Wireless Networks 11, 523–529, 2005
© 2005 Springer Science + Business Media, Inc. Manufactured in The Netherlands.

Increasing Deadline-Constrained Throughput in Multi-Channel ALOHA Networks via Non-Stationary Multiple-Power-Level Transmission Policies*

YITZHAK BIRK and YORAM REVAH†
Technion–Israel Institute of Technology, Haifa 32000, Israel

Abstract. Multi-channel Slotted ALOHA is currently used primarily in satellite-based networks for transaction processing (e.g., credit card payments at cash registers). For these applications, maximization of attainable throughput while adhering to a maximum-delay constraint with a required probability reflects both the user's requirements, the network owner's desires, and the non-deterministic nature of ALOHA. This paper explores the judicious use of multiple power levels as a priority mechanism; e.g., the last transmission attempt uses higher power. It focuses on the practical and relevant range of three transmission attempts, up to three power levels, and maximum values of the permissible probability of missing the deadline (P_e) in the range $10^{-5} < P_e < 10^{-2}$. Our scheme increases the attainable delay-constrained throughput by 84–355% (two power levels) and 140–762% (three levels) over classical slotted ALOHA. An optimized combination of multiple copies and two power levels outperforms classical slotted ALOHA by 144–1240%. The smaller P_e, the greater the improvement. The benefit of our schemes is thus dramatic, and far exceeds the contribution of power capture to (unconstrained) capacity of ALOHA.

Keywords: multi-channel ALOHA, delay-constrained throughput, multiple-access schemes, quality of service

1. Introduction

The ALOHA random access scheme was introduced by Abramson in the 1970s [1], and a slotted-time version was introduced in [12]. Since then, it has been studied extensively. The majority of the studies suggested improvements to the basic scheme in order to increase the attainable throughput or improve the mean delay for a given throughput. Other studies considered stability problems (e.g., [6,9]).

Presently, ALOHA is used almost exclusively for transmitting short messages in networks that utilize a shared channel when propagation delay is larger than message-transmission time. (The messages may carry user information, serve for network control, or serve for reservation requests.) In such an environment, short messages render reservation schemes ineffective, and the long round-trip delay precludes the effective use of channel-sensing access schemes.

One of the important uses of ALOHA nowadays is in satellite-based networks used for short transactions [5]. For these uses, [4] introduced a performance measure that reflects both the user's requirements, the network owner's desire to use the network efficiently, and the probabilistic nature of ALOHA: maximizing attainable throughput ("capacity") while adhering to a maximum permissible delay (deadline) constraint with (at least) a required probability.

*This work was supported in part by the Information Superhighway In Space (ISIS) consortium, administered by the office of the Chief Scientist of the Israeli ministry of industry and trade.
†Presently a PhD candidate at the Ben-Gurion University of the Negev, Israel.

The straightforward use of multi-channel Slotted ALOHA for this purpose entails transmission of a packet over a randomly chosen channel. If an acknowledgment is not received from the hub, the packet is retransmitted over a randomly chosen channel. These *transmission rounds* are repeated until receipt of an ACK or expiration of the deadline. We refer to this as the *baseline scheme* or as "standard multi-channel slotted ALOHA".

A general alternative approach for maximizing delay-constrained capacity was introduced in [4]: transmission policies whereby the maximum channel-resource expenditure per message is high, while the mean is kept low. The rationale is as follows:

- Spending a **large maximum** effort on a message before giving up on it reduces the probability of its failure to meet the deadline;
- The **low mean** resource expenditure minimizes "pollution", thereby allowing more active users (and thus higher throughput) at any given "working point" (offered load).

A class of policies that applies this approach entails spending a small amount of network resources on a message in the first transmission round, and increasing this amount in late rounds. Because late rounds are unlikely to take place (transmission ceases upon successful reception), a large amount of network resources can be spent on a message before it is abandoned, with little impact on the mean per-message resource expenditure [4].

