Optimal Protocols for 2-Party Contention Resolution

Dingyu Wang

September 27, 2021

Abstract

Contention Resolution is a fundamental symmetry-breaking problem in which n devices must acquire temporary and exclusive access to some *shared resource*, without the assistance of a mediating authority. For example, the n devices may be sensors that each need to transmit a single packet of data over a broadcast channel. In each time step, devices can (probabilistically) choose to acquire the resource or remain idle; if exactly one device attempts to acquire it, it succeeds, and if two or more devices make an attempt, none succeeds. The complexity of the problem depends heavily on what types of *collision detection* are available. In this paper we consider *acknowledgement-based protocols*, in which devices only learn whether their own attempt succeeded or failed; they receive no other feedback from the environment whatsoever, i.e., whether other devices attempted to acquire the resource, succeeded, or failed.

Nearly all work on the Contention Resolution problem evaluated the performance of algorithms *asymptotically*, as $n \to \infty$. In this work we focus on the simplest case of $n = 2$ devices, but look for *precisely* optimal algorithms. We design provably optimal algorithms under three natural cost metrics: minimizing the expected average of the waiting times (AVG), the expected waiting time until the first device acquires the resource (MIN), and the expected time until the last device acquires the resource (MAX). We first prove that the optimal algorithms for $n = 2$ are *periodic* in a certain sense, and therefore have finite descriptions, then we design optimal algorithms under all three objectives.

AVG. The optimal contention resolution algorithm under the AVG objective has expected cost $\sqrt{3/2} + 3/2 \approx 2.72474$.

MIN. The optimal contention resolution algorithm under the MIN objective has expected cost 2. (This result can be proved in an ad hoc fashion, and may be considered folklore.)

MAX. The optimal contention resolution algorithm under the MAX objective has expected cost $1/\gamma \approx 3.33641$, where $\gamma \approx 0.299723$ is the smallest root of $3x^3 - 12x^2 + 10x - 2$.[1]

1 Introduction

The goal of a contention resolution scheme is to allow multiple devices to eventually obtain exclusive access to some shared resource. In this paper[2] we will use often use the terminology of one particular application, namely, wireless devices that wish to broadcast messages on a multiple-access channel. However, contention resolution schemes are used in a variety of areas [16, 21, 13], not just wireless networking. We consider a model of contention resolution that is distinguished by the following features.

[1] We may also express γ in radical form: $\gamma = -\frac{1}{6}\left(1 - i\sqrt{3}\right)\sqrt[3]{13 + i\sqrt{47}} + \frac{4}{3} - \frac{1+i\sqrt{3}}{\sqrt[3]{13+i\sqrt{47}}}$.

[2] This is the full version of [24].

Discrete Time. Time is partitioned into discrete *slots*. It is the goal of every device to obtain exclusive access to the channel for exactly one slot, after which it no longer participates in the protocol. We assume that all n devices begin at the same time, and therefore agree on slot zero. (Other work considers an infinite-time model in which devices are injected adversarially [8, 10, 3], or according to a Poisson distribution [18, 23] with some constant mean.)

Feedback. At the beginning of each time slot each device can choose to either transmit its message or remain idle. If it chooses to idle, it receives no feedback from the environment; if it chooses to transmit, it receives a signal indicating whether the transmission was successful (all other devices remained idle). ("Full sensing" protocols like [18, 23, 8, 10, 3], in contrast, depend on receiving ternary feedback at each time slot indicating whether there was no transmission, some successful transmission, or a collision.)

Noiseless operation. The system is errorless; there is no environmental noise.

Anonymity. Devices are indistinguishable and run the same algorithm, but can break symmetry by generating (private) random bits.

There are many ways to measure the time-efficiency of contention resolution protocols. In infinite-time models, we want to avoid deadlock [2, 6, 7, 8, 5, 10], minimize the latency of devices in the system, and generally make productive use of a (large) constant fraction of the slots [8, 5, 10]. When all n devices begin at the same time [6, 7], there are still several natural measures of efficiency. In this paper we consider three: minimizing the time until the *first* successful transmission (MIN), the *last* successful transmission time (MAX, a.k.a. the *makespan*), and the *average* transmission time (AVG).

