Heterogeneous Vehicular Wireless Networking: A Theoretical Perspective

Hongwei Zhang
Wayne State University
hongwei@wayne.edu

Fan Bai, Xi Ju
GM Global Research & Development
{fan.bai, xi.ju}@gm.com

***Abstract*—Optimizing vehicle-centered mobility experience requires leveraging the heterogeneous wireless connectivities (e.g., cellular, WiFi, and VANET) between vehicles and the Internet. Towards a foundation for heterogeneous vehicular wireless networking, we investigate the mathematical formulation of the problem, and we analyze the impact of bandwidth aggregation on the problem formulation as well as the computational complexity of the problem. Our analysis shows that the problem can be solved in polynomial time if bandwidth aggregation is employed, but the problem becomes NP-complete and cannot be solved optimally in general if bandwidth aggregation is not employed. Our analysis is constructive such that it suggests efficient approaches to solving the problem in an optimal manner when bandwidth aggregation is employed, and it suggests both optimal and approximate approaches to solving the problem when bandwidth aggregation is not employed. The above analytical insight serves as a guidance on choosing the system architecture and algorithms for heterogeneous vehicular wireless networking.**

I. Introduction

Americans alone spend more than 500 million commute-hours per week in their automobiles, and this number keeps increasing. Thus vehicles have become an important part of people's mobile experience. To enable vehicle-centered seamless mobility, wireless networking is expected to play a central role in the next-generation vehicle systems. With the proliferation of wireless technologies and their pervasive deployments, it is expected that vehicles can leverage heterogeneous wireless networks at the same time [1], [2]. For instance, vehicles can be connected to the Internet through cellular networks, public WiFi networks, or vehicular ad hoc networks (VANETs). Optimal usage of heterogeneous vehicular wireless networks can reduce cost through free WiFi and VANETs, can enable high-throughput telematics applications through bandwidth aggregation, and can improve networking performance in the presence of dynamics and uncertainties in wireless communication.

Nonetheless, existing vehicular network applications are mostly developed using a single wireless network technology (e.g., cellular networks), and there exist major challenges in optimally utilizing multiple, heterogeneous wireless networks at the same time. *Firstly*, vehicular network applications differ in their requirements on network services, and it is non-trivial to optimally share heterogeneous wireless network resources among these heterogeneous applications. For instance, video streaming to vehicles requires high-throughput, real-time communication, whereas location-aware weather forecast can tolerate delay and low bandwidth in communication. *Secondly*, different wireless networks tend to have different properties (e.g., in throughput, delay, and reliability), and it is non-trivial to integrate them into a single system with a unified application interface. For instance, cellular links tend to have lower bandwidth but better service continuity; free WiFi links tend to have higher bandwidth but suffer from service discontinuity; VANET links/paths can have high bandwidth but is susceptible to vehicle mobility and assume complex dynamics and uncertainties.

Towards a foundation for addressing the aforementioned challenges, we focus on the theoretical aspect of heterogeneous vehicular wireless networking in this study. In particular, we mathematically formulate heterogeneous vehicular wireless networking as an optimization problem, and we analyze the impact of bandwidth aggregation on the problem formulation as well as the computational complexity of the problem. Our analysis shows that the problem is tractable and can be solved in polynomial time if bandwidth aggregation is employed, but the problem becomes NP-complete and cannot be solved optimally in general if bandwidth aggregation is not employed. On the other hand, the use of bandwidth aggregation requires deploying a proxy server between vehicles and content servers, thus introducing a different operation model than the case when bandwidth aggregation is not used. Besides understanding the impact of bandwidth aggregation on the tractability of the problem, our analysis suggests efficient approaches to solving the problem in an optimal manner when bandwidth aggregation is employed, and our analysis suggests both optimal and approximate approaches to solving the problem when bandwidth aggregation is not employed. To the best of our knowledge, this is the first work investigating the theoretical aspect of heterogeneous vehicular wireless networking as well as the associated tradeoffs between the ease of deployment, computational tractability, and optimality.

The rest of this article is organized as follows. We introduce the concept of bandwidth aggregation in Section II, then we study heterogeneous vehicular wireless networking with and without bandwidth aggregation in Sections III and IV respectively. We discuss related work in Section V, and we make concluding remarks in Section VI.

