Data Transmission Plan Adaptation for Connected Vehicles

Tobias Rueckelt[o*], Ioannis Stavrakakis[+], Tobias Meuser[*], Doreen Böhnstedt[*], Ralf Steinmetz[*]
[*]Multimedia Communications Lab (KOM), Technische Universität Darmstadt, Germany
[+]National and Kapodistrian University of Athens, Greece & IMDEA Networks & UC3M
[o]Adam Opel AG, Rüsselsheim, Germany
tobias.rueckelt@kom.tu-darmstadt.de

***Abstract*—Connected vehicles can nowadays be equipped with multiple network interfaces to access the Internet via a number of networks. Given a network availability map and a vehicle's route, a joint time-network selection plan could attempt to schedule delay-tolerant data transmission efficiently. Since the networks presence and available capacities change rapidly and in an unforeseen manner (because of changing conditions due the uncertainty in car movement, data transmission needs and network characteristics), the efficiency of such a planning may be severely compromised. Consequently, mechanisms that determine the deviation from the foreseen conditions are derived in this paper[1], which modify the transmission plan depending on the type of deviation observed. Simulation results show that the proposed mechanisms help maintain the benefits of a joint time-network planning under changing conditions.**

I. Introduction

Nowadays, mobile nodes typically integrate different wireless network interfaces. To improve connectivity performance, they may use these network interfaces in parallel to distribute their data traffic. Moreover, the connected vehicle use case provides an additional optimization potential, especially considering automated vehicles: Routes are usually known and, thus, movement can be predicted accurately. As a result, a vehicle can predict future network availability and characteristics using the so-called connectivity maps [6]. Furthermore, according to Sandvine [8], a major part of a mobile node's data traffic is delay-tolerant or heavy-tailed. Assuming networks and data traffic to be roughly known for a certain time horizon, we show in prior work [7] that a transmission planning can provide significant benefits. The transmission planning approach combines network selection [2] with a selection of the transmission time [3]. Accordingly, the approach plans ahead data transmission over multiple networks. However, the therein-presented approach assumes perfect prediction of vehicle movement, network characteristics and data to transmit. Such accurate prediction might not always be available. Accordingly, we present two contributions in this paper:

1) An investigation of the effects of erroneous prediction on the performance of transmission plan execution
2) A transmission plan adaptation that can mitigate a negative impact of erroneous prediction

[1]This work has been funded in part by the German Research Foundation (DFG) within the Collaborative Research Center (CRC) 1053 - MAKI

In Section II, we briefly outline our previous work on data transmission planning assuming static conditions and perfect knowledge of the environment and compare it to an Opportunistic Network Selection (ONS) (no planning or prediction). Furthermore, we introduce our prediction error models and show that the performance of the joint time-network selection approach degrades severely in the presence of prediction errors, due to its inability to react to changing conditions. In contrast, the opportunistic approach (although underperforming with respect to the transmission plan approach under static conditions and perfect knowledge) appears to deliver a constant performance in the presence of prediction errors. This provides the motivation for our proposed transmission plan adaptation mechanism, presented in Section III. The benefit of each planned transmission is re-evaluated and the plan is modified by invoking a *constrained* ONS taking into account the type and magnitude of condition changes (i.e., car movement, data flows or network characteristics); mechanisms detecting relevant condition changes are also introduced. In Section IV, we discuss the performance of our novel adaptation approach under various changing conditions, followed by a related work discussion in Section V. It turns out, that our approach can largely sustain the gain foreseen from long-term planning, under small to moderate changes in the environment.

II. Data Transmission Planning

The predictable movement of multi-homed mobile clients enables a transmission planning over networks and time. In our prior work [7], we demonstrate significant benefits of such a planning in comparison to state-of-the-art approaches. In this section, we summarize the approach, the evaluation metrics and results of this work. This constitutes the base for the adaptation approach proposed and evaluated in this paper.