In [4], an optimal *multi-copy* policy was developed: an increasing number of copies of the message are transmitted (over different, randomly chosen channels) in successive transmission rounds until success or deadline expiration. (The number of copies in each round is optimized.)

In [3], a *"multiple working points"* scheme was proposed: disjoint subsets of channels are allocated to different transmission rounds, with less loaded channels allocated to the later rounds. (The partitioning of the channels among the groups is optimized.)

The "multi-copy" and the *"multiple working points"* schemes achieve dramatic improvements relative to the baseline scheme (one copy over a randomly chosen channel in each transmission attempt). It was also discovered that multi-copy is superior, and an optimal combination of the two approaches only slightly outperforms it [3].

In [2], the multi-copy approach of [4] was generalized and extended to multi-slot packets: erasure-correcting codes computed over multiple same-packet fragments replaced replication; this allows any given "extra" (redundant) transmitted fragment to be substituted for **any** single fragment of the original packet fragments that was not received. The successful reception of any message fragment moreover made the hub aware of the need to send the remaining ones, and the hub allocated contention-free slots for them. This *"coding-reservation"* scheme can break the well-known $1/e$ capacity barrier of Slotted ALOHA even with a delay constraint.

The use of multiple power levels, selected randomly, has been shown to increase the capacity of various multiple-access schemes, in particular ALOHA [7,8,10,14]. The improvement is due to the power capture effect, which, in certain cases, permits successful reception of one packet despite the concurrent transmission of others on the same channel. Capacity (no delay constraint) increases of 43% and 70% were achieved with two and three power levels, respectively, and perfect capture [10].

In the current paper, we explore the benefits of an explicit, deterministic optimized use of multiple power levels as a priority mechanism in order to increase delay-constrained capacity. Specifically, transmission using a higher power level in the last round(s) is used to prefer the late-round transmissions of a message that is about to be dropped over early attempts of other messages. The method is developed and analyzed for "narrowband" channels and "conventional" receivers, which permit at most one successful reception per channel per time slot. We then jointly optimize the use of multiple power levels and multiple copies to obtain further performance improvements. We intentionally consider only a small number of power levels, a short permissible maximum delay, and simple hardware, so as to avoid any doubt as to the practicality of our schemes.

The remainder of the paper is organized as follows. In Section 2, we present the network model. In Section 3 we develop and analyze specific transmission policies using two and three power levels, focusing on the latter for brevity. In Section 4 we optimally combine multiple copies with two power levels. Section 5 offers concluding remarks.

2. Network model and preliminaries

2.1. Model and definitions

The network comprises ground stations that transmit single-slot messages over randomly chosen time-slotted channels. A hub monitors all channels and ACKs all successful receptions. ACKs are sent over separate, contention-free lossless channels. The lack of an ACK when it is expected indicates a collision.

We assume an infinite number of stations and a large number of channels. The number of transmissions over any given channel in any given time slot is modeled as a Poisson random variable, independent from slot to slot and from channel to channel. The analysis in this paper is carried out under such an independence assumption. The independence is between the fates of different transmissions of the same message. Strictly speaking, the fates of retransmissions depend on those of past transmissions, but this dependence diminishes as the number of channels is increased. This fact was confirmed by simulation [4]. Thus, the assumption of independence is a good approximation.

We employ a realistic first-order model of power capture, whereby a packet is captured by a receiver iff the ratio of its received power to the sum of the total received power of all other packets simultaneously received by the receiver plus receiver noise (S/(I + N)) is greater than a given capture ratio, denoted β. We also assume the use of "narrowband" transmissions (no coding or spreading gain), so $\beta > 1$. This, in turn, implies that at most one message can succeed on any given channel in any given time slot. In the analysis, we will assume that the noise is negligible relative to the received power of a packet. This is sensible because, in satellite communications, the power is set such that packets are receivable despite the noise, with proper link margins to guarantee this in changing weather conditions. We will thus focus on signal-to-interference ratio. The analysis can easily be extended to the case in which noise is not negligible.

