

Impact of Information Availability on Starvation Mitigation and Delay Minimal Delivery in ICRCNs

Ribal Atallah and Wissam Fawaz

Abstract—This paper looks into an Intermittently Connected Roadside Communication Network (ICRCN) scenario comprising two isolated source Stationary Roadside Units (SRUs) relying on mobile smart vehicles to relay data to a destination SRU. In this case, it was shown in [1] that the downstream source SRU may suffer from a significant starvation problem. As such, a Markov decision process framework was established therein to identify a suitable Bulk Release Decision Policy (BRDP). BRDP was then implemented within a Starvation Mitigation and Delay-Minimal (SMDM) delivery scheme.

In this paper, we investigate the impact of the level of information availability at the upstream non-starving node on the performance of the SMDM scheme. In particular, extensive simulations are conducted for the purpose of quantifying the ability of SMDM to jointly mitigate starvation and achieve minimal end-to-end bundle delivery delay under conditions of perfect, imperfect, and no information availability at the non-starving node.

Index Terms—ICRCN, SRU, Performance Evaluation, MDP.

I. INTRODUCTION

THE conception of ICRCNs consists of transforming mobile vehicles into smart and communication-enabled entities that are able to establish connectivity among isolated Stationary Roadside Units (SRUs). ICRCNs emerged as such as a terrestrial application of the Disruption-Tolerant Networking paradigm, [2]. Precisely, this conception consists of augmenting vehicles, particularly those restricted to navigable roadways, with computerized modules, finite buffers and wireless communication devices to inter-communicate as well as to communicate with the SRUs deployed along the roadways. Under such conditions, ICRCNs represent an effective and cost-minimal solution for bringing digital connectivity closer to isolated areas, especially when the cost of setting up such a networking infrastructure is deemed elevated [3].

This paper revolves around the networking scenario depicted in Figure 1, which involves three SRUs, namely two source SRUs referred to as S_1 and S_2 along with a destination SRU, D . All three SRUs are installed along a one-dimensional, uninterrupted highway segment. The three SRUs cannot directly communicate with one another as each one is located outside the respective coverage ranges of the two others. The two source SRUs S_1 and S_2 are assumed to be completely isolated, while D is privileged by a connection to the Internet. In the absence of a networking infrastructure connecting S_1 and S_2 to D , mobile vehicles act as store-carry-forward data carriers from both of S_1 and S_2 to D .

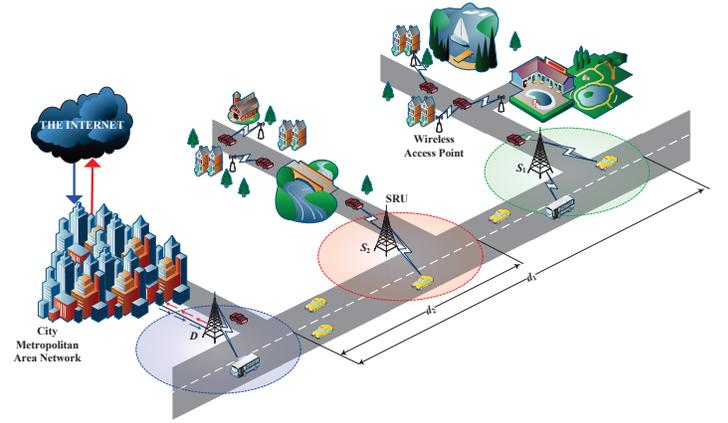


Fig. 1. ICRCN scenario.

Upon the occurrence of a vehicle arrival to either one of the two source SRUs S_1 or S_2 , the vehicle presents itself to the SRU as a bundle¹ release opportunity. In a point-to-point ICRCN with a single source SRU, the work of [3] assessed the contribution of such an opportunity to the minimization of the average end-to-end delivery delay of a singly released bundle. The authors of [4] addressed the limitations of [3] and demonstrated how the release of a bulk of bundles² from a source SRU's buffer significantly improves the performance in terms of the average end-to-end delivery delay.