A. Evaluation Metrics and Model of Forces

To assess the efficiency of our joint time and network selection approaches, we developed a performance rating function, that captures application QoS requirement satisfaction and monetary cost. We bisect the performance rating function into two components that are in effect in a mutually exclusive manner depending on whether data is allocated or not. We call the first component the attracting forces c_{attr}. It captures

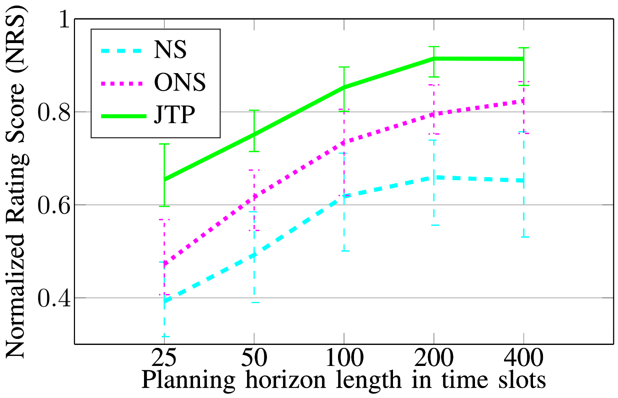

Fig. 1. Normalized Rating Score (NRS) of the transmission planners for different planning time horizon lengths

cost associated with data that is not allocated to a network, referred to as p^*, punishing the violation of a minimum throughput requirement. The second component, referred to as the repelling forces c_{rep}, captures cost associated with data that is allocated to a network, punishing the violation of the QoS requirements of the data flows (deadline, latency, jitter, preferred start time, etc.) or monetary transmission cost. The rating function in Equation (1) adds the two mutually exclusive components for a given transmission plan p.

$$c(p) = c_{rep}(p) + c_{attr}(p^*) \tag{1}$$

Minimizing this cost function leads to the following effects: Networks attract data for allocation in general through c_{attr}, creating attracting forces for each data flow according to its priority. In addition, the repelling forces push data away from networks and time slots that cannot satisfy the data flow's QoS requirements. This results in data allocation to the best matching networks at matching points in time over the complete planning time horizon. For the detailed model of the cost function, refer to [7].

As the absolute value of the cost in Equation (1) strongly depends on the scenario, a Normalized Rating Score (NRS) is introduced to allow for a meaningful comparison of multiple scenarios. *NRS describes a transmission plan's used share of the absolute optimization potential of the given scenario.* A value of 0.8 means that a transmission plan uses 80% of the scenario's optimization potential. To define the optimization potential, we employ an upper and a lower bound. As a lower cost bound, we use the cost of an optimal transmission plan. As an upper cost bound, we use the average cost of random transmission plans. We assume this as a reasonable lower bound because no transmission plan, which was created with intent, should perform worse than random.

B. Transmission Planners

Transmission planners steer data allocation to networks and time to optimize the data transmission characteristics. We analyze three transmission planners from [7] in this paper. All of them use the same ratings for network selection and data flow prioritization. However, they differ in the way they handle the time dimension.

The first is a **Network Selection (NS)** derived from state-of-the-art approaches. It allocates data to the currently available networks ignoring the time dimension. It prioritizes data flows and decides for each one, which currently available networks are best suited for its transmission. Finally, it allocates data according to these priorities.

As a second approach, we present an **Opportunistic Network Selection (ONS)**. It extends Network Selection considering an opportunistic component, which decides whether to transmit data (as the NS would dictate) or not. This decision is based on an estimated benefit for data allocation, defined as the difference between the repelling (in case data is transmitted) and the attracting (in case data is not transmitted) forces. Whenever the benefit exceeds some threshold c_{lim}, the approach allocates data to the network. This is shown in Equation (2), showing the cost difference for a specific data allocation of data flow f at time slot t to network n. Rejecting non-beneficial transmission at the current point in time amounts to waiting for a better opportunity to transmit. This leads to a statistical time selection.

