Solving the Coupon Collector's Problem for the Safety Beaconing in the IEEE 802.11p WAVE

Hyundoc Seo, Sangki Yun and Hyogon Kim
Korea University

***Abstract*—For the WAVE applications built on safety beacons, it is imperative that the neighboring vehicle information be collected as quickly and efficiently as possible. But the unreliability of broadcast transmission and the randomness of channel access in the IEEE 802.11p MAC hamper the collection process. Specifically, the process suffers from essentially a form of the classic Coupon Collector's Problem, where it takes longer and longer to obtain the remaining information. In this paper, we solve the problem by introducing the application-level acknowledgement of the safety beacons. We demonstrate that this optimization drastically reduces the collection completion time, eventually contributing to the safety and efficiency in WAVE-based systems.**

I. INTRODUCTION

In the looming IEEE 802.11p Wireless Access in Vehicular Environment (WAVE) systems [1], [2], the periodic broadcast messages from each vehicle reporting their position, speed, and direction [3] are expected to serve as the basis to build important traffic safety applications such as cooperative collision warning (CCW) [4] and safety message routing [5]. In each control channel interval (CCHI), these messages are transmitted on contention basis [6], at a frequency typically ranging from 1Hz [5] to 10Hz [4]–[7]. In the literature, this periodic broadcast generated by safety applications is frequently called the "beacon," but in order to avoid confusion with the IEEE 802.11 MAC layer beacon (a standard management frame), we will refer to it as the *safety beacon* in this paper.

Since the safety beacon is an application data, it is transmitted according to the IEEE 802.11e Extended Distributed Channel Access (EDCA) rule [6]. The contention-based access to the channel gives random and long-term fair transmission opportunities to the safety beacon senders. Since the safety beacons are independently scheduled at each vehicle, and the Sync Interval (SI) [2] configuration is independent of the safety beacon intervals, the safety beacons get to be randomly scattered over the SI. The SI is composed of two sub-periods, the aforementioned CCHI and the service channel interval (SCHI).

ACI	AC	CW_{min}	CW_{max}	AIFSN	TXOPLimit
1	Background	15	511	9	0
0	Best Effort	7	15	6	0
2	Video	3	7	3	0
3	Voice	3	7	2	0

TABLE I
EDCA PARAMETER SET USED ON THE CONTROL CHANNEL (CCH) IN IEEE 802.11P.

For the transmission during CCHI, the application must use WAVE Short Message Protocol (WSMP) to encapsulate the generated safety beacon and push it down to the 802.11p MAC layer [2]. If the safety beacon happens to be generated during a CCHI, it will be transmitted as soon as the channel is available in the same CCHI according to the EDCA rules. But if it is generated by the application during a SCHI, it has to wait until the next CCHI arrives. The message will be put in the MAC queue for the control channel, which as soon as the next CCHI arrives, the MAC will serve. This can cause a flash crowd of safety beacons attempting to access the channel at the beginning of every CCHI. Due to the flash crowd, many of the safety beacon broadcasts can be lost in collisions. The surviving safety beacons should be few, due to the small fixed contention window size (CW_{min}) defined for the control channel messages in the IEEE 802.11p standard [1] (Table I). Moreover, since the safety beacons are broadcast, there is neither MAC layer retransmission nor contention window backoff. Therefore, only the minimum contention window size CW_{min} is applied, which is 15 (slots) with $ACI = 1$. (Although not specified by the standard, we expect the safety beacons to be sent in Access Class Index (ACI) of 1, in order not to interfere with more urgent safety messages.)

