

This item is the archived peer-reviewed author-version of:

Impact of introducing road charging on supporting mobile data networks

Reference:

Vandenberghe W., Bergs Johan, Carels D., van den Wijngaert Nik, Van de Velde Erwin, Blondia Christian, Moerman I., Demeester P., Dedene N.- *Impact of introducing road charging on supporting mobile data networks*

Proceedings of the Ninth International Conference on ITS Telecommunications, Lille, France, 2009 - S.l., 2009

Handle: <http://hdl.handle.net/10067/845380151162165141>

Impact of Introducing Road Charging on Supporting Mobile Data Networks

W. Vandenberghe, D. Carels, I. Moerman, P. Demeester
INTEC Broadband Communication Networks
Ghent University
Ghent, Belgium

J. Bergs, E. Van de Velde, N. Van den Wijngaert, C. Blondia
Performance Analysis of Telecommunication systems
University of Antwerp
Antwerp, Belgium

N. Dedene
Division Traffic Centre
Department Mobility and Public Works
Flemish Government
Antwerp, Belgium

Abstract— This paper presents a model for assessing the impact of introducing nationwide road charging systems on supporting mobile data networks. It defines a model based on real traffic measurements, and establishes appropriate parameter values for capturing specific application characteristics such as the importance of data compression, truck-only versus all traffic charging, and dimensioning margins for handling peak days. Finally, this paper applies the developed model and defined parameter values to a Belgian case study.

Keywords - Road charging, impact study, Intelligent Transport Systems

I. INTRODUCTION

Road charging offers an alternative approach for dealing with mobility. In contradiction to classical road tax mechanisms, nationwide road charging systems can introduce fair pricing mechanisms based on the real driven distances. For the deployment of nationwide road charging systems, information systems based on vehicle probes using GPS receivers and data communication possibilities can offer a nationwide coverage without the need for substantial investments in infrastructure. This paper presents a model that can be used to estimate the impact of introducing such systems on the supporting mobile data networks. The implemented model is based on real traffic measurements, and uses this information to assess the amount of established connections that can be expected and the amount of data to be transferred.

In the following sections, this paper elaborates on the chosen model, explains the different datasets that were used and describes the correct values to be used for the model parameters. Finally, it concludes with the results of applying the model to three different road charging client scenarios (streaming thin client, store and forward thin client and thick client).

II. IMPACT ON MOBILE DATA NETWORKS

The data demand model was defined in two phases. In the first phase, the focus was to estimate how many vehicles will be present in one network cell at a given moment in time, based on the traffic measurements data. In the second phase, refinements were introduced to this model to capture more

specific characteristics of road charging, such as the possibility to make the distinction between road charging for trucks only or for all vehicles, the effects of compressing GPS coordinates, margins to handle days with more intense traffic than average, etc.

A. Simple impact model

The main idea behind the impact model is that starting from traffic intensity measurements, it is possible to estimate how many vehicles will be present in one network cell at a given moment in time. This information can be used to determine the amount of required connections in a cell, and the amount of data that will be transferred.

To estimate the amount of vehicles in one cell, we consider the measured traffic intensity i , moving at a corresponding vehicle speed v . We assume that the network cell diameter for cells along highways can be considered as a constant value, and represent this by the variable l . The behavior of the road charging application is, following the store and forward principle, defined by two parameters, d and c , where d is the distance between two connections to the road charging servers, and c is the coordinate inter-distance. These variables are depicted in Figure 1. The unit of i is vehicle per hour, v is kilometer per hour, and the units of l , d and c are kilometer.

The d and c parameters define the update and logging interval in terms of distance. Another possible approach would be to define the update interval in terms of time. This is not supported by the model, because when standing still in traffic jams, that approach produces unnecessary data and establishes unnecessary connections with the road charging servers. It also sets up superfluous connections containing none or very little data when the vehicle has (almost) not been used.