$$c_{attr}(p^*_{f,t,n}) - c_{rep}(p_{f,t,n}) > c_{lim} \tag{2}$$

As a third approach, we present the **Joint Transmission Planning (JTP)**, as introduced in [7]. Instead of considering currently available networks only, JTP selects the best transmission opportunities within the complete planning time horizon. It plans data allocation ahead to a time, in which the transmission is expected to be most beneficial. Hence, it represents a joint time-network selection, which handles the time dimension explicitly. However, the approach requires a prediction of network availability and characteristics, of client movement and of the data to transmit. Figure 2 visualizes the characteristic behaviors. It shows the benefit over time for allocating data as a black thin curve. Allocating data with higher benefit leads to a lower cost function value, representing a better transmission. In addition, the figure contains lines for the cost-benefit thresholds c_{lim} of ONS. ONS sets $c_{lim} = 0$ by default (magenta dotted line), thus allocating data at the earliest point in time offering a transmission benefit. In the example, this corresponds to the first benefit 'hill'. In contrast, JTP allocates data during the highest benefit (green); in the example the second 'hill'. Figure 1 presents the Normalized Rating Score results of the transmission planners for different planning time horizon lengths. JTP uses up to 91.5% of the scenario optimization potential, significantly outperforming the two state-of-the-art approaches, NS and ONS, which use only up to 65.2% and 82.3%, respectively.

The results for the JTP approach in Figure 1 are derived assuming perfect knowledge of the environment, i.e. network resources and demand characteristics. As this can hardly be the case in real environments, predictions about the state of the environment in the future will not be perfect. In this paper, we first introduce prediction error models and analyze the impact of errors on the performance of the planners. Then, we propose a transmission plan adaptation approach that takes into account prediction errors, which is the main contribution of this paper.

III. Adaptation of Transmission Plans

Transmission plans are applicable whenever prediction is correct. Nevertheless, what does happen if the prediction used for transmission plan creation is erroneous? In this section, we analyze prediction error types of the connected vehicle use case and design a novel adaptation approach with the goal of robustness against this kind of uncertainty.

A. Prediction Errors

In a connected vehicle environment, transmission plans are derived based on some predictions on the vehicle movement, the encountered network characteristics and the data to be transmitted. As such predictions may not be correct, it is important that resulting prediction errors are calculated and some adjustments in the transmission plan are made. The Symmetrical Mean Absolute Percentage Error (SMAPE) [4] is employed to measure those prediction errors. The movement prediction error mainly affects the availability of networks. For example, a vehicle, which moves faster than expected, may reach a small range network earlier and may spend less time in its covered area. We measure the error in the number of time slot drifts over the planning horizon. The network characteristics prediction error affects the throughput, latency and jitter of the networks over time. Finally, the data flow prediction error arises from canceling or pausing running data transmissions or from unexpected new data transmissions. Next, we present our adaptation approach handling these three types of errors.

B. Adaptation Approach

The idea of our adaptation approach is to use a *constrained* ONS whose decision threshold is determined according to environmental changes so that the data transmission plan that is actually implemented is still beneficial. First, we design a transmission plan execution algorithm, which constraints ONS to implement the initial transmission plan, when no environmental changes occur. Second, to allow ONS to adapt the plan as a reaction to environmental changes, we present three adaptation mechanisms that dynamically relax the constraints and modify parameters of the first algorithm.

Execution Algorithm: To follow the transmission plan, this algorithm suppresses each data transmission of ONS, which does not comply with the plan. Therefore, the mechanism increases the benefit threshold c_{lim} of ONS to the flow's maximum benefit value $c_{max}(f,t)$, defined as the supremum of its attracting force according to Equation (3). When, in contrast, data is allocated in the initial plan, it sets the threshold to the flow's minimum benefit value $c_{min}(f,t)$, defined as the infimum of its repelling forces, which is based on the highest requirement violations from the currently available networks N_0 according to Equation (4). To decide which of the two threshold values ONS should use, the execution algorithm compares the amount of released data $p^{rel}(f,t_0)$ from Equation (5) to the actually allocated data $s^{alloc}(f,t_0)$ of flow f from Equation (6) at the current time slot t_0 according

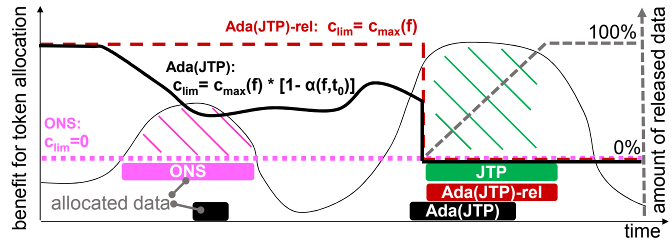

Fig. 2. Schematic of basic concept

to Equation (7). While the released data $p^{rel}(f,t_0,n)$ does not change over time within in the considered time slot, the value of the actually allocated data s^{alloc} is refreshed continuously, stopping the allocation as soon as the amount of the data planned for the current time slot is reached.