II. PROBLEM STATEMENT

A. Beacon flash crowd

The undesirable implication of all this is that those safety beacons generated in the SCHI will all attempt to transmit in the first 15 slots at the beginning of the CCHI. With a large number of vehicles generating safety beacons in the SCHI, the collision probability can be high with only a small number of safety beacons successfully delivered. For example, Fig. 1 shows the number of successful broadcast when there are $2n$ vehicles in mutual communication range where half of them generate the safety beacon in the SCHI. We assume that the sizes of SCHI and CCHI are the same, as is specified as the default in the IEEE 1609.4 standard. Under the 802.11e parameters for $ACI = 1$, the maximum achievable successes is 6 beacons when there are approximately 15 vehicles. As the number of vehicles increases, the number of successful beacons is even less.

B. Coupon Collector's Problem

Although the "beacon flash crowd" problem is certainly a pathological aspect of the safety beaconing under the current WAVE standards, it may *not* be entirely unavoidable and is not

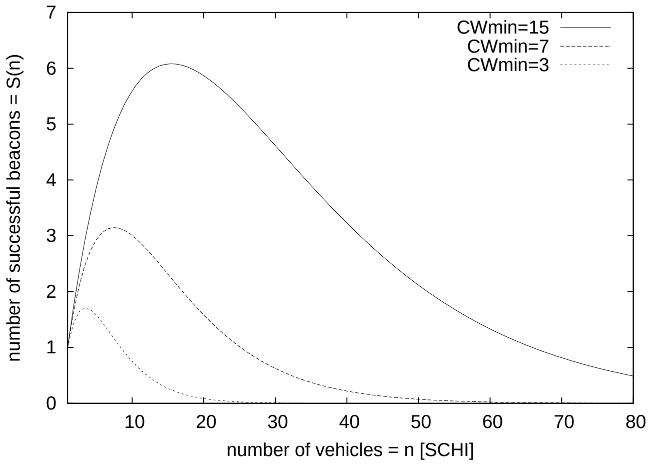

Fig. 1. Number of successfully transmitted beacons during initial CCHI contention, given n vehicles generate safety beacons in SCHI.

the main focus of this paper. For instance, we could consider synchronizing the safety beaconing application with the CCHI to avoid the concentration, although it would require cross-layer cooperation between the application and the MAC/PHY.

No matter how the safety beacons are distributed, however, the safety beaconing is still subject to a more fundamental problem. Since some of the safety beacons are inevitably lost[1] in each CCHI as long as there are collisions, the collection process should span multiple CCHIs. The problem is that the vehicular nodes that successfully transmitted safety beacons will contend for the channel on level ground with those that failed due to the collision or bad channel in subsequent CCHIs, as the 802.11p MAC is fair. It renders the whole beacon collection process into a "sampling with replacement" experiment. And a well known pathology in the sampling with replacement is the Coupon Collector's Problem [8]. As we will discuss below, it vastly elongates the time for collecting the neighboring vehicle information. This goes counter to the objective of the applications based on safety beaconing, where they need to gather the neighboring vehicle information in the shortest amount of time, or equivalently, the most exhaustively in the given amount of time.

The classic Coupon Collector's problem is about time T to collect all n coupons when each collection operation uniformly randomly samples with replacement one of the n coupons. It is well known that $T = n \cdot H_n$, where H_n is the harmonic number. Not only the collection efficiency drops towards the end of the collection process with a given n, but also the coupon collection time superlinearly increases with n. It is because H_n does not converge and keeps increasing with n. Namely, as the number of vehicles n grows, it takes increasingly longer to collect the information about those n vehicles (or part thereof).

The difference in our case, however, is that we get to choose b safety beacons instead of 1 in each CCHI, where b is a random variable, not a constant. It is determined as an output of the 802.11 MAC contention process throughout the CCHI. For instance, $\bar{b} = S(n)$ in Fig. 1. The closed form solution for this b-ary coupon collection is difficult to find, even when b is a constant [9]. A Monte-Carlo simulation in Fig. 2 shows the number of rounds for which n distinct vehicles are collected for a batch size of $\bar{b}$. Given the default 1609.4 SI length of 100ms, 10 rounds correspond to 1 second in the figure.