To define a first simple model, we introduce the following concepts. $U_{1,d}$, is the amount of updates to the road charging servers that one vehicle performs per kilometer. Then

$$U_{1,d} = 1 / d \text{ (updates/kilometer)/vehicle.} \quad (1)$$

The number of updates that one vehicle performs per hour, $U_{1,t}$ can be defined as

$$U_{1,t} = U_{1,d} * v = v / d \text{ (updates/hour)/vehicle.} \quad (2)$$

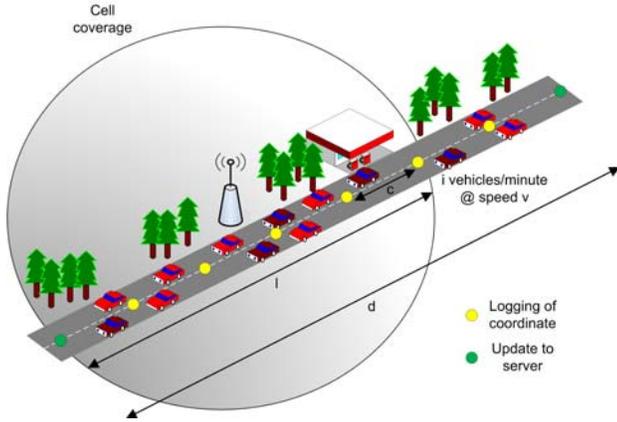


Figure 1: Variables simple model

The time that one vehicle spends in one network cell, $T_{1,c}$, is the length of the cell divided by the speed of the vehicle:

$$T_{1,c} = l / v \text{ hour.} \quad (3)$$

Then V_c , the average number of vehicles in one cell, can be defined as the product of the amount of vehicles entering the cell per hour and the number of hours that one vehicle stays in that cell. Thus

$$V_c = i * T_{1,c} = i * (l / v) \text{ vehicles.} \quad (4)$$

Finally, the number of update connections performed per hour in one cell, U_c , can be defined as the multiplication of the amount of vehicles in that cell with the amount of updates performed per hour per vehicle:

$$U_c = V_c * U_{1,t} = i * (l / v) * (v / d) = i * (l / d) \text{ updates/hour.} \quad (5)$$

To calculate D_c , the amount of road charging data that will be transferred per minute in one cell, it is necessary to define the amount of data from one update to the road charging server, D_1 . To define this parameter, the following concepts are introduced: the connection setup overhead O (in bytes per update), the bytesize of one GPS coordinate b (in bytes per coordinate) and the number of GPS coordinates that will be transferred to the road charging servers during one update, C_1 (in coordinates/update). This parameter is equal to the division of the distance between updates by the distance between two logged GPS coordinates:

$$C_1 = d / c \text{ coordinates/update.} \quad (6)$$

Taking these variables into account, the total data of one update D_1 can be defined as

$$D_1 = O + (C_1 * b) = O + ((d / c) * b) \text{ byte/update.} \quad (7)$$

Then D_c , the amount of road charging data that will be transferred per minute in one cell, can be defined as the multiplication of the amount of updates per minute in that cell, with the amount of data in one update:

$$D_c = U_c * D_1 = (i * (l / d)) * (O + ((d / c) * b)) \text{ byte/hour.} \quad (8)$$

B. Extended impact model

The above model is a good starting point, but it however lacks some important concepts for road charging. It should be able to model both road charging for trucks only, and road charging for all vehicles. Because the model uses traffic intensity measurements that are averaged out over many days, it should be able to include a dimensioning margin to support peak days. Similarly, the model should be able to compensate the difference in traffic intensity between older traffic measurement data and the current and future traffic situation. And because compression can have a significant impact on the amount of transferred data, this should also be included in the model.