$$c_{max}(f,t) = \sup_{t \in T} c_{attr}(p^*_{f,t,n}) \tag{3}$$

$$c_{min}(f,t) = \inf_{n \in N_0} c_{rep}(p_{f,t,n}) \tag{4}$$

$$p^{rel}(f,t_0,n) = p_{f,t_0,n} \tag{5}$$

$$s^{alloc}(f,t_0,n) = s_{f,t_0,n} \tag{6}$$

$$c_{lim} = \begin{cases} c_{min}(f), & s^{alloc}(f,t_0,n) < p^{rel}(f,t_0,n) \\ c_{max}(f), & else \end{cases} \tag{7}$$

1) Extended Data Release Mechanism: Our first adaptation mechanism addresses changes in the network characteristics. This corresponds, firstly, to changed transmission characteristics like latency and jitter and, secondly a differing throughput. Changes in network characteristics may affect the flow-network matching and preference. Strong performance degradation of networks might lead to the case in which transmission is not beneficial at all. To let the constrained ONS decide whether to transmit or not, we set the minimum threshold c_{min} to the ONS's default value 0 according to Equation (8). This restricts each data allocation to the cases for which ONS still considers a sufficient benefit. In addition, we relax constraints to employ ONS for re-evaluation of the flow-network matching and a re-selection using the actual network characteristics. Therefore, the adaptation stops distinguishing networks for releasing data, constraining all networks in the current time slot equally due to considering the sum of allocated data over all networks as shown in Equations (9) and (10).

To address unbiased throughput fluctuations, we relax the execution algorithm's limit for the amount of released data. Instead of focusing on the amount of data planned for transmission for each time slot separately, we redefine the released data $p^{rel}(f)$ according to Equation (9) to cover all data allocated in the initial plan until the current point in time t_0 plus the flow's data, which has not been allocated in the initial plan at all p^*_f. This helps the approach to cope with unbiased fluctuating network throughput and allows the constrained ONS to fill unexpected additional network resources opportunistically with initially non-allocated data p^*_f. As visualized with an example in Figure 2 using a gray dashed line at the right, it releases data for allocation according to the plan. In the case of no prediction errors, this results in setting the cost-benefit threshold c_{lim} according to the red dashed line, which lets this adaptation mechanism *Ada(JTP)-rel* (red) stay close to the initial plan.

$$c_{min}(f,t) = 0 \tag{8}$$

$$p^{rel}(f,t_0,n) = p^*_f + \sum_{t=0}^{t_0}\sum_{n\in N} p_{f,t,n} \tag{9}$$

$$s^{alloc}(f,t_0,n) = \sum_{t=0}^{t_0}\sum_{n\in N} s_{f,t,n} \tag{10}$$

2) Location reference mechanism: To cope with **movement prediction errors**, we present our corresponding adaptation mechanism, which refers to the initial plan by vehicle location instead of time. When a vehicle moves e.g. faster than predicted, it reaches and leaves short range networks earlier than expected. Compared to the prediction, location-dependent network characteristics move to another point in time. As a result, network availability is modified from the initial timeline, impacting on the network selection of the transmission plan. However, for delay-tolerant data transfers, the impact of network selection according to the plan dominates the impact of allocating data at the planned transmission time. To address this issue, we employ the following mechanism: *For delay-tolerant data flows, consider the spatial dimension of the transmission plan, i.e. the vehicle's location, and ignore the temporal one.* Referring to the spatial dimension is equivalent to a temporal offset $\epsilon_{move}(t_0)$ of the transmission plan. The location reference mechanism shifts data transmission in time by this temporal offset in order to preserve the initial network selection. However, for non-delay-tolerant data flows, e.g. interactive ones, this temporal transmission offset may lead to a requirement violation. Hence, we limit the temporal offset $\epsilon_{move}(t_0)$ to the maximum delay tolerance of the data flow, which our model from prior work [7] encodes in a throughput requirement window parameter $\Delta\widehat{t}_f^{min}$, c.f. Equation (11). Accordingly, we employ the time-limited spatial reference $t^{loc}(f,t_0)$ according to equation 12 to refer to the initial plan. This limited spatial reference preserves the initial network selection of the transmission plan for delay-tolerant data flows but accounts temporal requirements for non-delay-tolerant flows.