While these studies focused only on ICRCN scenarios consisting of a single source SRU, the work of [1] opted for the investigation of an ICRCN scenario where several source SRUs may be present along the way to the destination SRU. In addition, unlike the previous studies, the authors in [1] considered the case in which vehicle buffers are finite and thus unable to accommodate all the possibly emanating bundles from each one of the encountered source SRUs. Under such circumstances, it becomes highly likely that the source SRUs that are located further downstream be subject to severe starvation especially if the upstream SRU starts downloading too much data to passing by vehicles and thereby hogs the finite buffers offered by these vehicles.

More precisely, considering the highway segment illustrated in Figure 1, vehicles navigating towards D will first pass by S_1 before encountering S_2 . Both of S_1 and S_2 will try to load the arriving vehicles' finite buffers with as many as possible of their respective data bundles. Such a bundle release model

R. Atallah and W. Fawaz are with the ECE department of the Lebanese American University. E-Mail: r.atallah@gmail.com, wissam.fawaz@lau.edu.lb 978-1-4799-3060-9/14/\$31.00 ©2014 IEEE

¹Data and control signals are combined in a single atomic entity, called bundle, that is transmitted across a DTN-based ICRCN, [2].

²A group of bundles is referred to as a bulk of bundles.

eventually favors S_1 over S_2 . This is particularly true since by the time a vehicle has reached S_2 , its buffer might have been considerably loaded (ultimately fully exhausted) with S_1 's bundles. A node such as S_2 is typically called a *starving node*, since bundles would rapidly accumulate in S_2 's buffer. As a result, incoming bundles to S_2 will be subject to excessive queueing delays.

The study in [1] proposed to mitigate the severity of the starvation that a downstream SRU may suffer from through the deployment of the so-called Starvation Mitigation and Delay-Minimal (SMDM) delivery scheme. In the context of SMDM, the authors assumed that S_1 possesses real-time full knowledge (i.e., perfect information) about the number of bundles queueing at S_2 ³. Through the availability of perfect information at an upstream node, SMDM was shown to be able to resolve the starvation problem while at the same time achieving minimal end-to-end bundle delivery delay.

We refer in the sequel to the variant of SMDM where S_1 is equipped with perfect information about S_2 's queue length as SMDM-PI. In practice, it is impossible for S_1 to acquire such perfect information in a real-time manner. This limits the SMDM-PI scheme in terms of practicality and imprisons its ability to achieve optimal performance results within the frames of a purely theoretical study. Inspired by this observation, the present paper investigates a new variant of the SMDM scheme dubbed SMDM-II that differs from SMDM-PI in that imperfect information is assumed to be present at S_1 with respect to the status of S_2 's buffer. A simulation study is laid out in order to evaluate the impact of the existence of imperfect information at S_1 on the performance of the SMDM scheme.

The rest of this manuscript is organized as follows. In section II, the SMDM scheme proposed in [1] is described. Section III introduces the SMDM variant proposed in this paper. In section IV, extensive simulations are performed to analyze the performance of the bundle delivery schemes under study. Finally, section V concludes the paper.

II. STARVATION MITIGATION DELAY MINIMAL BUNDLE DELIVERY

A. Background

In the context of the ICRCN scenario illustrated in Figure 1, S_1 and S_2 are considered to be identical. Their respective coverage ranges are of length $d_C(meters)$ each. The distances that separate S_1 and S_2 from D are $d_1(meters)$ and $d_2(meters)$ respectively. Following the guidelines given in [1], [3], [4], the following assumptions are made:

- A1: The bundle arrivals follow a Poisson process with a rate of λ_b ($\frac{bundles}{second}$).
- A2: S_1 and S_2 are subject to equal bundle arrival rates.
- A3: The bundle sizes are fixed to b (*bytes*).
- A4: S_1 and S_2 have equal data transmission rates of T_R ($\frac{Mbits}{second}$).
- A5: S_1 and S_2 are equipped with infinite buffers.

