On Colliding First Messages in Slotted ALOHA

(Invited Paper)

Christian Bettstetter[*#], Günther Brandner[*], and Robert Vilzmann[+]

[*] University of Klagenfurt, Mobile Systems Group, Institute of Networked and Embedded Systems, Austria

[#] Lakeside Labs GmbH, Klagenfurt, Austria

[+] Technische Universität München, Institute of Communication Networks, Munich, Germany

***Abstract*—Considering n nodes performing random access using ALOHA with s slots, we study the probability that there occurs a non-colliding message in the first non-empty slot. If each node transmits with probability p in each slot and the number of slots is sufficiently large, a non-colliding first message occurs with probability $\Phi = 1 - np/2$ for large n and small np. If the number of slots is limited, the probability Φ is lower but can be maximized choosing an optimal p. To maximize Φ further, nodes can apply a slot-dependent transmit probability p_i with $i = 1, \ldots, s$. It is shown that a "slow start strategy," in which p_i is low for low i and increases with increasing i, is beneficial. Our main contribution is an equation for the p_i values that maximize Φ. We analyze how a higher probability of a non-colliding first message comes at the price of an increased delay of such a message. Besides being of interest for the theory of random access, the results are practically applicable to node selection protocols, such as relay selection in cooperative wireless networks.**

***Index Terms*—Medium access control, ALOHA, first message, collision probability, random access.**

I. Introduction and Motivation

Several techniques in wireless and wired networking require some form of node selection mechanism in which one out of multiple nodes is selected, in a distributed manner, to undertake a certain task. Such a mechanism is needed, for example, for cooperative relaying techniques in wireless networks to choose a "relay node" [1] and for data processing techniques in sensor networks to choose a "data gathering node" [2].

Node selection operates above the medium access control (MAC) layer and can be achieved in two steps: First, a set of candidate nodes is determined. Each node in this set must fulfill a certain criterion (or several criteria) that qualifies to serve as a selected node. For example, the node's battery level or/and its link quality to a destination must be above a certain threshold value. Second, all candidate nodes compete for random access on the shared medium ("channel"). The node that successfully accesses the channel first wins the selection process and acts as selected node. In wireless systems, the process typically starts with a query message broadcasted on the channel. For example, a node asks all its neighbors: "I need a relay with a bit-error-rate better than 10^{-4} to the destination." Each neighbor fulfilling this criterion tries to gain access to the channel, to send a positive reply to the querying node. The node that answers first will act as relay node.

It can be argued [3] that such a *first message* on the channel is more important than subsequent reply messages. Thus, from a MAC layer perspective, the probability of a non-colliding first message should be maximized, while the collision probability of later reply messages is less important [3].

This MAC design issue is the topic of this paper. Assuming slotted ALOHA, we discuss the following issues: What is the probability that there is a first message that does not collide? How can we maximize this probability? What is the tradeoff between this probability and the delay of the selection process?

Section II gives a formal description of the modeling assumptions and the problem. Section III derives the probability of a non-colliding first message and proposes medium access strategies that maximize this probability. Section IV analyzes the delay of a first message. Section V addresses related work.

II. Modeling Assumptions and Problem Statement

Consider a set of n nodes on an uplink channel. The MAC layer follows a slotted ALOHA scheme without carrier sensing [4]. All nodes compete for random access to s shared time slots. Each slot has the same duration τ. The duration of a message is assumed to fit into one slot.

Let p_{ij} be the probability that node j transmits a message in slot i, where $j \in \{1, \ldots, n\}$ and $i \in \{1, \ldots, s\}$. We assume that all nodes behave statistically the same; hence we skip the index j and write $p_i := p_{ij} \, \forall \, j$.

If two or more nodes transmit in the same slot i, a message collision occurs, making it difficult for the receiver to decode any of the sent messages during this slot. If a message does not suffer from a collision, it is called a non-colliding message.

A slot is empty if no node transmits during this slot. The first non-empty slot is the slot i in which at least one message is sent while previous slots $1, \ldots, i-1$ were empty. A message sent in the first non-empty slot is called a first message.

In this paper we discuss three major problems: What is the probability that, with these modeling assumptions, a non-colliding first message occurs? How should we choose the transmission probabilities p_i to maximize this probability? What is the tradeoff between a high probability and a low delay (an early slot) of a non-colliding first message?