$$t_f^{offset} = \min(\Delta\widehat{t}_f^{min}, \|\epsilon_{move}(t_0)\|) \tag{11}$$

$$t^{loc}(f,t_0) = \begin{cases} t_0 + t_f^{offset}, & \epsilon_{move}(t_0) > 0 \\ t_0 - t_f^{offset}, & else \end{cases} \tag{12}$$

Referring to the corresponding vehicle location in the transmission plan to release data for allocation causes one problem: whenever the car stops, no additional data is released. There is no progress in the vehicle's location and, thus, the transmission pauses. This effect impairs transmission similarly when the car moves slower than expected. To address this issue, our mechanism modifies the condition for the c_{lim} threshold selection of Equation (7) to that from Equation (13). *Whenever data is allocated within the initial plan in the reference time slot, release data for transmission.* Hence, this mechanism together with the extension in triggering conditions handles movement prediction errors up to a certain degree.

$$\begin{aligned} s^{alloc}(f,t_0,n) &< p^{rel}(f,t^{loc}(f,t_0),n) \\ \text{or } &\left(\sum_{n\in N} p_{f,t^{loc}(f,t_0),n}\right) > 0 \end{aligned} \tag{13}$$

3) Flow Prediction Error Handling: Finally, our third mechanism treats **flow prediction errors**. Flow prediction errors refer to additional data to be transmitted, time shifts in data transmission and canceled data transmission. For new data, there exists no reference in the existing transmission plan. Hence, we let the mechanism release new data completely for opportunistic transmission, handling it equivalently to the non-allocated data p^*_f. This way, the opportunistic algorithm automatically prioritizes active data flows correctly integrating the new ones into the ongoing transmission. However, prediction errors concerning planned transmissions have to be treated explicitly. This covers especially planned non-delay-tolerant data flows, whose transmission time differs from the predicted one. Hence, we consider the SMAPE flow prediction error $\epsilon_{flow}(f,t)$ over the time span of the past throughput window $\Delta\widehat{t}_f^{min}$ of the flow. However, for delay-tolerant flows, a pure opportunistic transmission might lead to a worse network selection. Hence, instead of setting c_{lim} to 0, we reduce c_{max} with rising error according to Equation (14) and (15). Thus, *when flow prediction errors occur, our approach does not suppress data allocation but restrict it to opportunities in which an error-dependent benefit threshold is reached.* We illustrate the threshold adaptation $\alpha(f,t_0)$ from Equation (14) in Figure 2 as a thick black line. In the example, this results in partial earlier data allocation for *Ada(JTP)* (black).

$$\alpha(f,t_0) = 1 - \sum_{t=t_0-\Delta\widehat{t}_f^{min}}^{t_0} \frac{\epsilon_{flow}(f,t)}{\Delta\widehat{t}_f^{min}} \tag{14}$$

$$c_{max}(f,t_0) = \sup_{t\in T} c_{attr}(p^*_{f,t,n}) \cdot \alpha(f,t_0) \tag{15}$$

Conclusively, our transmission plan adaptation approach combines the advantages of Opportunistic Network Selection and Joint Transmission Planning. Thus, it allows for opportunistic transmission when high prediction errors render parts of an initial plan infeasible but can exploit the superior transmission patterns in terms of time and network selection from long-term planning. We evaluate the effects of the execution and the three adaptation mechanisms within the next section.