Modeling road charging both with only trucks or with all vehicles is quite straightforward by introducing the parameter p , being the percentage of all vehicles that are equipped with road charging devices. Then U_c becomes

$$U_c = i * (l / d) * p \text{ updates/hour.} \quad (9)$$

The value of p has to be situated between 0 and 1. The p parameter can also be used to model other scenarios, e.g. a mandatory installation in all vehicles ($p = 1$), or a voluntary installation where the p will vary in time according to an appropriate adoption model.

The traffic intensity i used in the model originates from real traffic intensity measurements. It is possible, even likely that those values were averaged out over several days, weeks, months or even years. However, a peak traffic day can have a substantial larger amount of traffic than the average day. The supporting mobile data network should be dimensioned to support such busy days to an acceptable degree. Therefore we introduce the margin parameter m in the model. It is a factor that is multiplied with U_c , and should have a value greater than or equal to 1. U_c now becomes

$$U_c = i * (l / d) * p * m \text{ updates/hour.} \quad (10)$$

The consequence of using real traffic intensity measurements as an input for the model is that the results are only applicable for the moment in time that the data was gathered. In the case that only (relatively) old measurement data is available, the model should be able to compensate this, producing results that are representative for the current traffic situation. Similarly, it should also be possible to model future traffic situations, in case the model will be used to define more long term strategies related to road charging and its deployment. When taking the assumption into account that traffic will evolve similar over the entire road network, the model can support these requirements by introducing the compensation parameter co . Similar to the m parameter, it is a factor that is multiplied with U_c . Appropriate values for the co parameter can be derived from existing traffic evolution studies, taking the input dataset and scenario into account. U_c now becomes

$$U_c = i * (l / d) * p * m * co \text{ updates/hour.} \quad (11)$$

To model the effect of data compression, the zip factor z is introduced. This constant value is defined as the division of the

compressed file size by the original file size. Therefore, it will be a value smaller than or equal to 1. D_1 then becomes

$$D_1 = O + ((d / c) * b * z) \text{ byte/update.} \quad (12)$$

To summarize, in the extended impact model the number of update connections performed per minute in one cell is

$$U_c = i * (l / d) * p * m * c_o \text{ updates/hour,} \quad (13)$$

and the amount of road charging data that will be transferred per minute in one cell is

$$D_c = (i * (l / d) * p * m * c_o) * (O + ((d / c) * b * z)) \text{ byte/hour.} \quad (14)$$

III. BELGIAN CASE STUDY

A. Used data sets

The impact model requires traffic intensity measurements as an input. To assess the impact of introducing nationwide road charging in Belgium, two datasets from the Flemish Traffic Centre were used. The first one is publically available and contains for all included measurement locations the intensity values for every hour, averaged out over one year, 2007 [1]. The second dataset was a private set, containing the raw measurement data for all locations for every 10 minutes in the period October – November 2008. The goal of the private dataset was to validate the public dataset and to capture peak traffic behavior that is averaged out in the public dataset. This way, the model can rely on the public dataset, and can be used by other partners or for other studies (e.g. the determination of realistic traffic density requirements for ITS Cooperative Systems), for which the private dataset was not intended.

B. Parameter values

In the extended model, several variables were defined to represent specific road charging concepts. In this section, appropriate values for these different parameters are determined.

1) Compensation factor

As mentioned in section II.B, the compensation factor c_o enables the usage of (relatively) older datasets to model current or future network impact. The goal of this section is to define a correct value for the c_o parameter so that the public dataset (which contains average values over 2007) can be used to create results that are compliant with the private datasets (which contains values from October – November 2008). The only assumption that is taken is the fact that traffic evolves more or less equally in time over all locations.

To determine this appropriate value for the compensation factor c_o , the average traffic intensity over all locations was compared per hour between the public and the private dataset. When averaged out over all hours, the result is a value of 0.5%, meaning that the total traffic intensity (averaged out over all locations and all hours) of the 2008 private dataset is just 0.5% higher than that value for the 2007 public dataset. Thus at first sight, both datasets represent the same traffic situation, and can be interchanged freely, without the need for any compensation.