III. Probability of a Non-Colliding First Message

The number of messages sent in a given slot i can be described by a random variable M_i. A message sent in slot i does not collide if exactly one node transmits during this slot; the probability for this event being denoted by $\mathbb{P}[M_i = 1]$. A message sent in slot i is the first message to be sent if there was no message in previous slots; the probability for this event is $\prod_{w=1}^{i-1} \mathbb{P}\left[M_w = 0\right]$ for $i > 1$. Thus, the probability that there occurs within s slots a first message that does not collide is

$$\begin{aligned} \Phi(n, s, p_1, \ldots, p_s) &= \\ &= \mathbb{P}\left[M_1 = 1\right] + \sum_{i=2}^{s} \left(\prod_{w=1}^{i-1} \mathbb{P}\left[M_w = 0\right] \right) \mathbb{P}\left[M_i = 1\right]. \end{aligned} \tag{1}$$

With our assumptions, the probability that k messages are sent in a given slot i is given by the binomial distribution

$$\mathbb{P}\left[M_i = k\right] = \binom{n}{k} p_i^k \left(1 - p_i\right)^{n-k} \tag{2}$$

with $k \in \{0, \ldots, n\}$. Combining (1) and (2) yields

$$\begin{aligned} \Phi(n, s, p_1, \ldots, p_s) &= \\ &= n \sum_{i=1}^{s} \left(\prod_{w=0}^{i-1} \left(1 - p_w\right)^n \right) p_i \left(1 - p_i\right)^{n-1} \end{aligned} \tag{3}$$

with $p_0 := 0$.

A. Same Transmission Probability for Each Slot

We first analyze a scenario where the transmission probabilities are the same for each slot ($p_i = p\,\forall i$). Equation (3) simplifies to

$$\begin{aligned} \Phi(n, s, p) &= np \sum_{i=1}^{s} \left(1 - p\right)^{ni-1} \\ &= \frac{np\left(1 - p\right)^{n-1} \left(1 - \left(1 - p\right)^{ns}\right)}{1 - \left(1 - p\right)^n}. \end{aligned} \tag{4}$$

1) General Behavior: Let us discuss the qualitative influence of the three parameters on the probability Φ. Figure 1 plots (4) for selected values of s and n as a function of p.

If the product np is low, it is very likely that no message is transmitted at all within s slots. Hence, the probability Φ is low for $np \ll 1$ (see Figs. 1(a) and (b)). If we increase n or p or both, the likelihood that no message is transmitted decreases, which in turn increases Φ. For some (n, p)-pairs, Φ achieves a maximum, whose value also depends on s. Increasing n or p or both further decreases the probability Φ again, as the likelihood of message collisions increases. Finally, for very high np, the value of Φ approaches zero. In summary, two probabilities govern the behavior of Φ: the probability that a (first) message does actually occur within s slots and the probability that a (first) message collides.

The impact of s for given n is illustrated in Figs. 1(c) and (d). If only few slots are available, there is a high likelihood that no message is transmitted, yielding a low Φ. Increasing s decreases this likelihood and thus leads to a higher Φ. As we increase s further, the no-message-likelihood approaches zero and the probability Φ tends to an asymptotic limit, whose value depends on n and p.