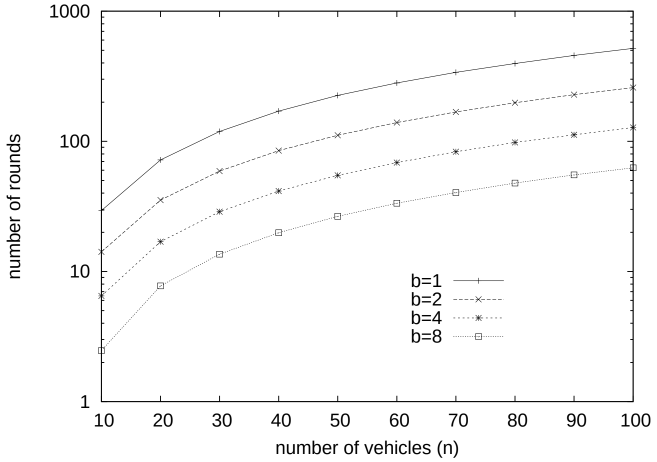

Fig. 2. Average number of rounds for complete coupon collection with a batch size of b.

The safety and/or efficiency implication of the Coupon Collector's problem can be significant. For instance, when a new vehicle merges to the ongoing vehicular traffic flow, it might be the victim of the safety beacon losses that other vehicles may recognize the vehicle much later than when it actually arrives. When there are a large number vehicles, the delay could easily reach a few tens or even hundreds of seconds, and the real-time applications such as CCW [4] may not safely function. And for applications such as safety message routing [5], it may take much longer to reflect the changes in the topology.

Below, we propose a solution approach to the Coupon Collector's Problem for the safety beaconing in the 802.11p WAVE environment. Our solution approach has several merits:

- It works mostly on application level, given a cross-layer support for contention level estimation. Therefore, no change in the current WAVE standard suit itself is required. We plan to remove the cross-layer component and make it work purely on application level in our future work.
- It works well even under severe beacon flash crowd situation as the current WAVE standards could be subject to. Again, this obviates the need to revise the WAVE standards to address the flash crowd problem.
- It drastically reduces the safety beacon collection time with small overhead in safety beacon size. For instance, for $n = 100$ vehicles, a factor of 50 reduction is achieved.

---

[1] Although we can think of using non-zero values for the IEEE 1609.4 *Repeats* parameter to obtain higher delivery probability[2], it is not without cost since will in turn increase the collision probability. Moreover, more urgent messages could be hindered by the increased number of safety beacons. In this paper, we assume that Repeats is set to 0.

III. MITIGATING THE COUPON COLLECTOR'S PROBLEM

The essence of the Coupon Collector's Problem for WAVE safety beaconing is that the vehicular nodes get equal channel access probability, regardless of whether the previous beacon transmission succeeded or failed due to collision or bad channel. The solution is simple, then: suppress the WAVE nodes that succeeded in the previous attempt from contending for the channel for the next few safety beacon intervals. By abstaining from contention, they will allow the failed nodes to have opportunities to transmit. It has two effects on the safety beaconing system.

- It transforms the safety beacon collection into the sampling *with* replacement. It essentially *eliminates* the Coupon Collector's Problem, as well as mitigates the application level unfairness.
- It reduces the contention hence the collisions. This is a good side-effect, and it further helps the previously failed nodes to succeed in subsequent CCHIs.

But there are two technical issues in implementing the application level feedback and safety beacon suppression. The first is how to let the successful nodes know that they succeeded and are required to suppress the safety beacon transmission for less fortunate ones. As there is no acknowledgement (ACK) support for broadcast in the 802.11p MAC, the successful nodes should get the feedback in some other way. So, we choose to introduce the application-level acknowledgement. But the application-level ACK is not sent to the successful node in unicast, but is piggybacked on the transmitted safety beacon of the beacon receiver.