1.1 Prior Work

Classic infinite-time protocols like ALOHA [1] and binary exponential backoff algorithms [15, 14] are simple but suffer from poor worst case performance and eventual deadlock [2, 6, 7], even under *non*-adversarial injection rates, e.g., Poisson injection rates with arbitrary small means. These are *acknowledgement-based* protocols which do not require constant (ternary) channel feedback. One line of work aimed to achieve deadlock-freeness under Poisson arrivals [9, 12, 18, 23], assuming ternary channel feedback. The maximum channel usage rate is known to be between 0.48776 [18, 23] and 0.5874 [17]. A different line of work aimed at achieving deadlock-freeness and constant rate of efficiency under *adversarial* injections and possibly adversarial *jamming*, also assuming ternary feedback. See [8, 10, 3] for robust protocols that can tolerate a jamming adversary. One problem with both of these lines of work is that all devices must monitor the channel constantly (for the ternary silence/success/collision feedback). Bender et al. [5] considered adversarial injection rates and showed that it is possible to achieve a constant efficiency rate while only monitoring/participating in $O(\log(\log^* n))$ time slots. This was later shown to be optimal [11].

When all n devices start at the same time slot (n unknown), we have a pretty good understanding of the AVG, MIN, and MAX objectives. Here there are still variants of the problem, depending on whether the protocol is full-sensing (requiring ternery feedback) or merely acknowledgement-based. Willard [25] and Nakano and Olariu [19] gave full sensing protocols for the MIN objective when n is unknown that takes time $O(\log\log n + \log f^{-1})$ with probability $1 - f$, which is optimal. The *decay* algorithm [4] is an acknowledgement-based protocol for the MIN objective that runs in $O(\log n \log f^{-1})$ time with probability $1 - f$, which is also known to be optimal [20]. When n is unknown, binary exponential backoff achieves optimal $O(n)$ time under the AVG objective, but suboptimal $\Theta(n \log n)$ time under the MAX objective [6, 7]. The *sawtooth* protocol of Bender et al. [6, 7] is optimal $O(n)$ under both AVG and MAX; it is acknowledgement-based.

1.2 New Results

In this paper we consider what seems to be the *simplest non-trivial symmetry breaking problem*, namely, resolving contention among two parties ($n = 2$) via an acknowledgement-based protocol. The *asymptotic* complexity of this problem is not difficult to derive: $O(1)$ time suffices, under any reasonable objective function, and $O(\log f^{-1})$ time suffices with probability $1 - f$. However, our goal is to discover *precisely* optimal algorithms.

We derive the optimal protocols for the AVG, MIN, and MAX objectives, in expectation, which are produced below. The optimal MIN protocol is easy to obtain using *ad hoc* arguments; it has expected cost 2. However, the optimal protocols for AVG and MAX require a more principled, rigorous approach to the problem. We show that the protocol minimizing AVG has expected cost $\sqrt{3/2} + 3/2 \approx 2.72474$, and that the optimal protocol minimizing MAX has expected cost $1/\gamma \approx 3.33641$, where $\gamma \approx 0.299723$ is the unique root of $3x^3 - 12x^2 + 10x - 2$ in the interval $[1/4, 1/3]$.

AVG-Contention Resolution:

Step 1. Transmit with probability $\frac{4-\sqrt{6}}{3} \approx 0.516837$. If successful, halt; if there was a collision, repeat Step 1; otherwise proceed to Step 2.

Step 2. Transmit with probability $\frac{1+\sqrt{6}}{5} \approx 0.689898$. If successful, halt; if there was a collision, go to Step 1; otherwise proceed to Step 3.

Step 3. Transmit with probability 1. If successful, halt; otherwise go to Step 1.

MIN-Contention Resolution:

- In each step, transmit with probability 1/2 until successful.

MAX-Contention Resolution:

Step 1. Transmit with probability $\alpha \approx 0.528837$, where α is the unique root of $x^3+7x^2-21x+9$ in $[0, 1]$. If successful, halt; if there was a collision, repeat Step 1; otherwise proceed to Step 2.

Step 2. Transmit with probability $\beta \approx 0.785997$, where β is the unique root of $4x^3 - 8x^2 + 3$ in $[0, 1]$. If successful, halt; if there was a collision, go to Step 1; otherwise proceed to Step 3.

Step 3. Transmit with probability 1. If successful, halt; otherwise go to Step 1.