To found this statement more, the comparison per hour was also further studied. These values are depicted in Figure 2. It can be seen that during the more traffic intense hours, after U8, the datasets are very similar, with a difference ranging between -5 and +5 percent. Larger differences can be found from U1 to U8. However, these are of less importance since their absolute traffic intensity values are lower than the rest of the day, and can thus be neglected since the model is used for network dimensioning targeting the traffic intense moments of the day. The fact that these hours have less intense traffic is depicted in Figure 3, which illustrates the averages per hour (over all locations) for both datasets. It can be concluded that when using the public dataset as an input to the model, the value for the compensation factor c_o should be 1 to model the period of the private dataset, October – November 2008.

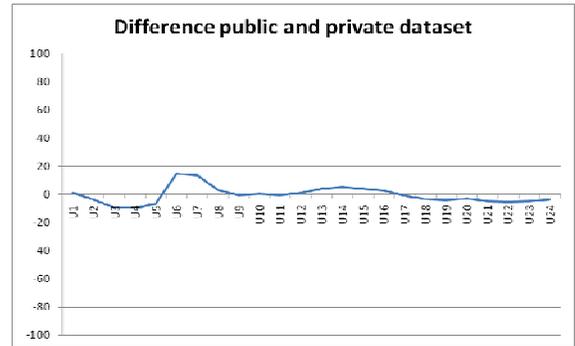


Figure 2: Difference between public and private dataset in terms of percentage

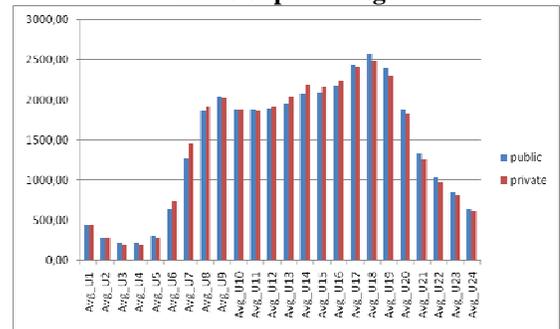


Figure 3: Average hour values public and private dataset

2) Busy day margin

The implemented model uses traffic intensity measurements that are averaged out over a longer period, per category (all days, weekdays, Saturdays or Sundays). This approach has the advantage that it gives a profound view of the traffic intensity that can be expected, less influenced by special circumstances. However, to an acceptable extent, the network should be able to support days with a more intense traffic than on the average day. Therefore, the busy day margin parameter m was introduced in the model (see section II.B).

To determine an appropriate value for the m parameter, an analysis was performed on the private dataset. Because this set contains the real measured traffic intensities with a granularity of 10 minutes, over a period of 2 months and over the entire Belgian road network, it offers all information necessary for a good assessment of the m parameter.

The first step in the analysis was to determine for each location and per hour, the difference in traffic intensity between its busiest day in the dataset, and the average over all days. These values were then averaged out over all locations. The result determines for every hour, the average extra traffic on the busiest day, compared to the traffic on the average day. This result is depicted in Figure 4.

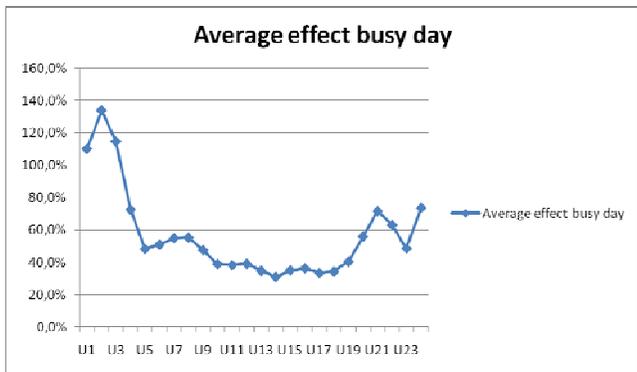


Figure 4: Average effect busy day

The average effects displayed in Figure 4 indicate that during the peak traffic hours and averaging out over all locations, the busiest day was approx. 35% more intense than the average day. This means that the network should be able to handle 1.35 times the average traffic.