A message that is not received by the deadline is dropped by the transmitting node. By a slight abuse of notation, we use P_e to denote both the failure-probability constraint (maximum permissible probability of not meeting the deadline) and the actual failure probability at a given working point with a given transmission policy. The intent should be obvious to the reader in each instance. (Note that, when operating at the maximum attainable throughput for a given scheme, the values of P_e in the two senses become equal.) We also adopt the distinction made in [4] between the generation rate of new messages, S_g, and the throughput S. Specifically, $S = S_g(1 - P_e)$. (In this work, we consider maximum permissible probabilities of missing the deadline and thus dropping a message in the range $10^{-5} \leq P_e \leq 10^{-2}$, so this distinction is mostly a formality.)

Remark. We speak of transmission power levels. In the analysis, however, these should be understood as the resulting received-power levels. Also, all the analysis is carried out for a single channel. The existence of many channels is only reflected in the independence among fates of any given packet in different rounds.

Stability. Multi-channel ALOHA with message discarding upon deadline expiration cannot "crash" due to overload, because the offered load never exceeds the product of the arrival rate and the maximum number of copies of a message that are transmitted before it is discarded due to deadline expiration. Nonetheless, it is bistable in certain load regions, with a lower bound on $S(S_g)$. The hub can detect such situations and "push" the network into the "good" stable point, namely one in which increasing G increases S. The hub can, for example, use the contention-free outbound channel to instruct stations to back off probabilistically. (Such a scheme does not require the hub to know the identities of the contending ground stations.) The analysis in this paper applies to "good" stable operation.

3. Maximum throughput with multiple power levels

In this section, we develop and analyze access schemes that entail the transmission of a single copy per round. The number of rounds is determined by the deadline and by the round time (round-trip propagation delay plus any processing time in the hub or in the satellite terminals). We have elected to focus on the case of three rounds, corresponding to a deadline of 2-3 sec in practical systems using geosynchronous satellites. The extension to a different number of rounds is obvious. Similarly, we only consider the use of two and three different power levels, so that the schemes are practical and relatively simple to implement. Once again, the schemes and their analysis can be extended with relative ease.

"Stationary" multi-channel ALOHA policies that use multiple power levels have been studied in the past. With these, power for each transmission is chosen independently with optimal probabilities from among a set of optimally-chosen power levels [10]. However, our interest is in attainable delay-constrained throughput rather than in capacity. We therefore begin by re-optimizing these policies for our performance measure. Next, we replace the "stationary" policies with "non-stationary" ones, whereby the power is selected optimally (and deterministically) per transmission round. (We refer to the latter as "deterministic" in view of the deterministic selection of transmission power, which is decided off-line for each round. Channels are still selected randomly per transmission.) Results are presented for both two and three power levels. For the sake of brevity, however, we omit the analysis for the case of two levels.

3.1. Optimized "stationary" policy

Let P_L, P_M and P_H denote the probability of transmitting with low, medium and high power, respectively. For facility of exposition and because we use these schemes primarily for comparison with our non-stationary schemes, we assume perfect power capture: a transmission succeeds if and only if there are no other transmissions with equal or higher power on the same channel. We also assume that receiver noise is negligible. The probability of success of a transmission is therefore

Table 1
Attainable throughput with stationary policies; three rounds.