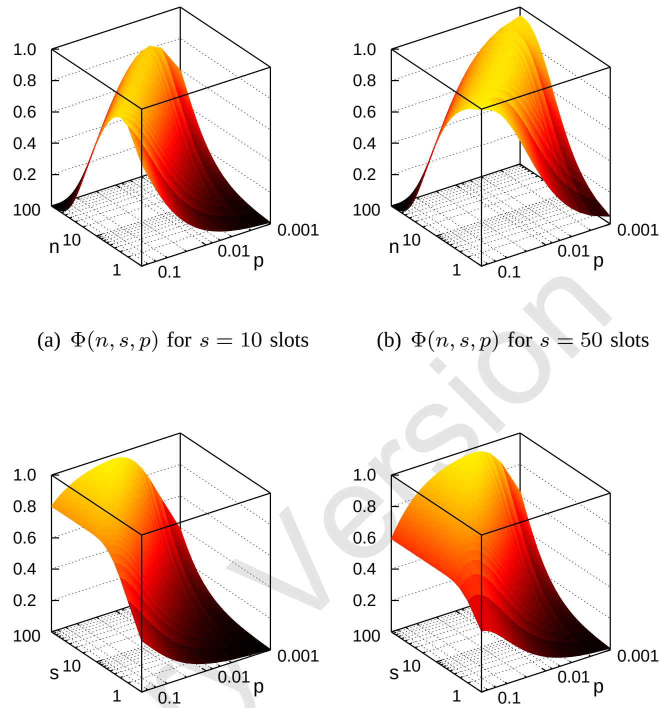

(a) $\Phi(n, s, p)$ for $s = 10$ slots (b) $\Phi(n, s, p)$ for $s = 50$ slots

(c) $\Phi(n, s, p)$ for $n = 5$ nodes (d) $\Phi(n, s, p)$ for $n = 10$ nodes

Fig. 1. Each of n nodes transmits with probability p in each of s time slots. The probability that there is a non-colliding first message is $\Phi(n, s, p)$.

2) Many Slots: Assuming that there are sufficiently many slots, such that a first message occurs for sure, yields

$$\Phi(n, p) := \lim_{s \to \infty} \Phi(n, s, p) = \frac{np\left(1 - p\right)^{n-1}}{1 - \left(1 - p\right)^n}. \tag{5}$$

Note that this term is equal to the probability that $M = 1$ message has been sent in an arbitrarily chosen time slot under the condition that at least one message has been sent in this slot ($M \geq 1$). Furthermore, recall that $\Phi(n, p)$ is an upper bound for $\Phi(n, s, p)$ for $s < \infty$.

Figure 2 shows some (n, p)-pairs leading to a non-colliding first message with probability $99\,\%$, $90\,\%$, $80\,\%$, or $50\,\%$.

Example: To achieve a non-colliding first message with a probability above $90\,\%$ on a shared channel with $n = 10$ nodes, each node must access a slot with probability lower than $p = 0.02$. If $n = 20$ nodes are present with the same p, the non-colliding first message probability decreases to roughly $80\,\%$. To retain a probability of $90\,\%$ with $n = 20$ nodes, the slot access probability must be decreased to $p = 0.01$.

3) Many Nodes: If n becomes large while the total traffic load $\lambda = np$ remains constant, we obtain

$$\Phi(\lambda) := \lim_{n \to \infty} \Phi(n, p) = \frac{\lambda}{e^\lambda - 1}. \tag{6}$$

It can be realized by inspection of (6) that if n is increased by a certain factor, we must decrease p by the same factor — and vice versa — to keep the same $\Phi(\lambda)$.

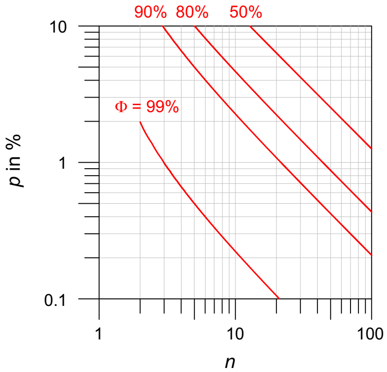

Fig. 2. Each of n nodes transmits with probability p in each of infinitely many time slots. The probability of a non-colliding first message is $\Phi(n, p)$. The figure shows some points of equal Φ (contour lines) in the n/p-plane.

Developing (6) into a series yields $\Phi(\lambda) = 1 - \frac{1}{2}\lambda + \frac{1}{12}\lambda^2 - \frac{1}{720}\lambda^4 + O(\lambda^6)$. This expression is useful for the practically relevant case if $\lambda \ll 1$. It gives us the following rule of thumb: for large n and small np, the probability for a non-colliding first message is

$$\Phi(n, p) \approx 1 - \frac{np}{2}\,. \tag{7}$$

Example: To achieve a non-colliding first message with a probability of at least $90\,\%$, the traffic load must be lower than $np = 0.2$ (also see Fig. 2).

4) Optimal Transmission Probability: Let us consider again a scenario with a finite number of nodes n and a finite number of slots s. In such a scenario, nodes can choose an optimal transmission probability p, such that the probability for a non-colliding first message according to (4) becomes maximum. This transmission strategy is called "slot-independent optimization strategy" in the following. Using numerical optimization, we obtain Figure 3(a), showing us the maximum possible Φ that can be achieved for a given (n, s)-pair. Figure 3(b) shows some contour lines of this plot, namely the (n, p)-pairs leading to $\max \Phi = 98\,\%$, 95%, $90\,\%$, $80\,\%$, $70\,\%$, or $50\,\%$, respectively. Table I gives some examples of optimal p-values for selected (n, s)-pairs.