One may naturally ask: what is the point of understanding Contention Resolution problems with $n = O(1)$ devices? The most straightforward answer is that in some applications, contention resolution instances between $n = O(1)$ devices are commonplace.[3] However, even if one is only interested in the asymptotic case of $n \to \infty$ devices, understanding how to resolve $n = O(1)$ optimally is essential. For example, the protocols of [9, 12, 18, 23] work by repeatedly isolating subsets of the n' active devices, where n' is Poisson distributed with mean around 1.1, then resolving conflicts within this set (if $n' > 1$) using a near-optimal procedure. The channel usage rate of

[3] For a humorous example, consider the Canadian Standoff problem https://www.cartoonstock.com/cartoonview.asp?catref=CC137954.

these protocols (≈ 0.48776) depends critically on the efficiency of Contention Resolution among n' devices, where $\mathbb{E}[n'] = O(1)$. Moreover, *improving* these algorithms will likely require a much better understanding of $O(1)$-size contention resolution.

Organization. In Section 2 we give a formal definition of the model and state Theorem 1 on the *existence* of an optimal protocol for any reasonable objective function. In Section 3 we prove another structural result on optimal protocols for $n = 2$ devices under the AVG, MIN, and MAX objectives (Theorem 2), and use it to characterize what the optimal protocols for AVG (Theorem 3), MIN (Theorem 4), and MAX (Theorem 5) should look like. Corollary 1 derives that **AVG-Contention Resolution** is the optimal protocol under the AVG objective, and Corollary 2 does the same for **MAX-Contention Resolution** under MAX. The proofs of Theorems 1 and 2 and Corollaries 1 and 2 appear in the Appendix.

2 Problem Formulation

After each time step the channel issues responses to the devices from the set $\mathcal{R} = \{0, 1, 2_+\}$. If the device idles, it always receives 0. If it attempts to transmit, it receives 1 if successful and 2_+ if unsuccessful. A *history* is word over $\mathcal{R}^*$. We use exponents for repetition and $*$ as short for $\mathcal{R}^*$; e.g., the history $0^3 2_+^2$ is short for 0002_+2_+ and $*1*$ is the set of all histories containing a 1. The notation $a \in w$ means that symbol a has at least one occurrence in word w.

Devices choose their action (transmit or idle) at time step $t \in \mathbb{N}$ and receive feedback at time $t + 0.5$. A *policy* is a function f for deciding the probability of transmitting. Define $\mathcal{F} = \{f : \mathcal{R}^* \to [0, 1] \mid \forall w \in \mathcal{R}^*, 1 \in w \implies f(w) = 0\}$ to be the set of all proper policies, i.e., once a device is successful ($1 \in w$), it must halt ($f(w) = 0$).[4] Every particular policy $f \in \mathcal{F}$ induces a distribution on decisions $\{D_{k,t}\}_{k\in[n],t\in\mathbb{N}}$ and responses $\{R_{k,t}\}_{k\in[n],t\in\mathbb{N}}$, where $D_{k,t} = 1$ iff the kth device transmits at time t and $R_{k,t} \in \mathcal{R}$ is the response received by the kth device at time $t + 0.5$. In particular,

$$\mathbb{P}(D_{k,t} = 1 \mid R_{k,0}R_{k,1}\cdots R_{k,t-1} = h) = f(h), \tag{1}$$

$$R_{k,t}(w) = \begin{cases} 0, & D_{k,t}(w) = 0 \\ 1, & (D_{k,t}(w) = 1) \wedge (\forall j \neq k, D_{j,t}(w) = 0) \\ 2^+, & (D_{k,t}(w) = 1) \wedge (\exists j \neq k, D_{j,t}(w) = 1) \end{cases} \tag{2}$$

Define X_i to be the random variable of the number of time slots until device i succeeds. Note that since we number the slots starting from zero,

$$X_i = 1 + \min\{t \geq 0 \mid R_{i,t} = 1\}.$$

Note that $\{X_i\}_{i\in[n]}$ are identically distributed but not independent. For example, minimizing the average of $\{X_i\}_{i\in[n]}$ is equivalent to minimizing X_1 since:

$$\mathbb{E}\frac{\sum_{i=1}^n X_i}{n} = \frac{\sum_{i=1}^n \mathbb{E}X_i}{n} = \frac{n\mathbb{E}X_1}{n} = \mathbb{E}X_1.$$

[4] A policy may have no finite representation, and therefore may not be an *algorithm* in the usual sense.