$S_3^S - S_1$ (%)	$S_2^S - S_1$ (%)	S_3^S	S_2^S	S_1	P_e
53	35	0.2917	0.2574	0.1904	10^{-2}
51	34	0.1435	0.1271	0.0948	10^{-3}
51	34	0.0682	0.0605	0.0453	10^{-4}
50	33	0.0320	0.0284	0.0213	10^{-5}

$$P_{\text{suc}} = P_L e^{-G} + P_M e^{-G(1-P_L)} + P_H e^{-G(1-P_L-P_M)}, \quad (1)$$

where G is the offered load (per channel). The probability of failing to meet the deadline (three rounds) is

$$P_e = (1 - P_{\text{suc}})^3. \quad (2)$$

The mean total number of transmitted copies per packet, including both successful and dropped messages, is

$$E = P_{\text{suc}} + 2(1 - P_{\text{suc}})P_{\text{suc}} + 3(1 - P_{\text{suc}})^2, \quad (3)$$

and throughput is given by

$$S = \frac{G}{E}(1 - P_e) = S_g(1 - P_e). \quad (4)$$

An optimal choice of probabilities for using the three power levels yields nearly identical values for P_L, P_M and P_H. (Note that, with perfect capture, the relative power levels do not matter.)

Table 1 depicts the attainable throughputs. Subscripts denote the number of power levels, and the superscript "S" stands for "stationary".

We see that the optimized stationary policy improves performance relative to the baseline scheme by approximately 33% and 50% with two and three power levels, respectively. Results with realistic, imperfect power capture are very similar. This improvement is actually smaller than the improvement in unconstrained capacity (no deadlines), which is 44% and 70% for two and three power levels, respectively [10].

3.2. "Deterministic, non-stationary" policies

A comparison of the attainable throughputs of all possible deterministic choices from among three power levels in each of three rounds (with relative power levels optimized for each choice of power-level sequence) reveals that the optimal choice of power levels for the three rounds with two power levels is LHH, and with three power levels it is LMH (Low, Medium and High power in the first, second and third rounds, respectively). Also, probabilistic non-stationary policies do not outperform the deterministic ones. We next present the analysis for the LMH policy, including the determination of the optimal relative power levels. Analysis for other policies is similar.

3.2.1. A practical model for power capture

Suppose that, in a given time slot, n messages are transmitted over a given channel with power levels W_i $(i = 1, 2, 3, \ldots, n)$. Without loss of generality, let us consider message n. Per the assumptions stated in the previous section, message n will be received successfully if and only if

$$\frac{W_n}{\sum_{i=1}^{n-1} W_i} \geq \beta > 1. \qquad (5)$$

A message can thus succeed only if no other messages are transmitted over the same channel in the same time slot with equal or higher power, and not too many are transmitted with lower power.

Let us next determine the three power levels. The Low power level is determined by the receiver's sensitivity, denoted $W_{\min}$, and is simply $W_L = W_{\min}$. The Medium power level is determined so as to enable the receiver to capture a Medium-power message in the presence of an integer number of Low-power messages. Therefore, $W_M = b \cdot \beta \cdot W_{\min}$, where b is some integer. Following the same logic, the High power level is determined so that the receiver is able to capture a message of High power level when it is transmitted along with one Medium-power message and possibly some combination of additional Low- and Medium-power messages. We therefore let $W_H = l \cdot \beta^2 \cdot W_{\min}$, with $l \geq b$. The meaningful increments of l are $1/\beta$, corresponding to the ability of a High-power message to tolerate an additional Low-power transmission.

It can readily be seen that increasing l increases the probability of success of a High-power transmission without altering those of lower-power ones. This suggests choosing infinite l. In practice, however, maximum transmission power is dictated by the maximum possible transmission power or by regulatory constraints (rounded down to the nearest meaningful value). Given β and $W_{\min}$, this determines l. It then remains to find $1 \leq b \leq l$ that achieves maximum throughput according to our measure. (b determines the Medium power level.)

After determining the power levels, we will compute the probability of a message to succeed when transmitted at the Low power level W_L.

Due to the independence assumption, the offered load G can be expressed as the sum of the means of three independent Poisson random variables. With one copy transmitted in each attempt (without regard to transmitted power), it follows that

$$G = S_g + P_{c_1} S_g + P_{c_1} P_{c_2} S_g, \qquad (6)$$

where P_{ci} denotes that probability of failure ("collision") in the ith attempt.