The second issue is how long the successful node should suppress the beacon transmission. In sampling with replacement, the total time required to collect all coupons is simply n/b, where b is the (average) number of collections per each period. If each successful node abstains for $n/b-1$ periods, the collection will proceed smoothly. So we apply this idea in the proposed scheme. The design of the modified safety beaconing algorithm discussed below addresses the two technical issues.

A. Feedback

Suppose a safety beacon from vehicle x, denoted $SB(x)$, is successfully delivered to vehicle y. The beaconing application at y attaches the sender address x in its next beacon $SB(y)$. Note that $SB(y)$ is not unicast to x, as it would cause ACK implosion at x. The ACK is piggybacked on normal safety beaconing message exchange. Thus each safety beacon not only notifies other vehicles of its movement information, but also contains the list of successful transmitters in the immediately preceding CCHI. We will denote this beacon carrying the identities of the recently successfully transmitting vehicles by $SB(y; x_1, x_2, \ldots, x_k)$, if there were k successful transmitters excluding y. The list of successfully received beacons is purged at the end of each CCHI. Namely, the list $(x_1, x_2, \ldots, x_k)$ is removed, but the number of detected beacons k is used to compute the number of beacon-suppressed CCHIs.

B. Beacon suppression

When a vehicle x receives a beacon $SB(y; x_1, x_2, \ldots, x_k)$ that acknowledges its previously transmitted safety beacon, i.e. $\exists_i \; x_i = x$, it suppresses the generation of the next beacon for a certain amount of time. The approach used in this paper is that given the estimate $\hat{n}$ of the number of vehicles in the neighborhood and the number of safety beacons k successfully delivered during one CCHI, then the acknowledged node x suppresses the next safety beacon for the next $D = \hat{n}/k - 1$ CCHIs. Unfortunately, estimating the precise number of neighboring vehicles is not straightforward. In this paper, we assume that we obtain the average number of idle slots from Clear Channel Assessment (CCA) to estimate the contending population in the manner of [10], although other APIs might be made available in the future to detect the idle time distribution [11] directly from the application, utilizing the IEEE 802.11k [12].

C. Pseudocode

Algorithms 1 and 2 are the pseudocode of the algorithm run on the sender and the receiver sides, respectively. Below, v is the vehicle under consideration, $\hat{n}(v)$ is the estimated number of vehicles in the neighborhood of v, k is the number of received safety beacons at v in the previous CCHI, t is the time index of the current CCHI, and L is the list of received safety beacon source addresses from previous CCHI.

Algorithm 1 Sender side: transmit a safety beacon $SB(v)$

```
1:  while CCHI do
2:    if $BeaconsToSend \neq 0$ then
3:      if $D = 0$ then
4:        Transmit $SB(v; L_v[t-1])$
5:        $BeaconsToSend \leftarrow BeaconsToSend - 1$
6:      else
7:        /* Suppress transmission */
8:        Drop beacon; $BeaconsToSend \leftarrow 0$
9:      end if
10:   end if
11: end while
12: $D \leftarrow D - 1$
```

In the sender side algorithm, the sender checks if it still has a beacon to send (2). If so, it checks if the beacon needs to be suppressed by checking D (3). Notice that a vehicular node waits D CCHIs after a successful transmission. When a beacon arrives from the application before the wait is over, the beacon is dropped (8). For this, we can use the MAC queue flush primitive provided for canceling the low-priority transmission in the IEEE 1609 standard [2]. But if this CCHI is where the beacon transmission(s) should take place, the vehicle attempts to transmit the given number of safety beacons (4). We assume in this paper that a single safety beacon is transmitted in every CCHI, but the given safety beaconing application may decide to use the IEEE 1609.4 *Repeats* parameter to control the number MAC layer retransmissions. Piggybacked on the beacon is the list L_v of successful transmitters from the previous CCHI period $(t-1)$ heard by v (4). Finally, the countdown of D takes place at the end of each CCHI (12).