The above analysis is based on the difference between the busiest day on a location, and its average day. The other way around, it is also possible to start from a set of given m values, and calculate the number of days which had even more intense traffic than the multiplication of the average value with the value of the m parameter. Such an analysis was performed with m values ranging from 1 to 1.5, increasing in steps of 0.05 and focusing on the peak hour of the day, 18 h. The results are depicted in Figure 5. With an m value of 1, meaning that only the average traffic intensities are supported, the network was not able to support the demanded data load in 45% of the days. This clearly indicates the need for the m parameter in the model. Again, a suitable m value is 1.35, in which case just 3% of all days knew more intense traffic. This corresponds to less than two days in two months.

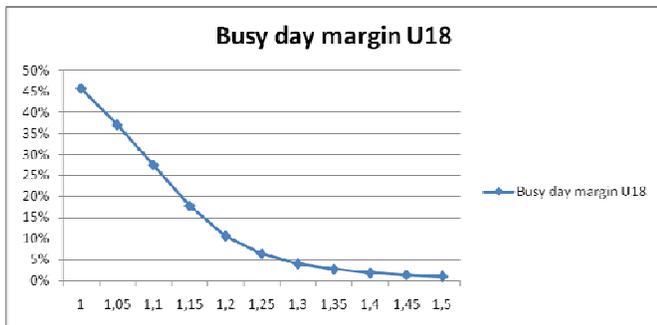


Figure 5: Percentage of busier days for different m values

3) Percentage equipped

As mentioned in section II.B, the model supports the distinction between equipping all vehicles with road charging

devices, or only trucks, by adjusting the percentage equipped parameter p . The same parameter can also be used to investigate different introduction scenarios, e.g. a mandatory installation in all vehicles ($p = 1$), a mandatory installation in new vehicles (p increases in time according to amount of new vehicles sold), or a voluntary installation where the p will vary in time according to an appropriate adoption model. This section determines the value for two scenarios: mandatory installation in all vehicles, or in trucks only. It is quite obvious that to model the mandatory installation in all vehicles, the value of p has to be 1. The other scenario requires a more profound analysis.

For this analysis, the Flemish Traffic Center provided a table that contains the percentage of trucks in the total traffic between 6h and 22h, for 100 of the most relevant locations. The first step in the processing of this data was to calculate some very simple statistics over all locations: the average, minimum and maximum. In average, 18.3% of the measured traffic at the different locations between 6h and 22h were trucks. The lowest value found over all locations was 7.6%, while the highest was 40.1%. Based on this average value, a possible choice for the value of p could be 0.18. However, because of the large discrepancy between the average, the minimum and the maximum, a more thorough analysis of the truck percentage dataset is required before concluding on the most appropriate p value.

In this analysis, for a set of p values, the percentage of locations with a truck intensity smaller than or equal to that p value was calculated. This indicates the percentage of locations where the trucks traffic is not underestimated. The results are depicted in Figure 6. It can be seen that for a p value of 0.18, the amount of truck traffic was not underestimated in just 60% of all locations, which is unacceptable. However, when using a p value of 0.3, truck traffic was not underestimated in 95% of all locations, while further increasing p only has a very small positive effect. This result implies that for the Belgian use case, the most suitable p value is 0.3 when modeling a rollout for trucks only.