II. Bandwidth aggregation

When a vehicle is connected to the Internet via multiple wireless networks, an application can simultaneously use the multiple wireless networks to leverage the aggregate

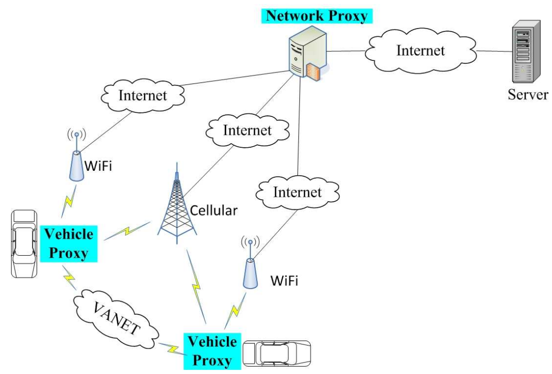

Fig. 1. System architecture: with bandwidth aggregation

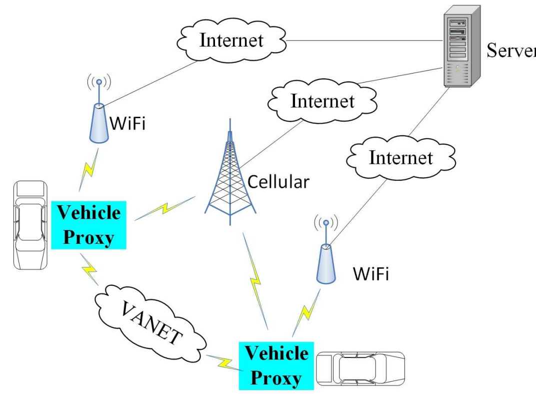

Fig. 2. System architecture: without bandwidth aggregation

bandwidths of all the wireless networks. For convenience, we regard this as bandwidth aggregation [3], and Figure 1 shows the system architecture for heterogeneous vehicular wireless networking when bandwidth aggregation is used. In the architecture, each vehicle has a vehicle proxy that communicates with content servers via a network proxy. The vehicle proxy and network proxy execute the bandwidth aggregation [3] algorithm to make the bandwidth of all the wireless networks available to the application. The vehicle proxy and the network proxy provide a unified, standard transport-layer socket interface (e.g., TCP interface) to the in-vehicle applications and content servers respectively so that the operation of bandwidth aggregation is transparent to them and the multiple wireless networks are represented as a single communication pipe/socket between the in-vehicle applications and content servers; this way, the in-vehicle applications and content servers can be developed without being aware of the underlying heterogeneity of wireless communication, and legacy applications and content servers can also be supported without any revision.

As we will present in Section III, bandwidth aggregation makes the optimal allocation of heterogeneous wireless network resources computationally tractable. On the other hand, bandwidth aggregation requires deploying not only a vehicle proxy at each vehicle but also network proxies in the network, and the requirement for deploying network proxies may slow down the adoption of heterogeneous networking since network proxies introduce additional cost that need to be addressed by innovative business models. An alternative approach to bandwidth aggregation is for each in-vehicle application to use one and only one of the wireless networks available. This way, we only need a vehicle proxy at the vehicle to assign applications to the individual wireless networks, and we do not need network proxies any more. Accordingly, the system architecture is as shown in Figure 2. Without network proxies, it is relatively easier to deploy the system shown in Figure 2 since only vehicle proxies need to be installed, as software packages, in vehicles by individual vehicle owners; on the other hand, the problem of optimal heterogeneous vehicular wireless networking becomes computationally intractable in general as we will show in Section IV. Therefore, there are tradeoffs between the ease of deployment, computational tractability, and optimality in different architectures of heterogeneous vehicular wireless networking, and a major objective of this study is to elucidate such tradeoffs.

III. Heterogeneous vehicular wireless networking with bandwidth aggregation

In what follows, we first mathematically formulate the problem of heterogeneous vehicular wireless networking with bandwidth aggregation, then we propose efficient approaches to solving the problem in an optimal manner.