Since a low-power message succeeds only if it is the only one transmitted in the slot, the probability of success of a message in the first attempt is:

$$P_{\text{suc}_1} = e^{-G}. \qquad (7)$$

The probability of success of a message transmitted with medium power W_M in the second attempt is:

$$P_{\text{suc}_2} = e^{-P_{c_1} S_g} e^{-P_{c_1} P_{c_2} S_g} \left(\sum_{k=0}^{b} (S_g)^k \frac{e^{-S_g}}{k!} \right). \qquad (8)$$

$e^{-P_{c_1} S_g}$ is the probability that Medium power messages will not appear in the same time slot (except the test message for which we are computing the probability of success).

$e^{-P_{c_1} P_{c_2} S_g}$ is the probability that no High-power message will appear in the same time slot.

$\sum_{k=0}^{b} (S_g)^k \frac{e^{-S_g}}{k!}$ is the probability to transmit no more than b Low-power messages in the same time slot.

After further algebraic manipulation, the probability of success for a message transmitted with Medium power level in the second attempt is:

$$P_{\text{suc}_2} = e^{-G} \left(\sum_{i=0}^{b} \frac{S_g^i}{i!} \right). \qquad (9)$$

In the third attempt (High power level), two conditions must be met in order for the message to succeed: (1) there are no other High-power messages, and (2) the total power of the other transmitted messages in the channel has to adhere to the following constraint:

$$N_M b \beta W_{\min} + N_L W_{\min} \leq \frac{W_H}{\beta} = l \beta W_{\min}. \qquad (10)$$

N_M and N_L are the number of transmitted messages with Medium and Low power, respectively.

Substituting (10) in (5) yields:

$$\frac{l \beta^2 W_{\min}}{N_M b \beta W_{\min} + N_L W_{\min}} \geq \beta. \qquad (11)$$

The probability of success of a message transmitted with High power in the third attempt is a sum of all the allowed events that fulfill inequality (10). So, the probability of success is:

$$P_{\text{suc}_3} = e^{-G} \left\{ \sum_{j=0}^{\lfloor (l - \lfloor \frac{l}{b} \rfloor b) \beta \rfloor} \left(\frac{(P_{c_1} S_g)^j}{j!} \left[\sum_{i=0}^{\lfloor (l - jb) \beta \rfloor} \frac{S_g^i}{i!} \right] \right) \right\} \qquad (12)$$

The probability of failure for a message in attempt i is:

$$P_{c_i} = 1 - P_{\text{suc}_i} \qquad (13)$$

Due to the independence assumption, the total probability for a message to fail in all three transmission attempts is:

$$P_e = P_{c_1} P_{c_2} P_{c_3} \qquad (14)$$

Let E denote the mean total number of copies of a message transmitted until success or failure. Then,

$$G = S_g E \qquad (15)$$

and channel throughput is:

$$S = \frac{G}{E}(1 - P_e) = S_g(1 - P_e) \qquad (16)$$

3.3. Results

The results were computed by solving (16) numerically. Given the values of l, β and P_e, we optimize b in order to achieve

Table 2
Attainable delay-constrained throughput ($D_R = 3$).

$S_3 - S_2$ (%)	$S_3 - S_3^S$ (%)	$S_3 - S_1$ (%)	$S_2 - S_2^S$ (%)	$S_2 - S_1$ (%)	S_3	S_2	S_1	P_e
30	67	140	38	85	0.456	0.354	0.191	10^{-2}
45	137	259	84	146	0.341	0.234	0.095	10^{-3}
65	269	457	150	237	0.251	0.152	0.045	10^{-4}
90	475	776	246	362	0.184	0.097	0.021	10^{-5}

maximum throughput. We computed results for $\beta = 2, 5, 10$ and 40.

For every β tested, throughput at $l = 4$ is already, is almost identical to that with perfect capture (where a packet succeeds in the presence of any number of lower-power packets). Also, $b_{\text{opt}} = 2$ across the range. We therefore only present one set of results.