Algorithm 2 Receiver side: process $SB(w; L_w)$

```
1:  while CCHI do
2:      Receive $SB(w; L_w[t-1])$
3:      if $v \in L_w$ then
4:          /* $SB(v)$ has been confirmed */
5:          $D_{new} \leftarrow \hat{n}/sizeof(L_v[t-1]) - 1$
6:      end if
7:      /* update the list */
8:      $L_v[t] \leftarrow L_v[t] \cup w$
9:  end while
10: if $D = 0$ then
11:     $D \leftarrow D_{new}$
12: end if
```

In the receiver side algorithm, when a vehicle v successfully decodes a safety beacon from some other vehicle w, it extracts L_w as well as add w to the list of heard vehicles in the current CCHI (8). In case it is in the acknowledged list of successful beacon senders (3), a suppression is scheduled (5), which will take effect at the end of the current CCHI (11) if D is not running.

IV. PERFORMANCE EVALUATION

In this section, we evaluate the performance of the proposed scheme through extensive simulation. We assume that the original beaconing frequency is $f = 10$ so that every vehicle transmits the safety beacon every 100ms. And following the IEEE 1609 standard default values [2], we assume that the SI is 100ms, which is equally split between the CCHI and the SCHI. We use the Qualnet 4.5 simulator, varying the number of vehicles n in the mutual communication range from 10 to 100.

A. Beacon collection time

Fig. 3 shows the time in seconds to identify all vehicles in the communication range. We observe that by transforming the safety beacon collection into sampling without replacement, the Coupon Collector's Problem disappears, which leads to the huge reduction in time to identify neighboring vehicles. Being the sampling without replacement, the collection process takes linear time, as confirmed by the simulation (inset). For the smallest vehicle densities such as $n = 10$, two schemes perform comparably since the safety beacons suffer few collision losses. From $n = 20$, however, the collection time of the unmodified scheme starts a stiff rise. When we reach $n = 90$ vehicles, for instance, it takes approximately 150 seconds to identify all of them whereas it takes less than 5 seconds in the modified scheme. For $n = 100$, the modified method still records less than 5 second, but the original beaconing method explodes to 280 seconds.

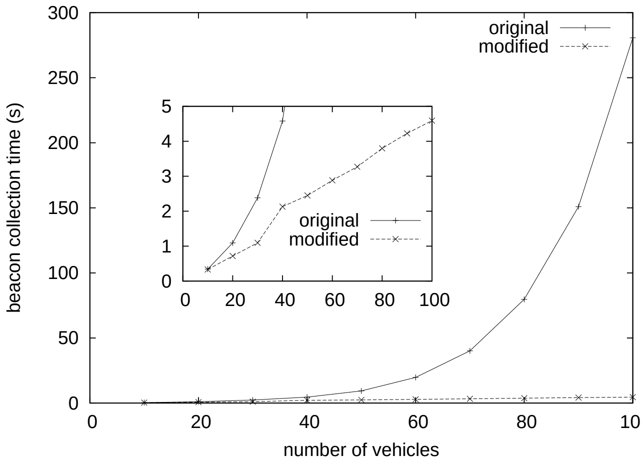

Fig. 3. Beacon collection time with and without the proposed modification.

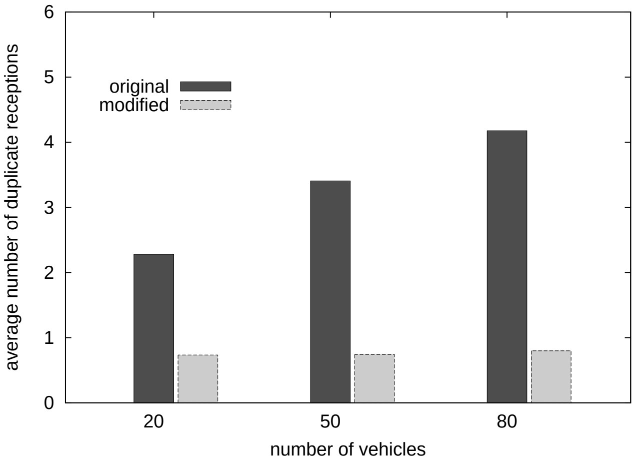

Fig. 4. Number of duplicate receptions.