A. Problem formulation

Given a set of applications to be supported by the multiple wireless networks available to a vehicle, the problem is to allocate the network bandwidth to the individual applications such that the overall application utility is maximized subject to the total available bandwidth. Given that the cost of using different wireless networks are different, we also consider the objective of minimizing cost.

Given a vehicle, each available wireless network provides a communication pipe connecting the vehicle to the Internet. If there are P wireless networks available, there are P communication pipes each of which has its own capacity and cost. For mathematically formulating the problem of optimal heterogeneous vehicular wireless networking, we use the notation shown in Table I. As used in most studies and

A:	total number of applications;
d_i:	bandwidth requirement (in Mbps) of the i-th application, $i = 1 \ldots A$;
$U_i(.)$:	the utility function of the i-th application;
P:	total number of pipes;
c_j:	capacity of the j-th pipe in Mbps, $j = 1 \ldots P$;
f_j:	cost of per-unit capacity of the j-th pipe, $j = 1 \ldots P$;
r_{ij}:	the amount of traffic (in Mbps) from the i-th application that are carried by the j-th pipe, $i = 1 \ldots A$, $j = 1 \ldots P$.

TABLE I
Notation

also reflecting the decreasing utility margin as the provided bandwidth to an application increases, the utility function $U_i(.)$ is assumed to be concave [4].

Then the problem can be modeled as the following optimization problem:

$$\begin{aligned}
\underset{r_{ij}}{\text{maximize}} \quad & \sum_{i=1}^{A} U_i(\sum_{j=1}^{P} r_{ij}) - \epsilon \sum_{j=1}^{P}(f_j \sum_{i=1}^{A} r_{ij}) \\
\text{subject to} \quad & \sum_{j=1}^{P} r_{ij} \leq d_i, i = 1 \ldots A \\
& \sum_{i=1}^{A} r_{ij} \leq c_j, j = 1 \ldots P \\
& r_{ij} \geq 0, i = 1 \ldots A, j = 1 \ldots P
\end{aligned} \tag{1}$$

where ϵ is a very small number such that the optimal solution maximizes application utility while minimizing cost as a secondary objective. In the formulation, the first constraint ensures that the total bandwidth allocated to the i-th application is no more than its demand, the second constraint ensures that the total amount of traffic sent to the j-th pipe is no more than its capacity, and the third constraint ensures the validity of the decision variables r_{ij}. Note that the utility functions $U_i(.)$'s can be chosen to prioritize certain applications by making their utility functions greater than others'.

Problem (1) can be extended to model, as follows, the case where there exists an upper bound F_0 on the total cost per second that a user can afford.

$$\begin{aligned}
\underset{r_{ij}}{\text{maximize}} \quad & \sum_{i=1}^{A} U_i(\sum_{j=1}^{P} r_{ij}) - \epsilon \sum_{j=1}^{P}(f_j \sum_{i=1}^{A} r_{ij}) \\
\text{subject to} \quad & \sum_{j=1}^{P} r_{ij} \leq d_i, i = 1 \ldots A \\
& \sum_{i=1}^{A} r_{ij} \leq c_j, j = 1 \ldots P \\
& \sum_{j=1}^{P}(f_j \sum_{i=1}^{A} r_{ij}) \leq F_0 \\
& r_{ij} \geq 0, i = 1 \ldots A, j = 1 \ldots P
\end{aligned} \tag{2}$$

where the third constraint ensures that the total cost is no more than the maximum affordable. Complexity analysis and algorithms for Problem (1) are readily extendable to Problem (2). For conciseness of presentation, therefore, the rest of this Section focuses on Problem (1) alone.

B. Optimal solutions

We solve Problem (1) using different approaches depending on the operation regimes of the system as follows.