TABLE I
TRANSMISSION PROBABILITIES p MAXIMIZING $\Phi(n, s, p)$. VALUES IN %.

$s \rightarrow$ / $\downarrow n$	1	2	5	10	20	50	100
2	50.0	39.42	25.29	16.59	10.29	5.14	2.94
5	20.0	15.35	9.57	6.19	3.82	1.90	1.10
10	10.0	7.62	4.72	3.04	1.87	0.94	0.54
20	5.0	3.80	2.34	1.51	0.93	0.46	0.27
50	2.0	1.52	0.93	0.60	0.37	0.18	0.11
100	1.0	0.76	0.46	0.30	0.18	0.10	0.05

Example: On a channel with $n = 10$ nodes and $s = 50$ slots, we can achieve a non-colliding first message with a probability of at most $95\,\%$. This performance is achieved setting $p = 0.94\,\%$; higher and lower p lead to worse performance. If only $s = 10$ slots are available, Φ can never be more than $83\,\%$.

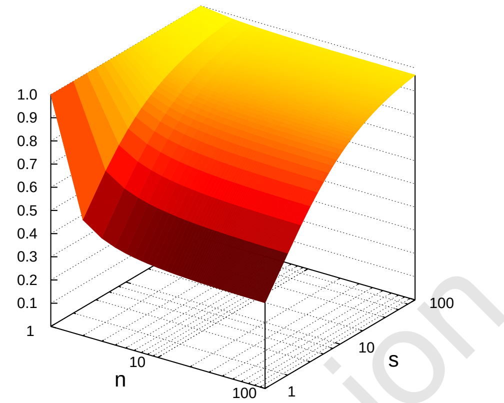

(a) $\max \Phi(n, s, p)$

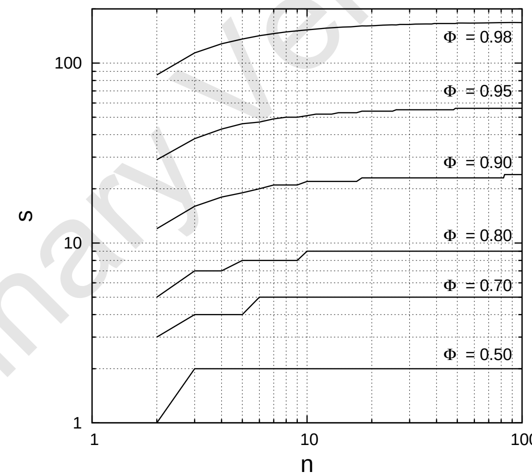

(b) Contour lines of $\max \Phi(n, s, p)$

Fig. 3. On a channel with n nodes and s time slots, each node transmits with probability p, where p is chosen to maximize the probability $\Phi(n, s, p)$ of a non-colliding first message.

On the other hand, $s = 10$ slots are also sufficient for as many as $n = 100$ nodes to yield about the same Φ.

B. Optimizing the Transmission Probabilities for Each Slot

The following questions arise: Can we improve Φ by choosing appropriate transmission probabilities p_i for each slot i individually? If so, what is a good (or even optimal) strategy for setting $\{p_1, p_2, \ldots, p_s\}$ for given n?

A key observation toward the solution of this optimization problem is as follows. Let us consider the last time slot $i = s$. The optimum transmission probability for this slot is independent of previous slots due to the following reasons:

- If a first message was already transmitted in a previous slot, the last slot is irrelevant anyway.
- If a first message has not been transmitted before, the system is the same as a system with a single slot ($s = 1$), hence has to be optimized in the same way.

This statement can be repeated in an iterative manner for all slots, starting from the last slot. Choosing an optimal transmission probability for the second last slot $i = s-1$, the previous $(s-2)$ slots are irrelevant — it is only important that there is another following slot. Hence, the optimal probability p_{s-1} is the same as for a two-slot channel. In general, the optimal transmission probability for slot $(s-k)$ with $k \in \{0, \ldots, s-1\}$ on a channel with s slots is the same as the optimal transmission probability of the first slot on a channel with $(k+1)$ slots. In other words, the optimal probabilities for the last $(k+1)$ slots on a channel with an arbitrary number of slots are identical to the optimal probabilities on a channel with $(k+1)$ slots.