³Thereafter, the number of bundles stored in S_2 's queue will be referred to as S_2 's queue length.

- A6: The vehicle arrivals follow a Poisson process with a rate of μ_v ($\frac{vehicles}{second}$).
- A7: The per-vehicle speed V_k is drawn from a uniform distribution defined over the range $[V_{min}; V_{max}]$ ($\frac{meters}{second}$).
- A8: The per-vehicle capacity C_k is drawn from a uniform distribution defined over the range $[C_{min}; C_{max}]$.
- A9: Each vehicle maintains a constant speed during its navigation over the highway segment $[S_1 D]$.
- A10: The IEEE 802.11p protocol is used for vehicle-to-SRU communication.

B. Theoretical foundation

On the occurrence of a vehicle arrival, either one of S_1 or S_2 has to decide on whether or not to use the arriving vehicle as a bulk transporter to the destination SRU D . This release decision process evolves over the sequence of vehicles $\{0, 1, 2, \dots, k, \dots\}$ with speeds $\{v_0, v_1, v_2, \dots, v_k, \dots\}$ arriving at times $\{t_0, t_1, t_2, \dots, t_k, \dots\}$. Based on assumption (A6), $\{I_k, k \geq 1\}$, the inter-arrival time intervals, form a sequence of independent and identical exponentially distributed values with an average $E[I] = \frac{1}{\mu_v}$. Let $N_{i,k}$ denote the number of bundles waiting in the i^{th} source SRU's buffer ($i = 1, 2$) at time t_k . Also let $Y_{i,k}$ represent the decision made by the i^{th} SRU ($i = 1, 2$) at time t_k and $X_{i,k}$ represent the size of the bulk released by the i^{th} source SRU at that time. In the context of the ICRCN scenario given in Figure 1, $Y_{i,k} \in \{0; 1\}$. It is important to highlight in this regard that whenever $N_{i,k} = 0$, $Y_{i,k} = 0$ is the only possible decision to be made and thus no bundles are released. As such:

$$N_{i,k+1} = \max\{0, N_{i,k} - X_{i,k} \cdot Y_{i,k}\} + A_{i,k} \quad (1)$$

where $A_{i,k}$ is the number of newly incoming bundles during the time interval $I_{k+1} = t_{k+1} - t_k$. $A_{i,k}$ has an average of $\bar{A}_i = E[A_{i,k}] = \lambda_b E[I]$. In view of the above, the state of the middle source SRU S_2 can be described by the two-tuple $(N_{2,k}; V_k)$ and its dynamics are governed by equation (1).

Now, define:

$$C_2(\ddot{n}_k, v_k, \ddot{y}_k) = \frac{\ddot{n}_k}{A_2} + f_2(v_k, \ddot{y}_k) \quad (2)$$

$C_2(\ddot{n}_k, v_k, \ddot{y}_k)$ is the single stage cost function corresponding to S_2 's state ($N_{2,k} = \ddot{n}_k; V_k = v_k$) and release decision $Y_{2,k} = \ddot{y}_k$. It is composed of two terms, namely: a) $\frac{\ddot{n}_k}{A_2}$ being the cost associated with S_2 's queue length (or the queueing delay) and b) $f_2(v_k, \ddot{y}_k)$ corresponding to the cost incurred by the exploitation of the vehicle of speed v_k as a bundle carrier from S_2 to D (i.e., the transit delay). Note that $f_2(v_k, \ddot{y}_k) = \frac{d_2}{A_2} \cdot \frac{\ddot{y}_k}{v_k}$.