Overprovisioning. When the total bandwidth required by the applications are no more than the total bandwidth provided by all the communication pipes, $\sum_{j=1}^{P} r_{ij} = d_i \; (i = 1 \ldots A)$ holds in the optimal solution to Problem (1), and bandwidth aggregation makes Problem (1) the problem of choosing the least-cost set of pipes that can satisfy the application bandwidth requirements. In this case, Problem (1) becomes the following problem:

$$\begin{aligned}
\underset{r_{ij}}{\text{minimize}} \quad & \sum_{j=1}^{P}(f_j \sum_{i=1}^{A} r_{ij}) \\
\text{subject to} \quad & \sum_{j=1}^{P} r_{ij} = d_i, i = 1 \ldots A \\
& \sum_{i=1}^{A} r_{ij} \leq c_j, j = 1 \ldots P \\
& r_{ij} \geq 0, i = 1 \ldots A, j = 1 \ldots P
\end{aligned} \tag{3}$$

This problem can be solved in an optimal manner using Algorithm 1.

Algorithm 1 Heterogeneous vehicular wireless networking with bandwidth aggregation and overprovisioning

Input: a set $\mathbb{A}$ of applications, with the i-th application $\mathcal{A}_i$ having bandwidth requirement d_i; a set $\mathbb{P}$ of communication pipes, with the j-th pipe $\mathcal{P}_j$ having capacity c_j and per-unit-capacity cost f_j.

Output: the set of pipes selected to carry application traffic, and the application traffic assignment specifying the amount of traffic, denoted by r_{ij}, from the i-th application that are carried by the j-th pipe ($i = 1 \ldots A, j = 1 \ldots P$).

/* Selecting the least-cost set of pipes satisfying the application bandwidth requirements */

$\mathbb{S} = \emptyset$;

Order the elements of $\mathbb{P}$ in the non-decreasing order of their per-unit-capacity cost such that $f_k \leq f_{k'}$ if $k < k'$;

$k = 1$; $\mathbb{S} = \{\mathcal{P}_1\}$;

while $\sum_{j=1}^{k} c_j < \sum_{i=1}^{|\mathbb{A}|} d_i$ **do**

  k = k+1;

  $\mathbb{S} = \mathbb{S} \cup \{\mathcal{P}_k\}$;

end while

/* Assigning application traffic to the selected pipes */

if $|\mathbb{S}| = 1$ **then**

  Assign all application traffic to pipe $\mathcal{P}_1$; generate the corresponding r_{ij}'s;

else

  Arbitrarily assign $\sum_{j=1}^{|\mathbb{S}|-1} c_j$ amount of application traffic to pipes $\mathcal{P}_1, \ldots, \mathcal{P}_{|\mathbb{S}|-1}$ so that each pipe's bandwidth is fully utilized; generate the corresponding r_{ij}'s;

  Assign the remaining $\sum_{i=1}^{|\mathbb{A}|} d_i - \sum_{j=1}^{|\mathbb{S}|-1} c_j$ amount of application traffic to pipe $\mathcal{P}_{|\mathbb{S}|}$; generate the corresponding r_{ij}'s;

end if

Return $\mathbb{S}$ and r_{ij} $(i = 1 \ldots A, j = 1 \ldots P)$.

In the algorithm, pipes are added to the selected set $\mathbb{S}$ in the non-decreasing order of their per-unit-capacity cost until the overall bandwidth provided by the pipes of $\mathbb{S}$ is no less than the total bandwidth requirements of all the applications.

Then the set $\mathbb{S}$ of the selected pipes are filled with application traffic in the non-decreasing order of their per-unit-capacity cost, with pipes of lower per-unit-capacity cost filled to the full before filling (potentially in a partial manner) the pipe of the largest per-unit-capacity cost.

It is easy to see that the time complexity of Algorithm 1 is $O(P+A)$. For the optimality of Algorithm 1, we have

Theorem 1: Algorithm 1 solves Problem (3) in an optimal manner.

Proof: We prove the theorem through contradiction.

Assuming that Algorithm 1 did not solve Problem (3) in an optimal manner, then one or both of the following cases would hold for an optimal solution:

- There exists a pipe $\mathcal{P}_{k'} \notin \mathbb{S}$ that should be included in the set of the selected pipes for carrying application traffic and the per-unit-capacity cost of pipe $\mathcal{P}_{k'}$ is different from those of the selected pipes in $\mathbb{S}$.
- $|\mathbb{S}| > 1$, and the bandwidth of a pipe $\mathcal{P}_j (j \in \{1, \ldots, |\mathbb{S}| - 1\})$ is not fully utilized, where the per-unit-capacity cost of $\mathcal{P}_j$ is different from that of pipe $\mathcal{P}_{|\mathbb{S}|}$.