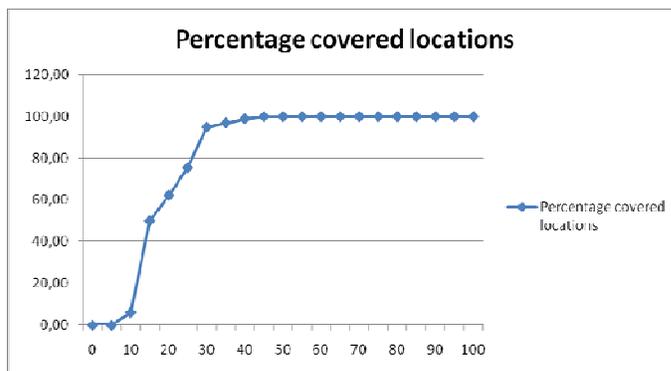


Figure 6: Percentage of locations not underestimating the amount of trucks for different p values

4) Bytesize and zip factor

To estimate the amount of data that would be generated by a vehicle equipped with an on-board road charging system, we started by making some GPS logs in NMEA format. The logs were generated by a portable GPS device and comprise of large

stretches of highway as well as city traffic, over a total length of approximately 92 km. Because data in NMEA format contains more information than required for the road charging application, we started by filtering the NMEA entries and extracting only the values needed for the road charging application: date, time, longitude and latitude. In this form one GPS log entry is a line of ASCII characters that looks like this: "070740,5112.724,N,00430.982,E,020209". In this example, the log was taken on February 2nd 2009, 7 minutes and 40 seconds past 7 am, with a latitude of 51 degrees and 12,724 minutes north, and a longitude of 4 degrees and 30.982 minutes east. In this format, the bytesize b of one coordinate is 37 bytes (36 characters + 1 newline character). This size could be further reduced by storing the logs in a binary form instead of the ASCII representation. E.g. the time could be represented by one value, the number of seconds between midnight and the current time, in the range between 0 and 86400. This value can be represented with 17 bits. For the date, the day needs 5 bits, the month 4 bits, and the year 12 bits. So the date can be stored using 21 bits. When separately storing the degrees, the minutes and the thousands of a minute, latitude needs 24 bits and longitude 25 bits. Adding this all up leads to the conclusion that the binary representation of one GPS coordinate has a size of 87 bits, or 11 bytes (with 1 padding bit). However, lossless compression of this binary representation will not be significant, since most compression algorithms are dictionary based, and perform bad on data with little recurrences. Therefore we chose to keep the ASCII representation, with a value for the bytesize parameter b of 37 bytes.

To compress this data, we defined an approach containing two steps. The first step is to transform the ASCII file, so that only differences in comparison with the previous entry are stored. The first value of the file is represented entirely, with a Unix timestamp, a longitude and a latitude value, while the following entries only store the difference in comparison with the previous entry. To avoid diversity between the full representation and the diffed representation of an entry, a full version is stored every nr_diffs entries. This new variable defines how the logged route will be split up in subroutes which are stored in separate files, each containing one complete value and a set of diffed values.

The next step is to compress all subroute files using standard lossless compression algorithms. Because the subroute diff files from the previous step should contain a lot of data recurrences, they should perform well. In our implementation, we used standard bzip. Besides the advantage that this approach leads to high zip factors, it also makes sense from a privacy point of view. It is not inconceivable that privacy regulations would demand mechanisms to ensure that no entity can have a view on the complete trips logged for road charging. This could be solved by having multiple billing centers, and one service broker. The different compressed subroutes could then be encrypted using different public keys corresponding with the different billing centers, combined in a single archive and sent to the service broker which distributes the different subroutes to the appropriate billing centers.