At this level, given that V_k is uniformly distributed over a finite range, the vehicle inter-arrival time is exponentially distributed, and the decision space $Y_{2,k}$ is finite, there exists a policy π_2^* that minimizes the following cost function at S_2 [3], [5]:

$$\bar{C}_2 = \frac{1}{A_2} \limsup_{k \rightarrow \infty} \left[\frac{1}{k} \cdot E \left[\sum_{i=0}^k \left(N_{2,i} + \frac{d_2}{V_i} \cdot Y_{2,i} \right) \right] \right] \quad (3)$$

The problem of identifying the policy π_2^* that minimizes (3) is a classical average cost Markov decision problem where at each time step, S_2 chooses an action, namely, a release/no-release decision, and incurs a per-stage cost similar to the one given by (2). Such problems can be solved via dynamic programming techniques like the ones discussed in [6].

Similarly to S_2 , a single stage cost function can be formulated for S_1 . However, in addition to S_1 's queue size and release decision, this cost function has to account for S_2 's queue size and release decision as well. It is, therefore, expressed as:

$$C_1(\dot{n}_k, \ddot{n}_k, v_k, \dot{y}_k, \ddot{y}_k) = \frac{\dot{n}_k}{A_1} + f_1(v_k, \dot{y}_k) + C_2(\ddot{n}_k, v_k, \ddot{y}_k) \quad (4)$$

Similar to (2), the term $\frac{\dot{n}_k}{A_1}$ in (4) is tied to the queue length of S_1 . Also, in the same spirit, $f_1(v_k, \dot{y}_k) = \frac{d_1}{A_1} \cdot \frac{\dot{y}_k}{v_k}$. Following similar arguments as above, there exists a deterministic Markov policy $\pi_1^* \equiv \pi_1^*(\dot{n}_k, \ddot{n}_k, v_k, \dot{y}_k, \ddot{y}_k)$ that minimizes the average cost function developed for S_1 and which is given by:

$$\overline{C}_1 = \frac{1}{A_1} \limsup_{k \rightarrow \infty} \left[\frac{1}{k} E \left[\sum_{i=0}^k N_{1,i} + \frac{d_1}{V_i} \cdot Y_{1,i} \right] \right] + \overline{C}_2 \quad (5)$$

In this way, the joint starvation mitigation/delay minimization problem at S_1 is formulated as a Markov Decision Process framework with an average cost criterion given by (5). Following the guidelines of [6], this problem can be resolved using a Dynamic Programming approach. The Bulk Release Decision Policy (BRDP) for the overall system of the two source SRUs is represented by $\pi^* \equiv (\pi_1^*, \pi_2^*)$. Next, BRDP is implemented within a Starvation Mitigation Delay-Minimal (SMDM) bulk bundle delivery scheme.

In the context of the SMDM delivery scheme, the bundle release decisions made by both source SRUs S_1 and S_2 are driven mainly by π^* . If a decision is made at time t_k to release a bulk of bundles to an arbitrary vehicle k with a free buffer capacity C_k , the number of bundles $X_{i,k}$ to be downloaded to that vehicle by the i^{th} SRU ($i = 1, 2$) is determined as follows. At S_1 , $X_{1,k}$ basically depends on: a) C_k being the amount of free buffer space that is available at the k^{th} vehicle as well as on b) $N_{2,k}$ being the number of bundles currently queueing at S_2 . More specifically, the number of bundles released by S_1 to a vehicle k selected by BRDP is governed by the following inequality:

$$X_{1,k} \leq \frac{N_{1,k}}{N_{1,k} + N_{2,k}} \cdot C_k \quad (6)$$

$N_{2,k}$ plays a major role in throttling S_1 to which, now, is allocated only a well defined portion of the free buffer space available at the vehicle chosen by BRDP. This, by itself, contributes to the reservation of some spare buffer space for S_2 . Consequently, S_2 will, in turn, beneficially exploit the more abundant remaining vehicle buffer capacity spared by S_1 and hence abandon the title of a *starving node*.