Fig. 4 compares the number of duplicately[2] received copies for each beacon in a single beacon collection time T. For instance, the original beaconing at $n = 80$ receives on average more than 5 beacons from the same vehicle. Since this is the average, for some vehicles the degree of redundancy should be much higher. In contrast, the modified scheme is stable at approximately 0.7. Ideally, the modified scheme should completely eliminate the duplicate copies. The source of this residual inefficiency is the error in the estimation of D, due to randomness in k. In particular, when D is underestimated, the previously successful node can prematurely send the next safety beacon before the current round of beacon collection completes, hence the duplicate receptions. In future work, we will refine the estimation of D so that the inefficiency is further reduced, accelerating the collection process. Fig. 4 shows that not only the original safety beaconing is exposed to higher collision probability (see Fig. 6 below), but even if the safety beacon is successfully delivered, many of them are redundant. And as the number of neighboring vehicles increase, the redundancy also grows.

Fig. 5 shows how the safety beacon collection proceeds over time. The original collection method clearly exhibits

---

[2] The "duplicate" copy is not really an identical safety beacon, since as long as the vehicle moves, the contents (e.g. GPS position) keeps changing. The meaning of duplicate here is that the same vehicular node generates multiple safety beacons in the same round of beacon collection.

the classic Coupon Collector's Problem, with its performance degrading quickly with the number of beaconing vehicles n. From the figure, we can also obtain the time required to collect the beacons of a fraction of the given vehicle population. For instance, if for some application 80% collection is acceptable, we can see that it takes over 20 seconds in the original scheme versus 2 seconds in the modified scheme.

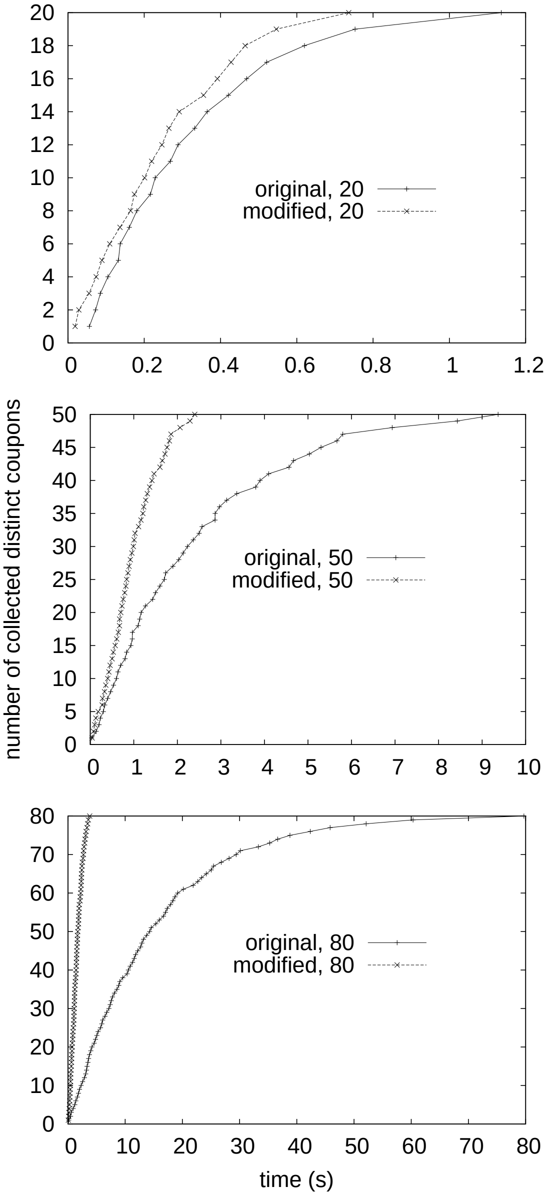

Fig. 5. Number of collected safety beacons: with $n = 20$ (top), $n = 50$ (middle), and $n = 80$ (bottom).