To determine the zip factor corresponding with this two-step approach, we investigated the effect of two variables: nr_diffs and c . The nr_diffs value influences the size of the

subroute files from the first step, while the c value, the distance between two logged coordinates, influences the data recurrence and thus the performance of the second step. Applying our compression algorithm on our logged dataset, with a fixed c value of 100 meter and a nr_diffs value varying between 10 and 1000, resulted in a zip factor between 0.4 and 0.15. When performing a power fit to these results, we obtained the following formula:

$$z = e^{(-0.1688127)} * nr_diffs^{(-0.3037247)} \quad (15)$$

To assess the influence of the c variable, we performed our compression on the same dataset, this time with a fixed value for the nr_diffs value (100), and a varying c value ranging from 50 to 500 meters. This resulted in a z value between 0.15 and 0.22, indicating a slight effect of c on the zip factor. To include this effect in (15), for every c value the ratio was calculated between the corresponding zip factor, and the zip factor for a c value of 100 meters (since (15) was calculated with a c value of 100). Then for c values between 50 and 500 meters, the zip factor was calculated for nr_diffs values between 10 and 300 by multiplying the result of (15) with the corresponding ratio. This result is depicted in Figure 7. It shows that lowering c has a positive impact on the compression, with most effects in the range between 50 and 200 meters. Increasing nr_chunks also has a positive effect, although once above a value of 100, this effect is not very important. It can be concluded that with a c value lower or equal to 200 m, and a nr_chunks value greater than or equal to 100, the zip factor will be situated between 0.1 and 0.3.

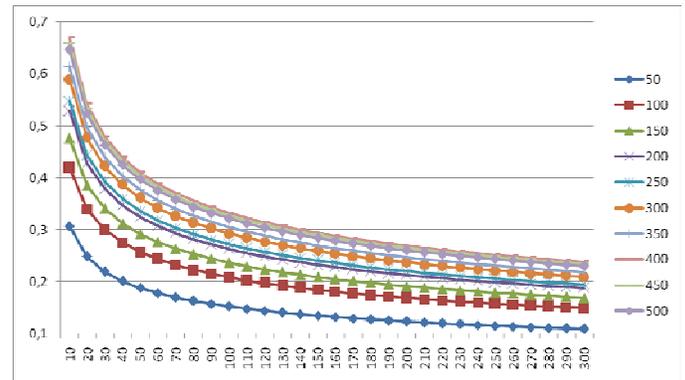


Figure 7: Influence of c and nr_diffs on z

C. Results

In this section the results are presented from running the model with the variables set defined in the previous section, for three different scenarios: a thin client streamer, a thin client store and forward client, and a thick client. The thin clients log GPS coordinates and send them to the road charging billing centers which will calculate the cost. The streaming variant uploads every coordinate immediately, while the store and forward variant saves a number of coordinates locally, and sends them to the server in one batch. The thick client calculates the cost itself, and only sends the final cost to the road charging billing servers.

For the thin client streamer, we used the following variable values: $l = 2.4$ km (an appropriate value for Flanders according to a network operator partner in the project), $d = 0.1$, $c = 0.1$, $b = 37$, $O = 1500$, $p = 0.3$ and l (trucks only and all vehicles), $z = 0.2$, $co = 1$ and $m = 1.35$. This resulted in an average of 50 000 connections and 600 Mbit per hour in one network cell at 18h when only trucks would be equipped, and 167 000 connections and 2000 Mbit with all vehicles. It is obvious that these values are too high, and that the thin client streamer would cause an unacceptable load on the network.

For the thin client store and forward, all values remained the same, except the d value which was varied between 1 and 500 km in steps of 25, and the z value which was adjusted accordingly (1-100 : 0.15, 100-200: 0.23, 200-300: 0.27, 300-500: 0.31). The results are shown in Figure 8 and Figure 9. The values range between an average of 17 000 connections and 207 Mbit per hour in one network cell at 18h, with a d value of 1, and 33 connections and 16 Mbit per hour with a d value of 500 for the scenario with all vehicles. Values for trucks only are about a third, the maxima instead of the averages are about the double.