In the first case, since the per-unit-capacity costs of the pipes in $\mathbb{S}$ are the lowest among all the pipes in $\mathbb{P}$, the per-unit-capacity cost of pipe $\mathcal{P}_{k'}$ is higher than those of the selected pipes in $\mathbb{S}$. Thus the cost of carrying the application traffic carried by pipe $\mathcal{P}_{k'}$ can be reduced using a pipe in $\mathbb{S}$, and any solution using a pipe whose per-unit-capacity cost is greater than those of the pipes in $\mathbb{S}$ leads to higher cost than the solution generated by Algorithm 1. Thus the first case does not exist in practice.

In the second case, the cost of vehicular Internet access can be reduced by moving traffic carried by pipe $\mathcal{P}_{|\mathbb{S}|}$ to fully fill pipe $\mathcal{P}_j$; this is because the per-unit-capacity cost of $\mathcal{P}_{|\mathbb{S}|}$ is greater than $\mathcal{P}_j$. Thus, the second case does not exist in practice either.

■

Underprovisioning. When the total bandwidth required by the applications exceeds the total bandwidth provided by all the communication pipes, $\sum_{i=1}^{A} r_{ij} = c_j \ (j = 1 \ldots P)$ holds in the optimal solution to Problem (1), and Problem (1) becomes the problem of controlling the traffic load of individual applications so that the total application-level utility is maximized. In this case, Problem (1) becomes the following problem:

$$\begin{aligned} \underset{r_{ij}}{\text{maximize}} \quad & \sum_{i=1}^{A} U_i(\sum_{j=1}^{P} r_{ij}) \\ \text{subject to} \quad & \sum_{j=1}^{P} r_{ij} \le d_i, i = 1 \ldots A \\ & \sum_{i=1}^{A} r_{ij} = c_j, j = 1 \ldots P \\ & r_{ij} \ge 0, i = 1 \ldots A, j = 1 \ldots P \end{aligned} \tag{4}$$

Problem (4) is a convex optimization problem and can be solved in $O((AP)^3)$ time using the interior point method [4], [5].

IV. Heterogeneous vehicular wireless networking without bandwidth aggregation

In what follows, we first mathematically formulate the problem of heterogeneous vehicular wireless networking without bandwidth aggregation, then we prove the NP-completeness of the problem and suggests potential approaches to solving the problem in an approximate manner.

A. Problem formulation

Without bandwidth aggregation, each application on a vehicle can use at most one and only one of the available wireless networks. To reflect this constraint, we introduce the indicator variable d_{ij} to denote whether the i-th application uses the j-th pipe ($i = 1 \ldots A, j = 1 \ldots P$). We reuse the notations presented in Table I. Then the problem of heterogeneous vehicular wireless networking without bandwidth aggregation can be modeled as the following optimization problem:

$$\begin{aligned} \underset{d_{ij}, r_{ij}}{\text{maximize}} \quad & \sum_{i=1}^{A} U_i(\sum_{j=1}^{P} d_{ij} r_{ij}) - \epsilon \sum_{j=1}^{P} (f_j \sum_{i=1}^{A} d_{ij} r_{ij}) \\ \text{subject to} \quad & \sum_{j=1}^{P} d_{ij} \le 1, i = 1 \ldots A \\ & r_{ij} \le d_{ij} c_j, i = 1 \ldots A, j = 1 \ldots P \\ & \sum_{j=1}^{P} r_{ij} \le d_i, i = 1 \ldots A \\ & \sum_{i=1}^{A} d_{ij} r_{ij} \le c_j, j = 1 \ldots P \\ & d_{ij} \in \{0, 1\}, r_{ij} \ge 0, i = 1 \ldots A, \\ & \qquad j = 1 \ldots P \end{aligned} \tag{5}$$

where ϵ is a very small number such that the optimal solution maximizes application utility while minimizing cost as a secondary objective. In the formulation, the first constraint ensures that an application uses no more than one pipe. The second constraint relates r_{ij} to d_{ij} such that they are consistent and that r_{ij} is 0 if d_{ij} is 0. The other constraints are similar to the ones on Problem (1) for the case when bandwidth aggregation is employed.