A consequence from this observation is as follows: To maximize the overall probability Φ for given n and given s, each transmission probability p_i can be optimized individually and independently of the transmission probabilities of other slots. We start by optimizing the transmission probability of the last slot, continue with the second last slot, and so on, until we have calculated all s values.

Let us apply this calculation algorithm using (1). Taking into account the last slot, it is straightforward to observe that the transmission probability $p_s = 1/n$ maximizes $\Phi(n, s, p_s)$. Taking into account the second last slot, the probability $\Phi(n, s, p_{s-1}, p_s)$ is maximized by the transmission probability

$$p_{s-1} = \frac{1}{n} \cdot \frac{n-1-n\alpha_1}{n-1-\alpha_1} \tag{8}$$

with $\alpha_1 := \left(\frac{n-1}{n}\right)^n$.

A general expression for the optimum transmission probability for slot $i = s-k$ with $k \in \{0, \ldots, s-1\}$ is not trivial to derive. Taking the derivates of (1) for higher values of s, setting them to zero, and analyzing the structure of the results gives us the following solution.

Theorem 1: Each of n nodes transmits with probability p_i in time slot $i \in \{1, \ldots, s\}$. The probability that there occurs a non-colliding first message within s slots is maximal if the p_i-values are set to

$$p_{s-k} = \frac{1}{n} \cdot \frac{(n-1)^k - n\,\alpha_k \beta_k}{(n-1)^k - \alpha_k \beta_k} \tag{9}$$

for all $k \in \{0, \ldots, s-1\}$ with $\alpha_k := \left(\frac{(n-1)^k}{n}\right)^n$ and the recursively defined term

$$\beta_k := \begin{cases} 0 & \text{for } k = 0 \\ \left((n-1)^{k-1} - \alpha_{k-1}\beta_{k-1}\right)^{1-n} & \text{else} \end{cases} . \tag{10}$$

This set of equations enables us to calculate all transmission probabilities that maximize Φ for given n and s.

Example: Consider $n = 5$ nodes on a channel with $s = 10$ slots. If we choose the same transmission probability p for each slot, the maximum achievable Φ is $84.05\,\%$, which is obtained for $p = 6.2\,\%$. Optimizing the transmission probability of each slot, i.e., choosing

$$\begin{aligned} \{p_1, \ldots, p_{10}\} \quad = \quad & \{3.51\,\%, 3.86\,\%, 4.28\,\%, 4.80\,\%, 5.48\,\%, \\ & 6.38\,\%, 7.65\,\%, 9.57\,\%, 12.86\,\%, 20\,\%\} \end{aligned}$$

increases Φ to $86.68\,\%$.

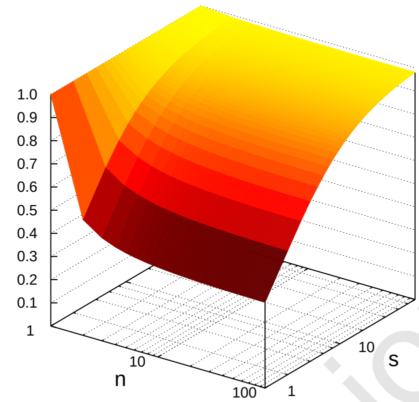

(a) $\max \Phi(n, s, p_1, \ldots, p_s)$

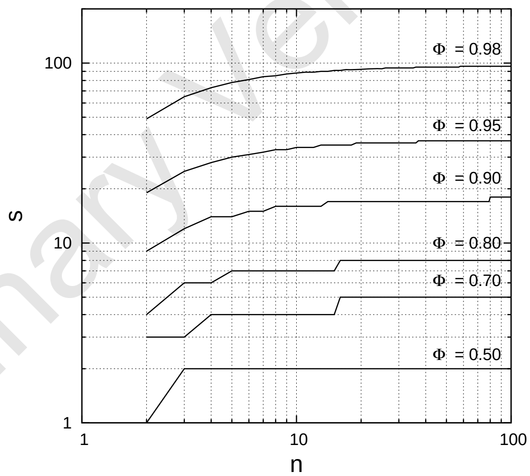

(b) Contour lines of $\max \Phi(n, s, p_1, \ldots, p_s)$

Fig. 4. On a channel with n nodes and s time slots, each node transmits with probability p_i in slot i, $i = 1, \ldots, s$, where p_i is chosen to maximize the probability $\Phi(n, s, p_1, \ldots, p_s)$ of a non-colliding first message.