C. Inadequacy of Perfect Information Assumption

It is clear that as described above, SMDM was built around the unrealistic assumption that the upstream SRU S_1 is equipped with accurate instantaneous information about the queue length at the downstream node S_2 . This assumption of perfect information in turn allowed SMDM to realize effective starvation mitigation. Obviously, a more realistic SMDM variant would be one that is based on imperfect information. Nonetheless, the more information is made available at S_1 about S_2 's buffer status, the higher the accuracy of S_1 's knowledge will be. It is thus indispensable for any SMDM variant based on imperfect information to identify practical means for feeding S_1 with as much information as possible about the status of S_2 's buffer. This observation established the foundation for the development of a new SMDM variant under which S_1 would only have partial imperfect information and yet be able to contribute to a great extent to the mitigation of the starvation problem. The aforementioned SMDM variant that we refer to as SMDM-II is proposed next.

III. SMDM WITH IMPERFECT INFORMATION

Recall that the objective of this study is to introduce a new variant of the SMDM delivery scheme that operates under the premise of the availability of imperfect information at S_1 with respect to the status of S_2 's buffer. The majority of today's highways are characterized by the use of dual carriageways running in opposite directions for the purpose of accommodating traffic plying between any two geographical points. As a result, it becomes possible to create a feedback channel connecting S_2 to S_1 by taking advantage of the vehicles navigating from S_2 to S_1 and by having these vehicles transport non-real time information about S_2 's buffer status to S_1 .

More formally, under the proposed SMDM-II, S_1 would maintain a variable called *Estimated* N_2 that represents the best current estimate by S_1 of S_2 's buffer length. When a new vehicle carrying a new *Sample* N_2 ⁴ from S_2 arrives at S_1 , S_1 simply updates *Estimated* N_2 as follows: *Estimated* $N_2 = \text{Sample}N_2$. In addition to having an estimate of S_2 's buffer length, S_1 is also required to keep track of the time instant at which *Estimated* N_2 was last updated. As such, each time *Estimated* N_2 gets updated, S_1 sets a variable called t_{update} to be equal to the time instant at which the update took place.

On the other hand, *Estimated* N_2 is updated differently when a vehicle navigating from S_1 to S_2 gets to S_1 . Particularly, when a vehicle k navigating towards S_2 enters the communication range of S_1 at time t_k , S_1 updates *Estimated* N_2 according to the following formula:

$$\text{Estimated}N_2 = \text{Estimated}N_2 + \lambda_2 \cdot (t_k - t_{update}) \quad (7)$$

where $\lambda_2 \cdot (t_k - t_{update})$ represents the average number of bundles that accumulates in S_2 's buffer over the time interval $[t_{update}; t_k]$.

So, with SMDM-II, S_1 's imperfect information about the queue size at S_2 takes the form of *Estimated* N_2 . Hence,

⁴The length of the queue at S_2 as seen by the newly arriving vehicle.

when applying the SMDM strategy as described in section II, S_1 would employ $EstimatedN_2$ instead of the $N_{2,k}$ parameter introduced earlier wherever the latter appears in the formulation of SMDM.

IV. SIMULATION AND NUMERICAL ANALYSIS

An ad-hoc Java-based in-house discrete event simulator is developed to simulate the scenario illustrated in Figure 1. The main purpose is to evaluate the performance of the proposed SMDM-II scheme and compare it to that exhibited by a number of benchmark delivery schemes including the SMDM-PI strategy.

A. Simulation Parameters

The following values were borrowed from [1] for the input simulation parameters: *a*) $\lambda_b = 30$ (bundles/s), *b*) $b = 1500$ (bytes), *c*) $T_R = 1$ (Mbits/s), *d*) $\mu_v \in [0.004; 0.033]$ (vehicles/s), *e*) $V_{min} = 10$ (m/s) *f*) $V_{max} = 50$ (m/s), *g*) $C_{min} = 0$, *h*) $C_{max} = 30000$, *i*) $d_C = 200$ (m), *j*) $d_1 = 2000$ (m) and *k*) $d_2 = 1000$ (m).