With the constraint of $r_{ij} \le d_{ij} c_j$, Problem (5) can be reformulated in a simpler form as follows where all the nonlinear terms $d_{ij} r_{ij}$ are replaced with the linear terms r_{ij} so that the formulation is more amenable to the complexity analysis and algorithm discussion in Section IV-B:

$$
\begin{aligned}
\underset{d_{ij}, r_{ij}}{\text{maximize}} \quad & \sum_{i=1}^{A} U_i(\sum_{j=1}^{P} r_{ij}) - \epsilon \sum_{j=1}^{P} (f_j \sum_{i=1}^{A} r_{ij}) \\
\text{subject to} \quad & \sum_{j=1}^{P} d_{ij} \leq 1, i = 1 \ldots A \\
& r_{ij} \leq d_{ij} c_j, i = 1 \ldots A, j = 1 \ldots P \\
& \sum_{j=1}^{P} r_{ij} \leq d_i, i = 1 \ldots A \\
& \sum_{i=1}^{A} r_{ij} \leq c_j, j = 1 \ldots P \\
& d_{ij} \in \{0,1\}, r_{ij} \geq 0, i = 1 \ldots A, \\
& \qquad\qquad\qquad\qquad j = 1 \ldots P
\end{aligned} \tag{6}
$$

Problem 6 can be extended to model the case where there exists an upper bound F_0 on the total cost per second that a user can afford, so can the complexity analysis and algorithm discussion in Section IV-B; for conciseness of presentation, however, here we only focus on Problem (6).

B. NP-completeness

First, we analyze the computational complexity of solving Problem (6) as follows:

Theorem 2: It is NP-complete to solve the pipe assignment problem (6).

Proof: Let's consider a special case of (6) where

$$
U_i(\sum_{j=1}^{P} r_{ij}) = \begin{cases} U_i & \text{if } \sum_{j=1}^{P} r_{ij} \geq d_i \\ 0 & \text{otherwise} \end{cases}
$$

That is, the bandwidth requirement is stringent such that the utility of an application is zero if its bandwidth requirement is not satisfied. This case applies to mission-critical applications such as active vehicle safety control through vehicle-to-infrastructure communication; this case may well apply to vehicle infotainment applications too since not satisfying application requirements causes undesirable application performance (e.g., garbled online music) which annoys/distracts drivers' attention and indirectly causes driving safety concerns.

For this special case, the optimal solution should be such that, if an application i is assigned to a data pipe j, then the pipe j should allocate d_i of its communication capacity to application i, i.e., $r_{ij} = d_i$. Therefore, the optimal solution to Problem (6) is the same as that to the following multiple knapsack problem:

$$
\begin{aligned}
\underset{d_{ij}, r_{ij}}{\text{maximize}} \quad & \sum_{i=1}^{A} U_i \sum_{j=1}^{P} d_{ij} - \epsilon \sum_{j=1}^{P} (f_j \sum_{i=1}^{A} r_{ij}) \\
\text{subject to} \quad & \sum_{j=1}^{P} d_{ij} \leq 1, i = 1 \ldots A \\
& r_{ij} = d_i d_{ij}, i = 1 \ldots A, j = 1 \ldots P \\
& \sum_{i=1}^{A} r_{ij} \leq c_j, j = 1 \ldots P \\
& d_{ij} \in \{0,1\}, r_{ij} \geq 0, i = 1 \ldots A, \\
& \qquad\qquad\qquad\qquad j = 1 \ldots P
\end{aligned} \tag{7}
$$

Since the multiple knapsack problem (7) is NP-complete [6], the general problem (6) is NP-complete.

■

Since Problem (6) is NP-complete and it is a nonlinear mixed-integer programming problem, one approach to solving it is the branch-and-bound method [7] if optimality is desired. In the standard branch-and-bound method framework [7], we can explore the feasible region by examining sub-regions as defined by the integer decision variables d_{ij} ($i = 1 \ldots A, j = 1 \ldots P$), and, given a sub-region defined by a specific configuration of the variables d_{ij}, Problem (6) becomes a convex optimization problem and can be solved through efficient methods such as the interior point method [5].

