. Monash University Accident Research Centre, Victoria, Australia, 2008.
- [2] BBC News. Edinburgh to roll out 20mph speed limit from April 2016. <http://www.bbc.co.uk/news/uk-scotland-edinburgh-east-fife-31841370>. Published: 11-05-2015. Accessed: 06-09-2015.
- [3] Michael Behrisch, Laura Bieker, Jakob Erdmann, and Daniel Krajzewicz. SUMO—Simulation of Urban MObility. In *The Third In-*



(a) Erlang+c



(b) 2-branch hyper-Erlang

Figure 5: Observation histograms for the time needed to complete the entire Waverley-airport journey based on a model in which each patch either has an Erlang+c or a 2-branch hyper-Erlang distribution, generated using UPPAAL. The hyper-Erlang distributions result in a stronger skew towards higher values.

- ternational Conference on Advances in System Simulation (SIMUL 2011), Barcelona, Spain, 2011.*
- [4] Gerd Behrmann, Alexandre David, and Kim G Larsen. A tutorial on UPPAAL. In *Formal methods for the design of real-time systems*, pages 200–236. Springer, 2004.
- [5] Wilco Burghout. Mesoscopic simulation models for short-term prediction. *KTH Royal Institute of Technology*, 2005.
- [6] Alexandre David, Kim G Larsen, Axel Legay, Marius Mikućionis, and Zheng Wang. Time for statistical model checking of real-time systems. In *Computer Aided Verification*, pages 349–355. Springer, 2011.
- [7] Edinburgh City Council. About 20mph for Edinburgh. http://www.edinburgh.gov.uk/info/20243/20mph_for_edinburgh/1024/about_20mph_for_edinburgh. Accessed: 06-09-2015.
- [8] Edinburgh News. Huge backlash to council’s 20mph speed limit plans. <http://www.edinburghnews.scotsman.com/news/transport/huge-backlash-to-council-s-20mph-speed-limit-plans-1-3658104>. Published: 12-01-2015. Accessed: 06-09-2015.
- [9] Gaston Gonnet and Ralf Scholl. *Scientific computation*. Cambridge University Press, 2009.
- [10] Kim G Larsen, Paul Pettersson, and Wang Yi. Uppaal in a nutshell. *International Journal on Software Tools for Technology Transfer (STTT)*, 1(1):134–152, 1997.
- [11] Living Streets Edinburgh. Response to Edinburgh Council’s 20mph consultation. <http://www.livingstreets.org.uk/sites/default/files/Living%20Streets%20Edinburgh%20Response%20to%20Edinburgh%20Council%E2%80%99s%2020mph%20consultation.pdf>. Published: 12-10-2014. Accessed: 06-09-2015.
- [12] Ludovica Luisa Vissat, Allan Clark, and Stephen Gilmore. Finding optimal timetables for Edinburgh bus routes. *Electronic Notes in Theoretical Computer Science*, 310:179–199, 2015.
- [13] James L. Pline. *Traffic engineering handbook*. Prentice-Hall, 1992.
- [14] R Development Core Team. R: A language and environment for statistical computing. 2005.
- [15] D. Reijbergen and R. Ratan. Probabilistic modelling of the impact on bus punctuality of a speed limit proposal in edinburgh. In *Proceedings of the 9th EAI International Conference on Performance Evaluation Methodologies and Tools (ValueTools)*, 2015.
- [16] Daniël Reijbergen, Stephen Gilmore, and Jane Hillston. Patch-based modelling of city-centre

bus movement with phase-type distributions. *Electronic Notes in Theoretical Computer Science*, 310:157–177, 2015.