It is also important to note that ρ , the overall load offered to the system made up of both S_1 and S_2 , is:

$$\rho = \frac{(\lambda_1 + \lambda_2)}{\mu_v \bar{C}} = \frac{2\lambda_b}{\mu_v \bar{C}} \quad (8)$$

where λ_i is the bundle arrival rate to the i^{th} SRU ($i = 1, 2$) and \bar{C} is the average vehicle spare storage capacity. It is finally important to highlight that the values of the input parameters were set to ensure that the resulting system of SRUs is stable with $0 < \rho \leq 1$.

B. Metrics and Benchmark Schemes

The following metrics are used as performance indicators:

- Average bundle queueing delay, denoted by $\overline{Q_D}$: Average amount of time a bundle spends in either one of the queues belonging to S_1 or S_2 .
- Average bundle transit delay, denoted by $\overline{T_D}$: Average amount of time required by a bundle to travel from either one of S_1 or S_2 to the destination D .
- Average bundle end-to-end delay, denoted by $\overline{E_D} = \overline{Q_D} + \overline{T_D}$

In each simulation run, 10^7 bundle arrivals per SRU are considered and each plotted value of the above-enumerated metrics is the average of several independent simulation runs to secure a very narrow 95% confidence interval.

The following bundle delivery schemes are used as benchmarks:

- SMDM with Perfect Information (SMDM-PI): under this scheme, S_1 is assumed to have perfect information about the status of S_2 's buffer length. Moreover, bulk of bundles are released to those vehicles that contribute the most to achieving the dual purpose of starvation mitigation and delay-minimal bundle delivery.
- The Greedy Bundle Release Scheme with Bulk Bundle Release (GBRS-BBR) proposed in [4]: under this strategy, S_1 and S_2 take identical actions by releasing bulk of bundles to every arriving vehicle regardless of the vehicle's speed.

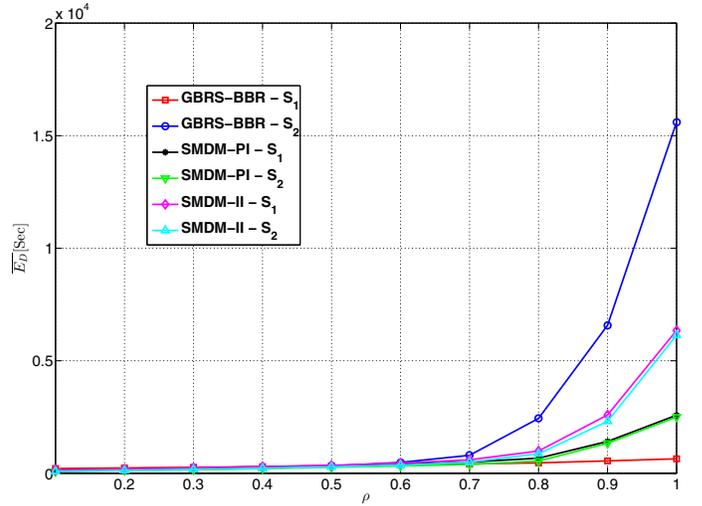


Fig. 2. Average end-to-end delay under the GBRs-BBR, SMDM-PI, and SMDM-II schemes.

C. Simulation Results

Figure 2 concurrently plots as a function of ρ the end-to-end delays at S_1 and S_2 as achieved by all of the GBRs-BBR, SMDM-PI, and SMDM-II schemes. This figure is a clear evidence of the amazing ability of the SMDM-PI and SMDM-II delivery schemes to resolve the starvation problem that the middle SRU S_2 suffers from under the GBRs-BBR scheme.