Table 2 depicts the numerical results for attainable throughput. Subscripts denote the number of power levels, and a superscript "S" denotes an optimized probabilistic stationary policy. S_1 thus refers to the baseline scheme. S_2 was obtained with an LHH power setting for the different rounds, and S_3 was obtained with LMH.

Clearly, the increase in attainable throughput with our optimized deterministic, non-stationary multi-power-level policies, even relative to optimized probabilistic stationary ones, is dramatic. Also, the contribution of a third power level is significant. Finally, the improvement is more pronounced for lower permissable P_e (i.e., stricter requirements).

4. An optimized multicopy scheme with power capture

In this section we combine the power capture scheme with the multi-copy scheme and analyze its performance using our measure. We confine ourselves to policies that are practical in terms of implementation. In order to reduce complexity, we only consider two power levels.

In [4], dynamic programming was used in order to find the number of copies that maximizes throughput under delay and error-probability constraints. Unfortunately, we could not find an efficient algorithm that enables the use of dynamic programming. Instead, we chose to tackle the probable using an intelligent search based on the insight that most effort should best be applied in later attempts. Specifically, we only investigated non-stationary policies that adhere to the following constraints: (1) one, low-power copy is transmitted in the first round, and (2) at least one high-power copy is transmitted in the third attempt. Having tested some 50 policies, we next present the best results (for our performance measure).

4.1. Analysis

The analysis is carried out under the assumption of perfect capture because, in a similar fashion to the previous section, the practical model converges to perfect capture already at $b = 3$.

Let N_i^L and N_i^H denote the number Low- and High-power copies of a given message transmitted in the ith attempt. A three-round, two-level, multi-copy transmission policy is specified as follows: $(N_1^L, N_1^H; N_2^L, N_2^H; N_3^L, N_3^H)$. Under the independence assumption, the mean number of Low-power transmissions is:

$$G_L = S_g\left(N_1^L + P_{c_1}N_2^L + P_{c_1}P_{c_2}N_3^L\right) \tag{17}$$

The mean number of High-power transmissions is:

$$G_H = S_g\left(N_1^H + P_{c_1}N_2^H + P_{c_1}P_{c_2}N_3^H\right) \tag{18}$$

(In the set of policies that we considered, $N_1^L = 1$ and $N_1^H = 0$).

The overall offered load on a channel (ignoring power levels) is:

$$G = G_L + G_H \tag{19}$$

The probability of failure for a message transmitted with low power is:

$$P_{C_L} = (1 - e^{-G}) \tag{20}$$

The probability of failure for a message transmitted with high power is:

$$P_{C_H} = (1 - e^{-G_H}) \tag{21}$$

Under the independence assumption, the probability of failure in transmission attempt $i = 1, 2, 3$ is:

$$P_{c_i} = \left(P_{c_L}\right)^{N_i^L}\left(P_{C_H}\right)^{N_i^H} \tag{22}$$

The probability of a message failing to be received by the deadline is given by (14).

For each P_e that was tested, there is a different optimal transmission policy. For example, for $P_e = 10^{-3}$, the optimal policy is: (1,0; 1,1; 1,3).

Substituting in equations (17)–(22), we obtain:

$$\begin{aligned} G &= S_g + 2P_{c_1}S_g + 4P_{c_1}P_{c_2}S_g \\ P_{c_1} &= (1 - e^{-G}) \\ P_{c_2} &= P_{c_L}P_{c_H} \\ P_{c_3} &= P_{c_L}\left(P_{c_H}\right)^3 \end{aligned} \tag{23}$$

The throughput is computed as in (16).

4.2. Numerical results

The equations were solved numerically and the results are depicted in Table 3. Subscripts denote the number of power levels, and a superscript "mc" denotes multi-copy.