2.1 Performance Metrics and Existence Issues

For our proofs it is helpful to assume the existence of an *optimal protocol* but it is not immediate that there *exists* such an optimal protocol. (Perhaps there is just an infinite succession of protocols, each better than the next.) In Appendix A.2 we prove that optimal protocols exist for all "reasonable" objectives. A *cost function* $T : \mathbb{Z}_+^n \to \mathbb{R}_+$ is one that maps the vector of device latencies to a single (positive) cost. The objective is to minimize $\mathbb{E}T(X_1, \ldots, X_n)$.

Definition 1 (Informal)**.** A function $T : \mathbb{Z}_+^n \to \mathbb{R}^+$ is *reasonable* if for any $s > 0$ there exists some $N > 0$ such that $T(x_1, \ldots, x_n) < s$ can be known if each of $x_1, x_2, \ldots, x_n$ is either known or known to be greater than N.

For example, $T_1(x_1, \ldots, x_n) = \frac{\sum_{k=1}^{n} x_k}{n}$ (AVG), $T_2(x_1, \ldots, x_n) = \min(x_1, \ldots, x_n)$ (MIN), and $T_3(x_1, \ldots, x_n) = \max(x_1, \ldots, x_n)$ (MAX) are all reasonable, as are all ℓ_p norms, etc.

Theorem 1. *Given the number of users n and a reasonable objective function T, there exists an optimal policy $f^* \in \mathcal{F}$ that minimizes $\mathbb{E}T(X_1, X_2, \ldots, X_n)$.*

3 Contention Resolution Between Two Parties

In this section we restrict our attention to the case $n = 2$. One key observation that makes the $n = 2$ case special is that whenever one device receives 2_+ (collision) feedback, it knows that its history and the other device's history are identical. For many reasonable objective functions the best response to a collision is to restart the protocol. This is proved formally in Theorem 2 for a class of objective functions that includes AVG, MIN, and MAX. See Appendix A.3 for proof.

Theorem 2. *Let $n = 2$, T be a reasonable objective function, and f be an optimal policy for T. Another policy f^* is defined as follows.*

$$f^*(0^k) = f(0^k), \quad \forall k \in \mathbb{N}$$
$$f^*(*2_+0^k) = f(0^k), \quad \forall k \in \mathbb{N}$$
$$f^*(*1*) = 0$$

If $T(x + c, y + c) = T(x, y) + c$ for any c (scalar additivity), then f^* is also an optimal policy for T.

Theorem 2 tells us that for the objectives that are scalar additive (including AVG, MIN, and MAX), we can restrict our attention to policies $f \in \mathcal{F}$ defined by a vector of probabilities $(p_i)_{i \geq 0}$, such that $f(w0^k) = p_k$, where w is empty or ends with 2_+, i.e., the transmission probability cannot depend on anything that happened *before* the last collision.

3.1 AVG: Minimizing the Average Transmission Time

Let $(p_k)_{k \geq 0}$ be the probability sequence corresponding to an optimal policy f for AVG. We first express our objective $\mathbb{E}X_1$ in terms of the sequence (p_k). Then, using the optimality of f, we deduce that $(p_k)_{k \geq 0}$ must take on the special form described in Theorem 3. This Theorem does not completely specify what the optimal protocol looks like. Further calculations (Corollary 1) show that choosing $N = 2$ is the best choice, and that AVG-**Contention Resolution** (see Section 1) is an optimal protocol.

Theorem 3. *There exists an integer $N > 0$ and $a_0, a_1, a_2 \in \mathbb{R}$ where $a_0 - a_1 + a_2 = 1$ and $a_0 + a_1 N + a_2 N^2 = 0$ such that the following probability sequence*

$$p_k = 1 - \frac{a_0 + a_1 k + a_2 k^2}{a_0 + a_1(k-1) + a_2(k-1)^2}, \quad 0 \le k \le N,$$

induces an optimal policy that minimizes $\mathbb{E}X_1$.

Remark. Note that defining $p_0, \ldots, p_N$ is sufficient, since $p_N = 1$ induces a certain collision if there are still 2 devices in the system, which causes the algorithm to reset. In the next time slot both devices would transmit with probability p_0.