Indeed, when the GBRs-BBR scheme is deployed, the bundles of S_2 suffer from excessive end-to-end delays as clearly indicated by Figure 2. This phenomenon can be demystified by closely examining the main components that make up the end-to-end delay, namely the average bundle queueing delay denoted by $\overline{Q_D}$ and the average bundle transit delay denoted by $\overline{T_D}$. As a matter of fact, according to Figure 3, the bundles released by S_2 under GBRs-BBR are expected to experience lower transit delays as compared to their S_1 's counterpart. This can be justified by the fact that the bundles of S_2 travel a shorter distance to the destination D as opposed to the ones

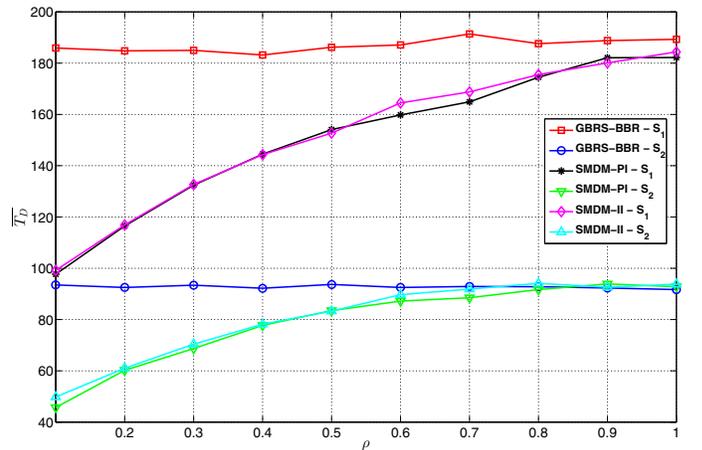


Fig. 3. Average transit delay under the GBRs-BBR, SMDM-PI, and SMDM-II schemes.

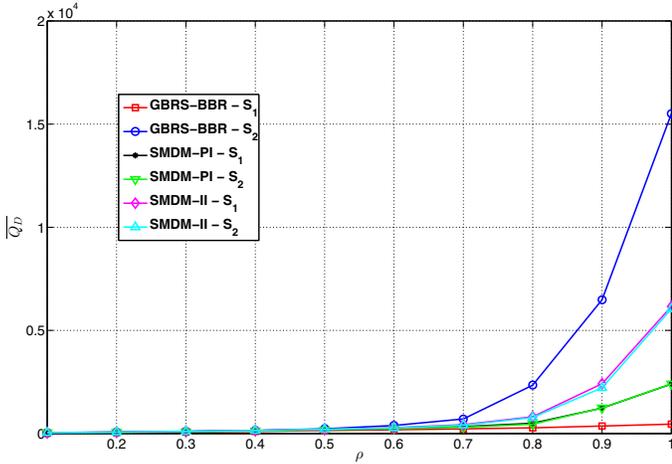


Fig. 4. Average queueing delay under the GBRS-BBR, SMDM-PI, and SMDM-II schemes.

generated by S_1 . However, Figure 4 reveals that the queueing delay experienced by the bundles released by S_2 is extremely high relative to the queueing delay of the bundles at S_1 . Knowing that both S_1 and S_2 are subject to the same bundle arrival rate, the only reason why bundles would accumulate at S_2 faster than they do at S_1 would be that S_2 is unable to clear out bundles as fast as S_1 is doing so. This is a direct implication from *Little's Theorem*, [7]. The fact that on the average, the queue at S_2 is bigger than the one at S_1 means that S_1 is monopolizing the capacity of the arriving vehicles and hence, leaving little spare room for the bundles queueing at S_2 .