The results confirm that attainable throughput is maximized by schemes that spend a monotonically non-decreasing amount of resources (power and number of copies) on a message as a function of round number. Multiple copies are more effective for lower permissible message-loss probabilities, whereas multiple power levels are more useful for high permissible ones. Finally, the combination of power capture and multiple copies substantially outperforms each

Table 3
Attainable throughput with jointly optimized multiple copies and power levels.

$S_2^{m.c} - S_1$ (%)	$S_2^{m.c} - S_2$ (%)	$S_2^{m.c} - S_1^{m.c}$ (%)	$S_2^{m.c}$	S_2	$S_1^{m.c}$ [6]	P_e
114	15	46	0.407 [1^L,1^H,3^L1^H]	0.354	0.279 [1,2,4]	10^{-2}
270	50	42	0.351 [1^L,1^L1^H,1^L3^H]	0.234	0.247 [2,3,7]	10^{-3}
605	110	37	0.32 [1^L,1^L1^H,4^H]	0.152	0.233 [2,3,10]	10^{-4}
1240	193	–	0.285 [1^L,1^L1^H,4^H]	0.097	–	10^{-5}

of them, probably because the schemes operate on different "dimensions".

5. Conclusions

This paper introduced the use of multiple power levels, alone or in conjunction with the use of multiple copies, to increase the attainable throughput of multi-channel slotted ALOHA subject to a given deadline and a (very small) maximum permissible probability of failing to meet it.

"Stationary" policies, whereby the power of each transmission is drawn from the same (optimized) probability function regardless of transmission round, increase attainable throughput by several tens of percents when compared with the baseline scheme of a single power level and a single copy per round. This increase is somewhat smaller than the known increase in unconstrained capacity with such (albeit differently optimized) schemes.

Our optimized deterministic, static non-stationary policies, whereby the power level is chosen based on the round number (but not on dynamic network conditions), are much more effective: they increase attainable throughput by 84–355% (with two power levels) and 140–762% (with three power levels) relative to the baseline scheme. These improvements were obtained for maximum permissible probabilities of failure to meet the deadline in the range $10^{-5} \leq P_e \leq 10^{-2}$, with greater improvement for smaller P_e.

The optimal hybrid policy, combining multiple power levels and multiple copies, outperforms the baseline scheme by 144–1240% when $10^{-5} \leq P_e \leq 10^{-2}$ using only two power levels, an improvement of 37–46% over the optimal multi-copy policy [4]. This is in stark contrast with the combination of multiple copies and multiple working points [3], which improves performance relative to multi-copy by only 1%. A common and interesting result for all the new schemes is that the advantage of the better schemes over the lesser ones increases as P_e is reduced (stricter compliance requirement).

Our focus in this paper has been on "narrowband" channels and "simple" receivers, representing the vast majority of current commercial satellite communication systems. With these, at most one message can be received successfully over any given channel in any given time slot. Therefore, the use of multiple power levels is a pure priority mechanism with no side effects (other than inter-channel interference, which is not a problem in most systems due to careful spectrum shaping). This, however, is no longer the case when one employs spread-spectrum techniques with CDMA. There, multiple messages can be received concurrently because the processing gain permits reception even when the signal-to-interference ratio is smaller than one. Transmitting a message with high power thus increases the interference seen by other messages, and can actually reduce the possible number of concurrently-received messages. For this reason, power-equalization is employed in many such systems. With our performance measure, there is thus a trade-off between the priority benefits offered by the multiple power levels and the harm that they cause. Indeed, we have seen that the overall benefit is small or even negative, depending on parameter values [13].

The use of sophisticated receivers, employing techniques such as multi-user detection [13], enhances the ability to receive weak signals in the presence of strong ones. The overall impact of the use of those on the effectiveness of transmission schemes employing multiple power levels for increasing delay-constrained throughput is unclear, and warrants further research.

In summary, we have introduced simple, practical transmission policies that offer dramatic performance enhancements with current equipment under a measure that is very well matched to the current uses of the ALOHA access scheme. The implications of next-generation receivers warrant further study.