When we can tradeoff optimality for shorter execution time, we can use approximation algorithms. For instance, Problem (6) is a type of nonlinear generalized assignment problem, and a greedy algorithm similar to that for nonlinear generalized assignment problems [8] can be used to solve Problem (6). Detailed study and implementation of such algorithms is an interesting topic worth pursuing in future work, but it is beyond the scope of this paper.

V. Related work

For improved bandwidth and reliability in data communication, both transport and network layer approaches have been proposed for enabling an application to leverage the communication capability from multiple networks [3], [9], [10], [11], [12]. Those work focused on the systems architecture and algorithmic issues for bandwidth aggregation, but they did not study the problem of how to optimally allocate bandwidth of heterogeneous wireless networks to different applications, nor did they examine the impact of bandwidth aggregation on the computational complexity of the problem. Koga et al. [13] and Higgins et al. [14] have studied the system architectures and programming abstractions for assigning the traffic of multiple applications to different networks. Focusing on systems architecture and implementation, however, those work did not investigate the optimization aspect of the problem, nor did they study the impact of bandwidth aggregation on the computational complexity of the optimization problem.

Handoff mechanisms have been extensively studied for vehicular wireless networks [15], [16] and cellular networks [17].

Those work mostly focused on how to enable an application to smoothly switch from one network to another network, and they did not study the problem of optimally allocating the traffic of multiple applications to multiple concurrent networks.

Many work have studied the problem of simultaneously leveraging different types of wireless networks (e.g., cellular networks and WiFi networks) for improved data dissemination, for instance, by optimally design the network topology of different types of networks and by optimally deciding which network to use during different phases of data communication [18], [19], [20], [21], [22]. Bhargava et al. [23] also examined the approach of using cellular networks for communicating control information (e.g., those for routing and localization) needed for effective data communication in mobile ad hoc networks. Those work did not consider the problem of optimally allocating the bandwidth of concurrently-operating wireless networks to multiple applications, nor did they study the impact of bandwidth aggregation on the computational complexity of the optimization.

VI. Concluding remarks

Towards a foundation for heterogeneous vehicular wireless networking, we have proposed the mathematical formulation of heterogeneous vehicular wireless networking, and we have analyzed the impact of bandwidth aggregation on the problem formulation as well as the computational complexity of the problem. Our analysis has shown that, with bandwidth aggregation, the problem can be solved in polynomial time and, without bandwidth aggregation, the problem becomes NP-complete and cannot be solved optimally in general. Our analysis is constructive such that it has suggested efficient approaches to solving the problem in an optimal manner when bandwidth aggregation is employed, and it has suggested both optimal and approximate approaches to solving the problem when bandwidth aggregation is not employed. The above analytical insight serves as a guidance on choosing the system architectures and algorithms for heterogeneous vehicular wireless networking, and a future direction is to implement and experimentally evaluate the different architectures and algorithms in real-world settings and understand the associated tradeoffs between the ease of deployment, computational tractability, and optimality. Our analysis in this paper focuses on the impact of throughput on application utility; an interesting future direction worth pursuing is to consider applications' delay requirements too, which are important for applications such as real-time multimedia streaming and networked vehicle control.

Acknowledgment

We thank Dan Zhang for his feedback to the formulation of our optimization problem. Zhang's effort in this work is supported in part by a grant from GM Research and by NSF awards CNS-1136007, CNS-1054634, GENI-1890, and GENI-1633.

References

[1] N. Navet, Y. Song, F. Simonot-Lion, and C. Wilwert, "Trends in automotive communication systems," *Proceedings of the IEEE*, vol. 93, no. 6, 2005.

[2] A. Vinel, "3GPP LTE versus IEEE 802.11p/WAVE: Which technology is able to support cooperative vehicular safety applications?" *IEEE Wireless Communications Letters*, vol. 1, no. 2, 2012.

[3] K. Chebrolu and R. R. Rao, "Bandwidth aggregation for real-time applications in heterogeneous wireless networks," *IEEE Transactions on Mobile Computing*, vol. 5, no. 4, 2006.