We observe that the optimal transmission probabilities correspond to a "slow start strategy": All nodes transmit with a very low probability in the first slot, they gradually increase this transmission probability until they transmit with probability $1/n$ in the last slot.

Figure 4(a) shows the maximum possible probability of a non-colliding first message in the n/s-plane, obtained by setting the p_i values according to (9). Figure 4(b) shows corresponding contour lines.

Comparing these results with Figure 3(b) shows that a slow start strategy, optimizing each p_i individually, leads to a (slightly) higher Φ for all n and s than the the slot-independent strategy, where a transmission probability p is chosen collectively for all slots that maximizes Φ. Equivalently, for given n, fewer slots are needed to achieve the same Φ.

IV. Delay of a Non-Colliding First Message

As observed in the previous section, the probability that there occurs a non-colliding first message can be increased by increasing the number of slots and/or using a slow start strategy. Both measures do not however come for free but have to be traded off against a longer delay until the first message actually occurs. This section thus analyzes some stochastic properties of the slot number of the first message.

Representing this slot number by the random variable D ("delay"), we are interested in two metrics: What is the expected delay $\mathbb{E}[D]$ of the first message? What is the maximum delay that can be guaranteed in $90\,\%$ of all cases?

Using these metrics, we compare the transmission strategies from the previous section: (a) the slot-independent optimization strategy of Section III-A4 and (b) the slot-dependent slow start strategy of Section III-B.

A. *Expected Delay*

The ith slot is the first non-empty slot, if at least one message is transmitted in slot i ($M_i > 0$) and no messages have been transmitted in all previous slots ($M_w = 0$ for $w = 1, \ldots, i-1$). The probability for this event is

$$\mathbb{P}[D = i] = \left(1 - (1-p_i)^n\right) \prod_{w=0}^{i-1} (1-p_w)^n \qquad (11)$$

with $p_0 := 0$. The expected delay is then defined by $\mathbb{E}[D] = \sum_{i=1}^{s} i\,\mathbb{P}[D = i]$ under the condition that at least one message is actually transmitted.

Let us analyze $\mathbb{E}[D]$ for the two transmission strategies, i.e., for each (n, s)-pair we always choose (a) the p-value or (b) the p_i-values, respectively, that maximize Φ. Figure 5 shows the normalized resulting expected delay $\mathbb{E}[D]/s$ as a function of s for $n = 5, 10$, and 100 nodes. For the slow start strategy also the expected delay curve for infinitely many nodes is plotted.

We observe the following behavior:

- The expected delay is almost independent of n. For the slow start strategy, it can be expressed as $\mathbb{E}[D] = \mu(s) + \epsilon(n, s)$, where $\mu(s)$ is the limiting value of the expected delay as $n \to \infty$, and $\epsilon(n, s) \in \mathbb{R}$ is a difference that is positive for $s \leq 3$ and negative for $s > 3$ and that tends (quickly) to zero as n becomes large. The difference between the curves with 100 and infinitely many nodes is invisible. It can be shown analytically that the limiting value exists. For the slot-independent strategy, the same general behavior can be observed. As here the optimal probabilities are computed numerically, the results are however no convergence proof for $n \to \infty$.
- The expected delay of the slow start strategy is always higher than that of the slot-independent strategy. For example, for $s = 10$ slots, the first message is expected to occur in the 3rd slot for the slot-independent strategy, while it is expected to occur at least one slot later using the slow start strategy.

For the slot-independent strategy, we can again consider the case of sufficiently many slots. If $s \to \infty$, the expected delay is