IV. Evaluation

To analyze the performance of the transmission plan adaptation mechanism (Ada), we assess its performance under controlled variation of the prediction errors with the above presented Normalized Rating Score (NRS) and compare it to that of the Opportunistic Network Selection (ONS) and the pure plan Execution (Exec). As additional performance reference, we show the results of Joint Transmission Planning

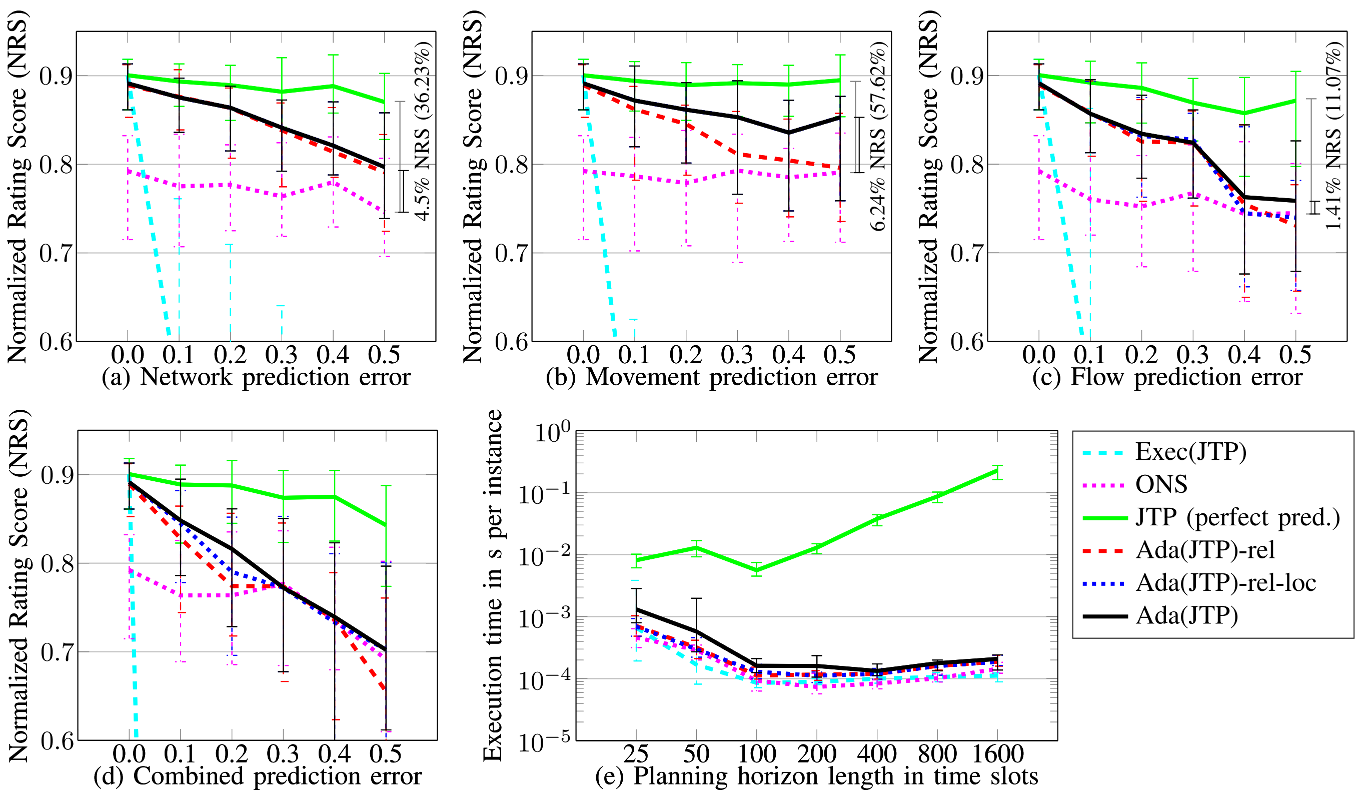

Fig. 3. Planners' NRS over SMAPE: movement, network, data flow, combined and execution duration in seconds per instance