[4] R. Srikant, , and L. Ying, *Communication Networks: An Optimization, Control and Stochastic Networks Perspective*. Cambridge University Press, 2014.

[5] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.

[6] C. Chekuri and S. Khanna, "A polynomial time approximation scheme for the multiple knapsack problem," *SIAM Journal on Computing*, vol. 35, no. 3, 2003.

[7] M. Pioro and D. Medhi, *Routing, Flow, and Capacity Design in Communication and Computer Networks*. Morgan Kaufmann, 2004.

[8] T. C. Sharkey and H. E. Romeijn, "Greedy approaches for a class of nonlinear generalized assignment problems," *Discrete Applied Mathematics*, vol. 158, pp. 559–572, 2010.

[9] K. Chebrolu, B. Raman, and R. R. Rao, "A network layer aproach to enable TCP over multiple interfaces," *Wireless Networks*, vol. 11, pp. 637–650, 2005.

[10] H.-Y. Hsieh and R. Sivakumar, "A transport layer approach for achieving aggregate bandwiths on multi-homed mobile hosts," in *ACM MobiCom*, 2002.

[11] L. Magalhaes and R. Kravets, "Transport level mechanisms for bandwidth aggregation on mobile hosts," in *IEEE ICNP*, 2001.

[12] D. S. Phatak and T. Goff, "A novel mechanism for data streaming across multiple ip links for improving throughput and reliability in mobile environments," in *IEEE INFOCOM*, 2002.

[13] H. Koga, S. Kashihara, Y. Fukuda, K. Lida, and Y. Oie, "A quality-aware VoWLAN architecture and its quantitative evaluations," *IEEE Wireless Communications*, February 2006.

[14] B. D. Higgins, A. Reda, T. Alperovich, J. Flinn, T. Giuli, B. Noble, and D. Watson, "Intentional networking: Opportunistic exploitation of mobile network diversity," in *ACM MobiCom*, 2010.

[15] K. Shafiee, A. Attar, and V. C. M. Leung, "Optimal distributed vertical handoff strategies in vehicular heterogeneous networks," *IEEE JSAC*, vol. 29, no. 3, pp. 534–544, 2011.

[16] J.-M. Chung, M. Kim, Y.-S. Park, M. Choi, S. Lee, and H. S. Oh, "Time coordinated V2I communications and handover for WAVE networks," *IEEE JSAC*, vol. 29, no. 3, pp. 545–558, 2011.

[17] C. Sun, E. Stevens-Navarro, V. Shah-Mansouri, and V. W. Wong, "A constrained MDP-based vertical handoff decision algorithm for 4G heterogeneous wireless networks," *Wireless Networks*, vol. 17, pp. 1063–1081, 2011.

[18] J. Ahn, M. Sathiamoorthy, B. Krishnamachari, F. Bai, and L. Zhang:, "Optimizing content dissemination in vehicular networks with radio heterogeneity," *IEEE Transactions on Mobile Computing*, vol. 13, no. 6, pp. 1312–1325, 2014.

[19] J. Ahn, Y. Wang, B. Yu, F. Bai, and B. Krishnamachari, "RISA: Distributed road information sharing architecture," in *IEEE INFOCOM*, 2012.

[20] A. Balasubramanian, B. N. Levine, and A. Venkataramani, "Enhancing interactive web applications in hybrid networks," in *ACM MobiCom*, 2008.

[21] L. K. Law, K. Pelechrinis, S. V. Krishnamurthy, and M. Faloutsos, "Downlink capacity of hybrid cellular ad hoc networks," *IEEE/ACM Transactions on Networking*, vol. 18, no. 1, pp. 243–256, 2010.

[22] H. Wu, C. Qiao, S. De, and O. Tonguz, "Integrated cellular and ad hoc relaying systems: iCAR," *IEEE JSAC*, vol. 19, no. 10, pp. 2105–2115, 2001.

[23] B. Bhargava, X. Wu, Y. Lu, and W. Wang, "Integrating heterogeneous wireless technologies: A cellular aided mobile ad hoc network (CAMA)," *Mobile Networks and Applications*, vol. 9, pp. 393–408, 2004.