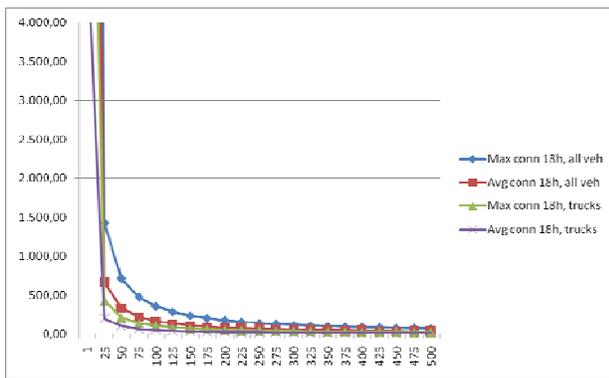


Figure 8: Connections per cell at 18h for varying d values

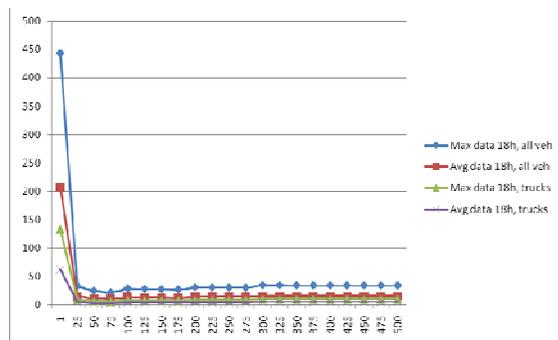


Figure 9: Data per cell at 18h for varying d values

It is quite obvious that the chosen d value has a significant impact on resulting network load. From the perspective of the amount of connections, d should be situated somewhere between 50 and 200 km. A lower value causes too much load, while a higher value has almost no extra gain and introduces an unnecessary delay between logging and updating to the server. From the perspective of the amount of data to be transferred over the network, the only restriction is that d should be greater than or equal to 25 km. A lower value introduces too much

data, while a higher value has almost no extra gain. Thus, it can be concluded that for the thin client store and forward scenario, the value of d should lie between 50 and 200 km, resulting at 18h in one cell in an average amount of connections between 80 and 330 per hour, together transferring between 11 and 14 Mbit of data per hour.

The thick client calculates the cost for a trip locally, and only uploads this one value to the road charging servers. To emulate this behavior in the model, the following variable values were used: $c = d$, $b = 8$ (because one double variable is 8 bytes) and $z = 1$ (it makes no sense compressing such a low amount of data). For d , the same values were used that were recommended for the store and forward client, between 50 and 200 km, all other variables remained unchanged. This resulted in the same amount of connections as for the store and forward client. This is logical, since we used the same intervals between updates to the server, d , for both cases. The only difference between them is that the thick client will need to transfer less data over each connection. Indeed, the model for the thick client resulted in the need to transfer in average between 1 and 4 Mbit of data in one cell in one hour, at 18h. This lower amount of data is the main advantage of the thick client in comparison with the store and forward version. The downside is that if the logic for calculating the cost of a trip needs to be changed, all clients should receive software updates in a coordinated and timely manner, what could turn out to be a challenge.

IV. CONCLUSION

In this paper we presented a model to assess the impact of introducing nation wide GPS based road charging systems on the supporting mobile data networks. This model was applied in a Belgian use case, taking real traffic intensity measurements as an input. For this use case, the appropriate values for the different variables in the model were derived, and applied to three different kind of road charging applications: a streaming thin client, a store and forward thin client, and a thick client. The results showed that only the streaming thin client is not feasible. They also indicated that when defining an implementation approach, the resulting impact is mostly influenced by three design decisions: the selection of an appropriate interval between two updates to the server, the choice between a store and forward thin client and a thick client, and the choice to introduce road charging for all vehicles, or for trucks only.

ACKNOWLEDGMENT

The authors would like to thank the Flemish Interdisciplinary institute for Broadband technology (IBBT) for defining the NextGenITS projects that funded the research presented in this paper. The authors would also like to thank the Flemish Traffic Centre for their extensive support.

REFERENCES

[1] Verkeerstelling 2007, Agentschap wegen en verkeer, <http://wegen.vlaanderen.be/documenten/tellingen/20>