Proof. Assume we are using an optimal policy f^* induced by a probability sequence $(p_i)_{i=1}^{\infty}$. Define $S_1, S_2 \ge 0$ to be the random variables indicating the *index* of the first slot in which devices 1 and 2 first transmit. Observe that S_1 and S_2 are i.i.d. random variables, where $\mathbb{P}(S_1 = k) = \mathbb{P}(S_2 = k) = p_k \prod_{i=0}^{k-1}(1 - p_i)$.[5] We have

$$\begin{aligned} \mathbb{E}X_1 &= \sum_{k=0}^{\infty} \left[\mathbb{P}(S_1 = k)(k + 1 + \mathbb{P}(S_2 = k) \cdot \mathbb{E}X_1)\right] \\ \iff \mathbb{E}X_1 &= \sum_{k=0}^{\infty} \left[p_k \left(\prod_{i=0}^{k-1}(1-p_i) \right) \cdot \left(k + 1 + p_k \left(\prod_{i=0}^{k-1}(1-p_i) \right) \cdot \mathbb{E}X_1 \right) \right]. \end{aligned} \tag{3}$$

Define $m_k = \prod_{i=0}^{k}(1-p_i)$ to be the probability that a device idles in time steps 0 through k, where $m_{-1} = 1$. Note that $p_k m_{k-1} = m_{k-1} - m_k$ is true for all $k \ge 0$. We can rewrite Eqn. (3) as:

$$\begin{aligned} \mathbb{E}X_1 &= \sum_{k=0}^{\infty} \left[(m_{k-1} - m_k)(k + 1 + (m_{k-1} - m_k) \cdot \mathbb{E}X_1)\right] \\ \iff \mathbb{E}X_1 &= \mathbb{E}X_1 \cdot \sum_{k=0}^{\infty}(m_{k-1} - m_k)^2 + \sum_{k=0}^{\infty}(m_{k-1} - m_k)(k+1) \\ \iff \mathbb{E}X_1 &= \frac{\sum_{k=0}^{\infty} m_{k-1}}{1 - \sum_{k=0}^{\infty}(m_{k-1} - m_k)^2}. \end{aligned} \tag{4}$$

By definition, $(m_k)_{k=-1}^{\infty}$ is a non-increasing sequence with $m_{-1} = 1$ and $m_k \ge 0$. There is no optimal policy with $m_{k-1} = m_k \ne 0$ (meaning $p_k = 0$), since otherwise we can delete m_k from the sequence, leaving the denominator unchanged but reducing the numerator. This implies $(m_k)_{k=-1}^{\infty}$ is either an infinite, positive, strictly decreasing sequence or a finite, positive, strictly-decreasing sequence followed by a tail of zeros. Pick any index $k_0 \ge 0$ such that $m_{k_0} > 0$. We know $m_{k_0 - 1} > m_{k_0} > m_{k_0+1}$. By the optimality of f^*, m_{k_0} must, holding all other parameters fixed, be the optimal choice for this parameter in its neighborhood. In other words,

$$\begin{aligned} &\frac{\partial \mathbb{E}X_1}{\partial m_{k_0}} = 0 \\ \iff &\frac{1 - \sum_{k=0}^{\infty}(m_{k-1} - m_k)^2 + \sum_{k=0}^{\infty} m_{k-1}(-2(m_{k_0-1} - m_{k_0}) + 2(m_{k_0} - m_{k_0+1}))}{(1 - \sum_{k=0}^{\infty}(m_{k-1} - m_k)^2)^2} = 0 \end{aligned}$$

[5] We use the convention that $\prod_{i=0}^{-1} a_k = 1$, where $(a_k)_{k=0}^{\infty}$ is any sequence.

Therefore we have for any $k_0 \geq 0$ such that $m_{k_0} > 0$,

$$\begin{aligned} 2m_{k_0} - m_{k_0+1} - m_{k_0-1} &= C \\ \iff m_{k_0} - m_{k_0+1} &= m_{k_0-1} - m_{k_0} + C \end{aligned} \tag{5}$$

where $C = \frac{\sum_{k=0}^{\infty}(m_{k-1}-m_k)^2-1}{2\sum_{k=0}^{\infty} m_{k-1}}$ is a real constant. Note that $C = -\frac{1}{2\mathbb{E}X_1} < 0$. Fix any $k_1 \geq 0$ such that $m_{k_1} > 0$. By summing up Eqn. (5) for $k_0 = 0, 1, \ldots, k_1$ and rearranging terms, we have