$$\mathbb{E}[D] = \frac{1}{1-(1-p)^n}\,. \qquad (12)$$

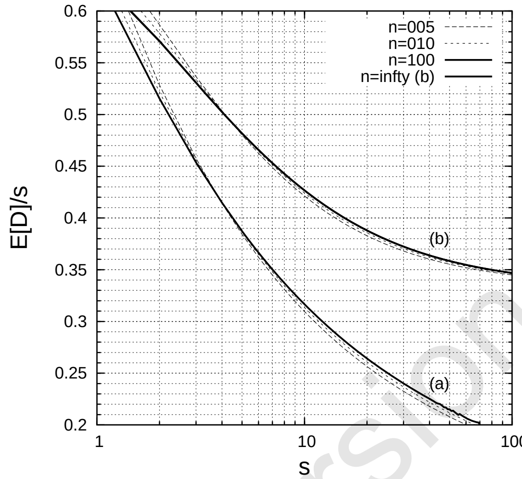

Fig. 5. Expected delay of a first message on a channel with n nodes and s slots. Using the slot-independent strategy (a), each node transmits with probability p in each slot. Using the slow start strategy (b), each node transmits with probability p_i in slot i, where $i = 1, \ldots, s$. The values p or p_i are chosen to maximize Φ.

If $n \to \infty$ and $p \to 0$ while $\lambda = np = \text{const}$, we get

$$\mathbb{P}[D \leq i] = 1 - \mathrm{e}^{-i\lambda} \quad \text{and} \quad \mathbb{E}[D] = \frac{1}{1-\mathrm{e}^{-\lambda}}\,. \qquad (13)$$

B. *Probabilistic Maximum Delay*

We now ask: Requiring that the maximum delay of a first message is k, with a probability of at least $90\,\%$, how small can we choose k? In other words, what is $\inf\{k \in \mathbb{N} : \mathbb{P}[D \leq k] \geq 0.9\}$? Again, we analyze this question in the context of optimizing Φ. This problem corresponds to a system design question in which a certain delay constraint must be fulfilled, and the optimization goal is to maximize Φ. The special case $k = s$ requires that at at least one message is sent with a probability of at least $90\,\%$.

If the same p is used in each slot, the probability that a first message occurs no later than in the ith slot ($i = 1, \ldots, s$) is

$$\mathbb{P}[D \leq i] = \sum_{w=1}^{i} \mathbb{P}[D = w] = 1 - (1-p)^{ni}\,. \qquad (14)$$

Figure 6 shows k/s as a function of s for $n = 5, 10$, and 100 nodes. We observe a similar qualitative behavior as with the expected delay:

- the maximum delay k is almost independent of n, and
- the slow start strategy has a higher or equal maximal delay k as the slot-independent strategy.

Furthermore, for low s (say $s \leq 5$), the maximum delay is not lower than the total number of slots, i.e., $k/s = 100\,\%$. If we increase s, the maximum delay can also be decreased relatively. The zigzag behavior is a consequence of the fact that a certain value of k usually holds for more than one value of s; then however, for some s, the value of k has to be increased.

Example: On a channel with $s = 20$ slots and $n = 10$ nodes, a first message occurs with high probability (in at least

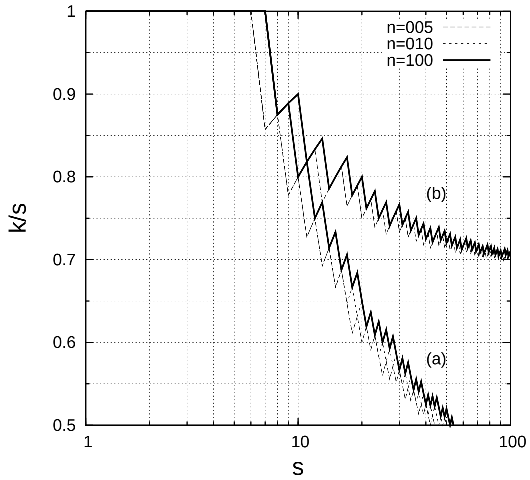

Fig. 6. A first message occurs within the first k time slots with probability $90\,\%$. The transmission strategies intend to maximize Φ according to Figure 5.

$90\,\%$ of all cases) within the first 15 slots using the slow start strategy, and it occurs within the first 13 slots using the slot-independent strategy.

V. Related Work

The paper [3] also addresses the collision probability of first messages assuming a node selection protocol in which the node answering first is selected. The selection process works, however, differently in that paper. While we assume here that all nodes fulfilling a certain *threshold value* with respect to a selection metric (e.g., battery level, link quality to destination) have *equal* chances to access the channel, the goal in [3] is to select the node that has the *best value* with respect to a selection metric. Based on this, the paper [3] tries to find a good mapping function between the metric and a backoff time, this time being in general different for each node.