B. Collision probability

Fig. 6 shows that as a fallout of the beaconing mechanism modification, the MAC level contention is reduced for all n. Due to the lower collision probability, the safety beacon collection process is further accelerated in the modified scheme. However, the collision probability is still surprisingly high. It raises an interesting question whether the current beacon suppression period $D = \hat{n}/k$ is sufficiently long. So in our future work, we will investigate the impact of using even larger D to decrease the collision probability, and see if it reduces the beacon collection time T further.

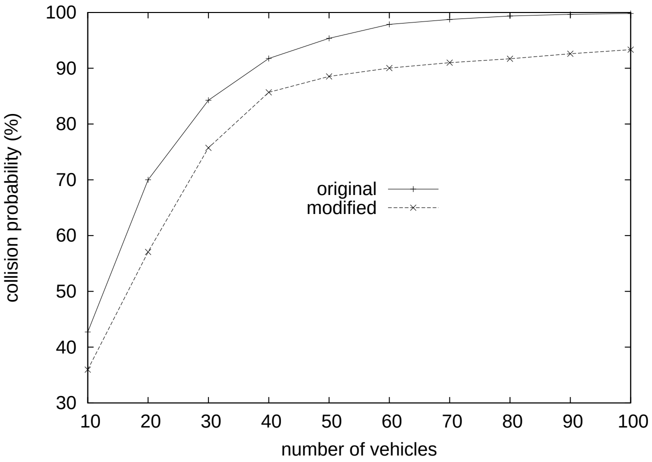

Fig. 6. Comparison of collision probability.

C. Fairness

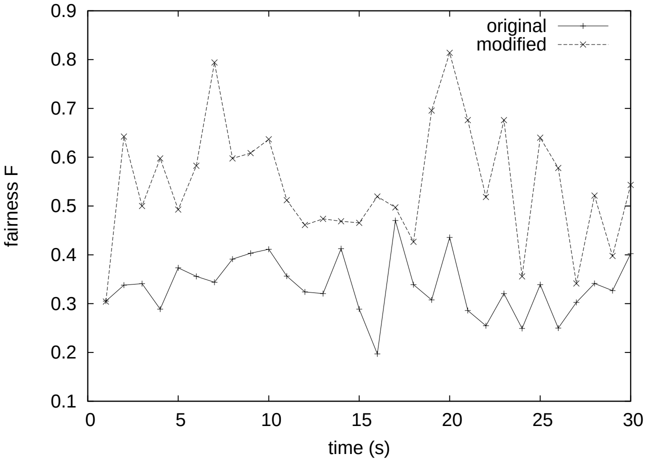

Fig. 7. Comparison of the Jain fairness index, $n = 50$, $\tau = 1$s.

On the MAC level, the 802.11p channel access is fair. However, on the application level, there are vehicles that fail to deliver their safety beacons in the given CCHI. If they are forced to compete with other vehicles with equal channel access probability in the subsequent CCHIs, it is unfair. Therefore, preferentially treating them on the application level contributes to the fair access. Here we show that the proposed scheme indeed improves fairness. Given the number of successful transmissions $S_v(n;\tau)$ for each vehicle v in a time period τ, we can compute the Jain fairness index

$$F(n;\tau) = \frac{\left(\sum_v S_v(n;\tau)\right)^2}{n \cdot \left(\sum_v S_v(n;\tau)^2\right)}$$

for the two compared schemes. For instance, let $\tau =1$ second and $n = 50$. Fig. 7 clearly shows the fairness of the modified scheme is always better than that of the original safety beaconing. This confirms from another angle that the modified safety beaconing scheme more efficiently uses the wireless bandwidth so that more vehicles can report their movement information to their neighbors. Finally, the generally low fairness index for both schemes stems from the absolute shortage of survivors from the 802.11p MAC contention. See Fig. 5 for $n = 50$, and notice that in $\tau = 1$ second the two schemes can only collect approximately 15 and 30 (i.e. less than n) safety beacons, respectively. With the nodes with no successful transmissions and some with redundant transmissions (see Fig. 4), the fairness drops.