$$m_{k_1} - m_{k_1+1} = (k_1 + 1)C + m_{-1} - m_0. \tag{6}$$

Fix any $k_2 \geq 0$ such that $m_{k_2} > 0$. By summing up Eqn. (6) for $k_1 = 0, 1, \ldots, k_2$, we have

$$m_{k_2+1} = (m_0 - m_{-1})(k_2 + 2) + m_{-1} - \frac{(k_2+1)(k_2+2)}{2}C \tag{7}$$

$$= -\frac{C}{2}k_2^2 + \left(-\frac{3C}{2} + m_0 - m_{-1}\right) k_2 + 2m_0 - m_{-1} - C. \tag{8}$$

Recall that $C < 0$ and $m_k \in [0, 1]$. This rules out the possibility that the sequence $(m_k)_{k=0}^{\infty}$ is an infinite strictly decreasing sequence, since a non-degenerate quadratic function is unbounded as k goes to infinity. As a result, there must be a positive integer $N \geq 1$ for which $m_{N-1} > 0$ and $m_N = 0$. Also note that Eqn. (7) is not only true for $k_2 = 0, 1, \ldots N - 1$, but also true for $k_2 = -1$ and -2. (This can be checked by directly setting $k_2 = -1$ and -2.) We conclude that it is possible to write (m_k) as

$$m_k = a_0 + a_1 k + a_2 k^2, \quad -1 \leq k \leq N,$$

for some constants a_0, a_1, a_2 satisfying

$$\begin{aligned} m_{-1} &= a_0 - a_1 + a_2 = 1 \\ m_N &= a_0 + a_1 N + a_2 N^2 = 0. \end{aligned}$$

Writing $p_k = 1 - \frac{m_k}{m_{k-1}}$ gives the statement of the theorem. □

Based on Theorem 3, we can find the optimal probability sequence for each fixed N by choosing the best a_2. It turns out that $N = 2$ is the best choice, though $N = 3$ is only marginally worse. The proof of Corollary 1 is in Appendix A.4.

Corollary 1. **AVG-Contention Resolution** *is an optimal protocol for $n = 2$ devices under the* AVG *objective. The expected average time is $\sqrt{3/2} + 3/2 \approx 2.72474$.*

3.2 MIN: Minimizing the Earliest Transmission Time

It is straightforward to show $\mathbb{E}\min(X_1, X_2) = 2$ under the optimal policy. Nonetheless, it is useful to have a general closed form expression for $\mathbb{E}\min(X_1, X_2)$ in terms of the (m_k) sequence of an arbitrary (suboptimal) policy, as shown in the proof of Theorem 4. This will come in handy later since $\mathbb{E}\max(X_1, X_2)$ can be expressed as $2\mathbb{E}X_1 - \mathbb{E}\min(X_1, X_2)$.

Theorem 4. *The policy that minimizes $\mathbb{E}\min(X_1, X_2)$,* **MIN-Contention Resolution**, *transmits with constant probability 1/2 until successful. Using the optimal policy, $\mathbb{E}\min(X_1, X_2) = 2$.*

Proof. By Theorem 2 we can consider an optimal policy defined by a sequence of transmission probabilities $(p_k)_{k\geq 0}$. Let $H_{j,k}$ be the transmission/idle history of player $j \in \{1,2\}$ up to time slot k. Then we have

$$\begin{aligned}
\mathbb{E}\min(X_1, X_2) &= \sum_{k=0}^{\infty} \mathbb{P}(H_{1,k} = H_{2,k} = 0^k 1)(k+1+\mathbb{E}\min(X_1,X_2)) \\
&\qquad + \sum_{k=0}^{\infty} \mathbb{P}(\{H_{1,k}, H_{2,k}\} = \{0^k 1, 0^k 0\})(k+1) \\
&= \sum_{k=0}^{\infty} \left(\prod_{i=0}^{k-1} (1-p_i)^2 \cdot p_k^2 (k+1+\mathbb{E}\min(X_1,X_2)) \right) \\
&\qquad + \sum_{k=0}^{\infty} \left(\prod_{i=0}^{k-1} (1-p_i)^2 \cdot 2p_k(1-p_k)(k+1) \right) \\
&= \sum_{k=0}^{\infty} \left(\prod_{i=0}^{k-1} (1-p_i)^2 \cdot \left((1-(1-p_k)^2)(k+1) + p_k^2 \cdot \mathbb{E}\min(X_1,X_2) \right) \right).
\end{aligned}$$