(JTP) with perfect prediction. JTP uses this perfect prediction for all modified scenarios, independent from the defined error on the x-axis. It defines an upper bound for reference. We apply the different approaches to scenarios with 100 time slots, covering 2 cellular networks and 6 WiFi networks, which are available within the scenario's planning horizon. The number of data flows is initially 8 and varies due to flow prediction errors. We apply the different approaches to 50 randomized scenarios per run. For execution of each instance, we use a single core of a server machine with Intel Xeon E5-2643 v3 @ 3.4GHz and 512 GB RAM. To show the typical performance and its distribution, we give the $Q_{25\%}$, $Q_{50\%}$ (median) and $Q_{75\%}$ quantiles. We vary the prediction errors (SMAPE) separately for movement, network characteristics, data traffic and finally for a combined one between 0.0 and 0.5. We let the adaptation approach follow the plan of JTP and abbreviate the different mechanisms with (1) Ada(JTP)-rel: ONS with the data release mechanism (2) Ada(JTP)-rel-loc: with additional location reference mechanism and (3) Ada(JTP): the final approach covering all three mechanisms.

The results are presented in Figure 3, showing the graphs for each of the error types and the combined case. First of all, we notice that the pure execution of the plan Exec(JTP) suffers substantially from robustness. It sinks far below the performance of ONS even for small prediction errors.

In contrast, the performance of our adaptation approach stays approximately between those of JTP and ONS, starting at the JTP's and converging towards the ONS's performance for rising errors. Increasing the network prediction error (a) shows that our basic data release mechanism is able to handle this error type well. Even at an error of 0.5, the adaptation sustains a gain from long-term planning of about 4.5% NRS over ONS which corresponds to 36.23% of the performance margin between JTP and ONS. However, the basic data release Ada(JTP)-rel algorithm's performance decreases fast with rising movement prediction error (b). Adding the spatial reference algorithm Ada(JTP)-rel-loc, which we designed to cope with this error type, resolves this issue. The performance loss from movement prediction error is even less significant than for the network prediction error. It still reaches a performance surplus of 6.24% NRS over ONS, which represents 57.62% of the margin. Thus, our mechanisms are able to cope well with network and movement prediction errors. Accordingly, conserving decisions for network selection and delaying data purposefully with the data release mechanism provide effective means to keep a significant share of the planning performance gain. However, data flow prediction errors (c) impose a tough challenge. According to our third mechanism, unplanned data is transmitted opportunistically. Furthermore, the adaptation transmits data for which the desired transmission times change partially opportunistic with an error dependent threshold. Since we cannot treat new nor canceled data transmissions, the effect of this error handling is rather small. However, while the above-mentioned mechanisms drop at the level of ONS, the third mechanism is able to keep a performance benefit of 1.41% NRS, corresponding to 11.07% of the margin between ONS and JTP. Finally, we combine the prediction errors in graph (d). A value of 0.2 represents a prediction error of 0.2

for each error type at the same time. The performance loss from the three error types nearly seems to sum up and lead to a convergence to the performance of ONS at a combined error of 0.3 for the final adaptation approach Ada(JTP).

Unlike the pure execution or the partial adaptation models, our final adaptation model never falls significantly below the performance of ONS. This confirms the validity of our designed mechanisms for following a transmission plan and allowing opportunistic allocation. Furthermore, for small and medium prediction errors, our adaptation mechanism is able to preserve a major share of the performance surplus that Joint Transmission Planning promises.

The last graph (e) in Figure 3 shows the execution time per instance over the planning horizon length. We expect a performance-optimized version of our approaches to reach similar execution times on automotive target hardware. Except JTP, all approaches are online methods, which plan transmission only for the next time slot. Hence, we normalize them by the number of time slots. After a first initialization overhead, their average execution time per time slot sinks below half a millisecond on the long run. In contrast, JTP always plans the complete time horizon. Hence, its execution time rises linearly with the number of time slots. This gives motivation for the following interaction concept between JTP and our adaptation Ada(JTP): After a long-term planning of JTP, the plan is implemented using Ada(JTP). As soon as certain prediction error levels are reached, e.g. through user interaction, unexpected movement or network characteristics, planning through JTP should be triggered in the background to update the transmission plan using fresh prediction values. After initialization of the new instance of Ada(JTP), it takes over the data allocation from the previous instance. Thus, heavy prediction errors can be treated within about a second, while reaction on small and moderate unexpected events happens within less than a millisecond through our novel transmission plan adaptation.