The paper [5] also derives a transmission strategy with the goal to maximize the probability of a successful first slot access in a channel with a finite number of slots. The modeling assumptions are however different, as [5] requires $\sum_{i=1}^{s} p_i = 1$. This constraint influences the transmission strategy and the resulting performance. Although a slow start strategy similar to the one presented here is applied, the p_i-values of [5] are higher; especially, the transmission probability in the last slot, p_s, is much higher. Our optimization according to (9) can yield a better success probability than the one in [5]. This holds e.g. for all parameters of Table I in [5], the main reason for this performance difference being a high collision probability in the last slot using [5].

The authors of [6] derive transmission probabilities for a slotted channel, also suggesting a slow start mechanism. Instead of maximizing the probability of obtaining a non-colliding first message, it maximizes the probability of obtaining at the end of s slots one selected node (“survivor”) from a set of n nodes, and uses the result for medium access contention resolution. Thus, in [6] collisions do not harm as long as the contention is resolved after s slots.

General work on the analysis of slotted ALOHA from a MAC layer perspective can be found, for example, in the landmark papers [4] and [7] as well as in [8] and [9]. These papers mainly focus on the overall system performance in terms of throughput and delay. No distinction is made between first messages and other messages.

VI. Conclusions and Outlook

Analyzing message collisions in time-slotted random access, we presented a slow start strategy for maximizing the probability of a non-colliding first message and investigated the involved collision-delay tradeoff.

The results are relevant to node selection problems, such as relay and sink selection in wireless systems. They can also be applied to advanced random access, for instance, when s slots are used as a contention window, and the node that starts transmitting first can continue transmitting beyond the s slots.

From this theoretical contribution, various protocol-related issues arise. Further work is needed on techniques to estimate the number of nodes n and to distribute the transmission probabilities p_i in an efficient manner.

Acknowledgments

Christian Bettstetter was inspired by a discussion with Mischa Dohler. Günther Brandner has been funded in part by the European Regional Development Fund and the Carinthian Economic Promotion Fund (KWF) under grant 20214/15935/23108. The work of Robert Vilzmann was performed in part during a visit at the University of Klagenfurt. The authors would like to thank Udo Schilcher for discussions. One of the reviewers brought to our attention the related work [5].

References

[1] J. N. Laneman, D. N. C. Tse, and G. W. Wornell, “Cooperative diversity in wireless networks: Efficient protocols and outage behavior,” *IEEE Trans. Inform. Theory*, vol. 50, pp. 3062–3080, Dec. 2004.

[2] J. Chou, D. Petrovic, and K. Ramachandran, “A distributed and adaptive signal processing approach to reducing energy consumption in sensor networks,” in *Proc. IEEE Infocom*, (San Francisco, CA), Mar. 2003.

[3] T. Watteyne, I. Augé-Blum, M. Dohler, and D. Barthel, “Reducing collision probability in wireless sensor network backoff-based election mechanisms,” in *Proc. IEEE GLOBECOM*, (Washington, DC), Nov. 2007.

[4] N. Abramson, “The throughput of packet broadcasting channels,” *IEEE Trans. Commun.*, vol. 25, no. 1, pp. 117–128, 1977.

[5] Y. Tay, K. Jamieson, and H. Balakrishnan, “Collision-minimizing CSMA and its applications to wireless sensor networks,” *IEEE J. Select. Areas Commun.*, vol. 22, pp. 1048–1057, Aug. 2004.

[6] J. A. Stine, G. de Veciana, K. H. Grace, and R. D. Durst, “Orchestrating spatial reuse in wireless ad hoc networks using synchronous collision resolution,” *J. Interconnection Networks*, vol. 3, pp. 167–198, Sept. 2002.

[7] L. Kleinrock and F. Tobagi, “Packet switching in radio channels: Part I–Carrier sense multiple-access modes and their throughput-delay characteristics,” *IEEE Trans. Commun.*, vol. 23, no. 12, pp. 1400–1416, 1975.

[8] T. Saadawi and A. Ephremides, “Analysis, stability, and optimization of slotted ALOHA with a finite number of buffered users,” *IEEE Trans. Automat. Contr.*, vol. 26, no. 3, pp. 680–689, 1981.

[9] Y. Jenq, “On the stability of slotted ALOHA systems,” *IEEE Trans. Commun.*, vol. 28, no. 11, pp. 1936–1939, 1980.