V. CONCLUSION

In this paper, we show that the randomness of the channel access in the IEEE 802.11p MAC and the unreliable nature of the safety beacon broadcast lead to a form of Coupon Collector's Problem. By introducing the application level acknowledgement piggybacked on the safety beacons, we transform the beacon collection process from sampling with replacement to sampling without replacement. It eliminates the Coupon Collector's Problem and significantly accelerates the collection process, e.g. more than 50 times speedup with 100 vehicles. As the safety beaconing provides the basis for many safety-related applications in the looming WAVE systems, we expect the proposed solution will have far reaching impact on the safety and efficiency of vehicular networking.

ACKNOWLEDGMENT

This work was supported in part by the Korea Research Foundation Grant 2009-0080413 and Ubiquitous Computing and Network(UCN) Project, Knowledge and Economy Frontier R&D Program of the Ministry of Knowledge Economy(MKE) in Korea as a result of UCN's subproject 09C1-C1-20S

REFERENCES

[1] IEEE 802.11 WG, "Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications, Amendment 7: Wireless Access in Vehicular Environments," IEEE P802.11p/D5.0, 2008.

[2] IEEE 1609 WG, "IEEE Trial-Use Standard for Wireless Access in Vehicular Environments (WAVE) — Multi-channel Operation," IEEE 1609.4-2006.

[3] M. Torrent-Moreno, P. Santi, and H. Hartenstein, "Fair Sharing of Bandwidth in VANETs," in proceedings of ACM VANET workshop, 2005.

[4] T. ElBatt, S. Goel, G. Holland, H. Krishnan, J. Parikh, "Cooperative Collision Warning using Dedicated Short Range Wireless Communications," in proceedings of ACM VANET workshop, 2006.

[5] O. Tonguz, N. Wisitpongphan, F. Bai, P. Mudalige, and V. Sadekar, "Broadcasting in VANET," in proceedings of IEEE Infocom MOVE workshop, 2008.

[6] IEEE 802.11 WG, "Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications, Amendment 8: Medium Access Control (MAC) Quality of Service Enhancements," IEEE 802.11e-2005, 2005.

[7] Q. Chen, D. Jiang, V. Taliwal, L. Delgrossi, "IEEE 802.11 based Vehicular Communication Simulation Design for NS-2," in proceedings of ACM VANET workshop, 2006.

[8] N. Alon, J.H. Spencer and P. Erdos, *The Probabilistic Method*, John Wiley & Sons, Inc., NY, 1992.

[9] Y. M. Minsky and F. B. Schneider, "Tolerating malicioius gossip," *Distributed Computing*, 16(1), pp. 49–68, 2003.

[10] M. Heusse, F. Rousseau, R. Guillier, A. Duda, "Idle Sense: An Optimal Access Method for High Throughput and Fairness in Rate Diverse Wireless LANs," in proceedings of ACM SIGCOMM, 2005.

[11] S. Mangold, Z. Zhong, G. R. Hiertz and B. Walke, "IEEE 802.11e/802.11k wireless LAN: spectrum awareness for distributed resource sharing," Wireless Communication and Mobile Computing, Vol. 4, pp. 881–902, Dec. 2004.

[12] IEEE 802.11 WG, "Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications, Amendment 1: Radio Resource Measurement of Wireless LANs," IEEE 802.11k-2008, 2008.