- [17] Philipp Reinecke, Tilman Krauß, and Katinka Wolter. Cluster-based fitting of phase-type distributions to empirical data. *Computers & Mathematics with Applications*, 64(12):3840–3851, 2012.
- [18] Philipp Reinecke, Tilman Krauß, and Katinka Wolter. HyperStar: Phase-type fitting made easy. In *Proceedings of the Ninth International Conference on the Quantitative Evaluation of Systems (QEST)*, pages 201–202, 2012.
- [19] The Scottish government. Bus Punctuality Improvement Partnerships (BPIPs) guidance document. 2009.
- [20] STV Edinburgh. Why 2700 support campaign against 20mph speed limits. <http://edinburgh.stv.tv/articles/309605-opinion-from-grant-maccusker-founder-of-campaign-say-no-to-20mph/>. Published: 09-02-2015. Accessed: 06-09-2015.
- [21] The Scotsman. Capital bus firm fined for running early. <http://www.scotsman.com/news/capital-bus-firm-fined-for-running-early-1-819225>. Published: 29-07-2010. Accessed: 06-09-2015.
- [22] Axel Thümmler, Peter Buchholz, and Mikl’os Telek. A novel approach for phase-type fitting with the EM algorithm. *Dependable and Secure Computing, IEEE Transactions on*, 3(3):245–258, 2006.