The ability of the SMDM schemes to resolve the starvation problem incurred by GBRS-BBR can be explained as follows. SMDM enables S_1 to regulate its bundle release rate as a function of the number of bundles it believes to be queueing at S_2 . Clearly, under SMDM-PI, S_1 is privileged with an exact information to this end and as such it can optimally regulate its sending rate with a view to reserving an appropriately dimensioned part of the arriving vehicle's storage space to be filled by S_2 's bundles. So, S_1 no longer selfishly exhausts the storage capacity of the arriving vehicles and as a result enables S_2 to abandon the title of a starving node. At this level, as far as SMDM-II is concerned, what is relevant to the present study is the fact that similar conclusions can be drawn about the SMDM-II scheme. In other words, under SMDM-II, S_1 is instructed to regulate its bundle release rate based on imperfect information about the length of S_2 's buffer. This justifies why the end-to-end delay achieved by SMDM-II is slightly worse than that resulting from SMDM-PI, as clearly supported by Figure 2. Further insights into the behavior of SMDM-II can be obtained through a careful examination of $\overline{Q_D}$ and $\overline{T_D}$, the main components underlying $\overline{E_D}$. In fact, it is clear from Figure 3 that the transit delays as achieved by both source SRUs S_1 and S_2 under both the SMDM-PI and SMDM-II schemes are equivalent. This can be simply justified by the fact that both of SMDM-PI and SMDM-II select those vehicles that best contribute to the minimization of the average transit

delay to the destination SRU. The difference between the end-to-end delay achieved by SMDM-PI and the one incurred by SMDM-II stems from the sub-optimal behavior that SMDM-II exhibits with respect to starvation mitigation. This observation is corroborated through the results reported in Figure 4 that indicate a sub-optimal handling of the starvation problem by SMDM-II. Under this condition, unlike in the case of SMDM-PI, the amount of buffer space spared by S_1 under SMDM-II will be far from being perfectly dimensioned. In consequence, fewer bundles can be released by S_2 and S_2 's queue gets destabilized leading to longer queueing delays. Nonetheless, this increase in terms of queueing delay remains acceptable.

In light of the above, one can confidently argue that SMDM-II is able to tackle convincingly the dual challenge of starvation mitigation and delay-minimal bundle delivery while eliminating the need for the inappropriate assumption of perfect information availability.

V. CONCLUSION

SMDM is a bundle delivery strategy that was developed for the purpose of dually reducing the starvation problem and achieving delay-minimal bundle delivery in the context of ICRCNs. At the heart of SMDM though is the non-pragmatic assumption of perfect information availability at the upstream SRU. In a bid to relax this unrealistic assumption, this paper proposed to equip the upstream source SRU with imperfect information about the buffer status of other interposed source SRUs. This gave rise to the so-called SMDM-II (SMDM with Imperfect Information) variant. A simulation study revealed the ability of SMDM-II to achieve near-optimal performance when it comes to starvation resolution as well as delay-minimal bulk delivery. The development of delivery mechanisms based on imperfect information that are more sophisticated than SMDM-II is possible but the study presented herein is a small humble first step in the right direction. This is especially true since our numerical results proved the practical realizable SMDM-II to be an appealing alternative to the theoretical SMDM-PI delivery scheme.

ACKNOWLEDGMENT

The authors would like to thank Dr. Maurice Khabbaz for his valuable comments on the technical content of the present paper and for his help in the setup of Figure 1.

REFERENCES

- [1] W. Fawaz (*et al.*), "A First Step Towards the Resolution of the Starvation Problem in Multi-Point-to-Point ICRCNs", *IEEE Communications Letters*, 17:11, November 2013.
- [2] V. Cerf (*et al.*), "Delay-Tolerant Network Architecture", *IETF*, RFC 4838, April 2007.
- [3] V. Ramaiyan (*et al.*), "Delay-Optimal Schedule for a Two-Hop Vehicular Relay Network", *ACM/Springer Mobile Networks and Applications*, 15:1, May 2009.
- [4] M. Khabbaz (*et al.*), "Probabilistic Bundle Relaying Schemes in Two-Hop Vehicular Delay Tolerant Networks", *IEEE Communications Letters*, 15:3, March 2011.
- [5] R.A. Berry (*et al.*), "Communication over Fading Channels with Delay Constraints", *IEEE Transactions on Information Theory*, 48:5, May 2002.
- [6] D. Bertsekas "Dynamic Programming and Optimal Control", *Athena Scientific*, 2012, vol. I and II.
- [7] J. Little, "A proof of the queueing formula: $l = \lambda w$ ", *Operations Research* 9, vol. 3, 1961.