V. Related Work

The topic of transmission planning covering joint network and time selection is barely investigated so far. Existing work in time selection reduces network selection to the WiFi-preferred principle and application QoS satisfaction to holding a deadline [5], [1]. In contrast, network selection approaches with detailed application QoS models do not consider the time dimension [2], [10]. Due to these simplifications, these approaches can apply a continuous re-planning. An adaptation of plans is not required and a handling of prediction errors gets obsolete. Nevertheless, we can learn from Bui et al. [1] that it is beneficial to separate long-term and short-term mechanisms in transmission planning. However, they apply this concept to prediction only but not to planning itself. Furthermore, it is in the nature of online network selection to apply light-weight algorithms for fast reaction to environmental changes [9]. We apply this principle also for our adaptation. The fact, which distinguishes our adaptation concept from existing work, is that we use information extracted from a long-term plan in order to control the transmission. Therefore, we develop mechanisms that recognize whether following the plan is feasible, infeasible or requires modifications, which are applied automatically.

VI. Conclusion

In this article, we investigate the impacts of environmental changes on transmission plan execution and design a plan adaptation to mitigate them. From our prior work, we know that transmission planning can create substantial benefits over state-of-the-art transmission approaches. However, in this paper, we identified that a direct execution of these plans is ineffective due its inability to react to environmental changes. To this end, we designed a novel transmission plan *adaptation* that uses an opportunistic online transmission algorithm, ensuring that the approach implements the plan whenever possible and adapting parts of the plan if new or alternative opportunities appear to be better in the *actual* environment. The performance of the adaptation shows a substantial gain of up to 10% over state-of-the-art approaches for small and medium prediction errors. With rising prediction errors, it converges towards the performance of the opportunistic approach and, unlike direct execution, does never fall significantly below its performance. Conclusively, the adaptation approach exploits the additional optimization potential from transmission planning using prediction without the risk of performing worse than state-of-the-art approaches. Thus, using the presented approach, connected vehicles can benefit from predictive transmission planning in order to improve their perceived Internet access performance.

References

[1] N. Bui and J. Widmer, "Mobile Network Resource Optimization under Imperfect Prediction," in *Proceedings of the IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM)*, 2015.

[2] M. A. Khan, U. Toseef, S. Marx, and C. Goerg, "Game-Theory Based User Centric Network Selection with Media Independent Handover Services and Flow Management," in *Proceedings of the IEEE Annual Conference on Communication Networks and Services Research (CNSR)*, 2010.

[3] J. Lee, K. Lee, C. Han, T. Kim, and S. Chong, "Resource-Efficient Mobile Multimedia Streaming With Adaptive Network Selection," *IEEE Transactions on Multimedia*, vol. 18, no. 12, 2016.

[4] S. Makridakis, "Accuracy measures: theoretical and practical concerns," *International Journal of Forecasting*, vol. 9, no. 4, 1993.

[5] F. Mehmeti and T. Spyropoulos, "Performance Modeling, Analysis, and Optimization of Delayed Mobile Data Offloading for Mobile Users," *IEEE/ACM Transactions on Networking*, vol. 25, no. 1, 2016.

[6] T. Poegel and L. Wolf, "Optimization of Vehicular Applications and Communication Properties with Connectivity Maps," in *Proceedings of the IEEE Local Computer Networks Conference WS (LCN)*, 2015.

[7] T. Rueckelt, D. Burgstahler, F. Jomrich, D. Böhnstedt, and R. Steinmetz, "Impact of Time in Network Selection for Mobile Nodes," in *Proceedings of the ACM International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM)*, 2016.

[8] Sandvine, "Global Internet Phenomena Asia-Pacific & Europe," Tech. Rep., 2015.

[9] L. Wang and G.-S. G. Kuo, "Mathematical Modeling for Network Selection in Heterogeneous Wireless Networks - A Tutorial," *IEEE Communications Surveys & Tutorials*, vol. 15, no. 1, 2013.

[10] Q. Wu, S. Member, Z. Du, S. Member, and P. Yang, "Traffic-Aware Online Network Selection in Heterogeneous Wireless Networks," *IEEE Transactions on Vehicular Technology*, vol. 65, no. 1, 2016.