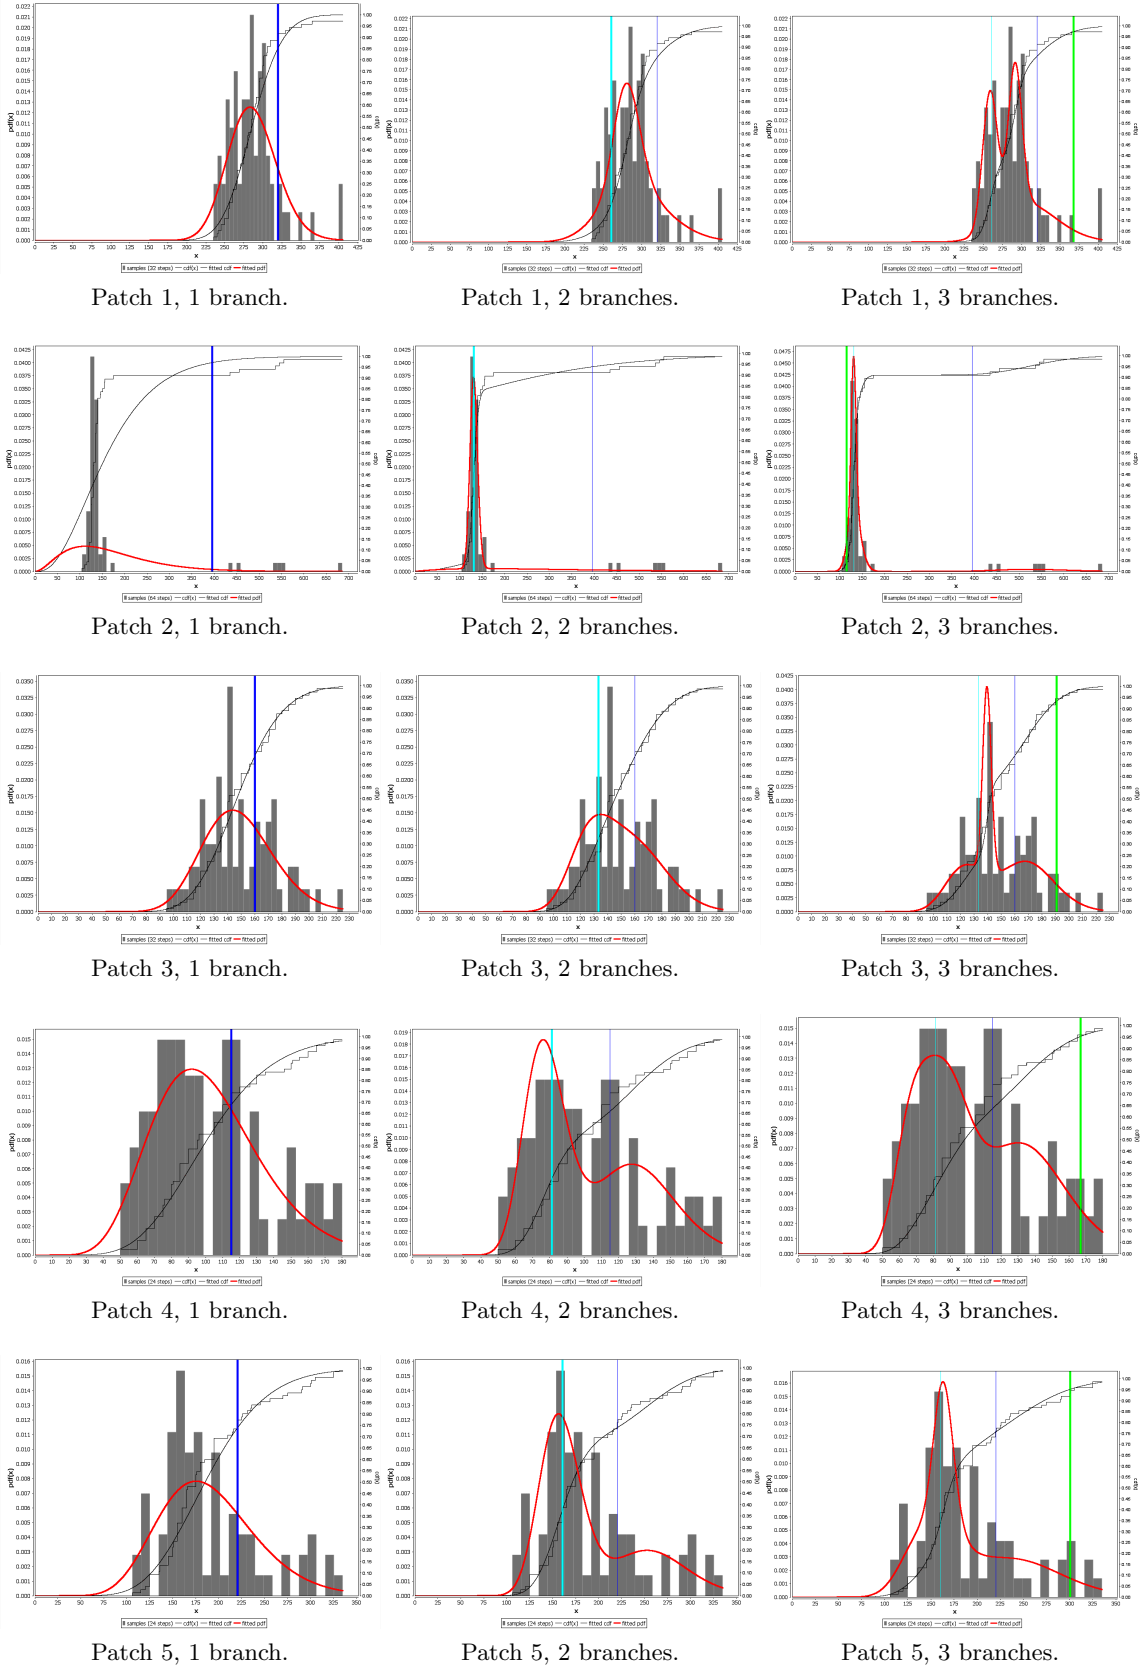





Figure 6: Hyper-Erlang parameter fitting results for Patches 1-5 of the Airlink route.




Patch 6, 1 branch.




Patch 6, 2 branches.




Patch 6, 3 branches.




Patch 7, 1 branch.




Patch 7, 2 branches.




Patch 7, 3 branches.




Patch 8, 1 branch.




Patch 8, 2 branches.




Patch 8, 3 branches.




Patch 9, 1 branch.




Patch 9, 2 branches.




Patch 9, 3 branches.



Patch 10, 1 branch.



Patch 10, 2 branches.



Patch 10, 3 branches.

Figure 7: Hyper-Erlang parameter fitting results for Patches 6-10 of the Airlink route.