References

[1] N. Abramson, Development of the ALOHANET, IEEE Trans. IT 31(3) (1985) 119–123.

[2] D. Baron and Y. Birk, Coding schemes for multislot messages in multichannel ALOHA with deadlines, IEEE Trans. Wireless Commun. 1(2) (2002) 292–301.

[3] D. Baron and Y. Birk, Multiple working points in multichannel ALOHA with deadlines, ACM/Baltzer Wireless Networks 8(1) (2002) 5–11.

[4] Y. Birk and Y. Keren, Judicious use of redundant transmissions in multichannel ALOHA networks with deadlines, IEEE JSAC 17(2) (1999) 257–269.

[5] R. Heiman and G. Kaplan, Issues in VSAT Networks, Lecture in *The 19th Convention of Electrical and Electronics Engineers in Israel, Jerusalem* (Nov. 5–6 1996).

[6] L. Kleinork and S.S. Lam, Packet switching in a multiaccess broadcast channel: Performance evaluation, IEEE Trans. on Comm. COM-23(4) (1975) 410–423.

[7] R.O. Lamaire, A. Krishna and M. Zorzi, On the randomization of transmitter power levels to increase throughput in multiple access radio systems, Wireless Networks 4(3) (1998) 7 263–277.

[8] C.C. Lee, Random signal levels for channel access in packet broadcast networks, IEEE J. Select. Areas of Comm. SAC-5(6) (1987) 1026–1034.

[9] Y.W. Leung, Generalized Multicopy ALOHA, Electronics Letters 31(2) (1995) 82–83.

[10] J.J. Metzner, On Improving Utilization in ALOHA Networks, IEEE Trans. on Comm. COM-24 (1976) 447–448.

[11] Y. Revah, Utilization of power capture for delay-constrained throughput maximization in multi-channel ALOHA networks, M.Sc. Dissertation, Electrical Engr. Dept, Technion (2002).
[12] L.G. Roberts, ALOHA packets, with and without Slots Capture, Comp. Comm. Review 5 (1975) 28–42.
[13] S. Verdú, *Multiuser Detection* (Cambridge University Press, 1998).
[14] W. Yue, The effect of capture on performance of multichannel slotted ALOHA systems, IEEE Trans. on Comm. 39(6) (1991) 818–822.

Yitzhak Birk has been on the faculty of the Electrical Engineering Department at the Technion since October 1991, and heads its Parallel Systems Laboratory. He received the B.Sc. (cum laude) and M.Sc. degrees from the Technion in 1975 and 1982, respectively, and the Ph.D. degree from Stanford University in 1987, all in electrical engineering. From 1976 to 1981, he was project engineer in the Israel Defense Forces. From 1986 to 1991, he was a Research Staff Member at IBM's Almaden Research Center, where he worked on parallel architectures, computer subsystems and passive fiber-optic interconnection networks. From 1993 to 1997, he also served as a consultant to HP Labs in the areas of storage systems and video servers, and was later involved with several companies.

Dr. Birk's research interests include computer systems and subsystems, as well as communication networks. He is particularly interested in parallel and distributed architectures for information systems, including communication-intensive storage systems (e.g., multimedia servers) and satellite-based systems, with special attention to the true application requirements in each case. The judicious exploitation of redundancy for performance enhancement in these contexts has been the subject of much of his recent work.
E-mail: birk@ee.technion.ac.il

Yoram Revah received the B.Sc. (Cum Laude) and M.Sc. degrees in Electrical Engineering from the Technion, Israel in 1997 and 2001, respectively. He worked as a VLSI Engineer and research assistant in the Electrical Engineering department at the Technion until 2003. He is presently a Ph.D. candidate in the Communication Systems Department at Ben Gurion University of the Negev, Israel, where he holds a Kreitman Foundation Fellowship.
E-mail: revahyo@bgumail.bgu.ac.il