Defining $m_k = \prod_{i=0}^{k}(1-p_i)$ as before, we have

$$= \sum_{k=0}^{\infty} \left((m_{k-1}^2 - m_k^2)(k+1) + m_{k-1}^2 p_k^2 \cdot \mathbb{E}\min(X_1,X_2) \right)$$

As $m_{k-1}p_k = m_{k-1} - m_k$, we can write $\mathbb{E}\min(X_1,X_2)$ in closed form as

$$\begin{aligned}
\mathbb{E}\min(X_1, X_2) &= \frac{\sum_{k=0}^{\infty} (m_{k-1}^2 - m_k^2)(k+1)}{1 - \sum_{k=0}^{\infty} (m_{k-1} - m_k)^2} \\
&= \frac{\sum_{k=0}^{\infty} m_{k-1}^2}{1 - \sum_{k=0}^{\infty} (m_{k-1} - m_k)^2} \\
&= \frac{\sum_{k=0}^{\infty} m_{k-1}^2}{2\sum_{k=0}^{\infty} (m_{k-1} - m_k) m_k} \\
&\geq \frac{\sum_{k=0}^{\infty} m_{k-1}^2}{2\sum_{k=0}^{\infty} \left(\frac{m_{k-1}}{2}\right)^2} = 2 \qquad ((m_{k-1} - m_k)m_k \text{ maximized when } m_k = m_{k-1}/2.)
\end{aligned} \tag{9}$$

$\mathbb{E}\min(X_1,X_2)$ attains minimum 2 if and only if for all $k \in \mathbb{N}$, $m_{k-1} - m_k = m_k$, i.e. $m_k = \frac{m_{k-1}}{2}$ and $m_0 = \frac{1}{2}$. Thus $p_k = 1 - \frac{m_k}{m_{k-1}} = \frac{1}{2}$ for all k. This constant probability sequence corresponds to the constant policy with sending probability $\frac{1}{2}$ (i.e., MIN-**Contention Resolution**). ☐

3.3 Max: Minimizing the Last Transmission Time

Before determining the optimal policy under the MAX objective, it is useful to have a crude estimate for its cost.

Lemma 1. *Let f be the optimal policy for the* MAX *objective and X_1^f, X_2^f be the latencies of the two devices. Then $\mathbb{E}\max(X_1^f, X_2^f) \in [3,4]$.*

Proof. The optimal policy under the MIN objective, f^*, sends with probability 1/2 until successful. It is easy to see that $\mathbb{E}\max(X_1^{f^*}, X_2^{f^*}) = 4$, so f can do no worse. Under f (or any policy), $\mathbb{E}\max(X_1^f, X_2^f) \geq 1 + \mathbb{E}\min(X_1^f, X_2^f)$. By the optimality of f^* for MIN, $\mathbb{E}\min(X_1^f, X_2^f) \geq \mathbb{E}\min(X_1^{f^*}, X_2^{f^*}) = 2$, so $\mathbb{E}\max(X_1^f, X_2^f) \geq 3$. ☐

Theorem 5. *Let f be the optimal policy for the* MAX *objective and define $1/\gamma = \mathbb{E}\max(X_1^f, X_2^f)$ to be its expected cost. Let x_1, x_2 be the roots of the polynomial*

$$x^2 - (2-\gamma)x + 1. \tag{10}$$

There exists an integer $N \geq 0$ and reals C_1, C_2 where $C_1 x_1^{-1} + C_2 x_2^{-1} = 0$ and $C_1 x_1^{N+1} + C_2 x_2^{N+1} = -1$, such that the following probability sequence

$$p_k = 1 - \frac{C_1 x_1^k + C_2 x_2^k + 1}{C_1 x_1^{k-1} + C_2 x_2^{k-1} + 1}, \quad 0 \leq k \leq N+1,$$

induces an optimal policy that minimizes $\mathbb{E}\max(X_1, X_2)$.

Remark. Note that $p_{N+1} = 1$, thus it is sufficient to only define $p_0, p_1, \ldots, p_{N+1}$.

Proof. Assume the optimal policy f is characterized by the probability sequence $(p_k)_{k=0}^{\infty}$. Using the derived expressions (Eqn. (4) and Eqn. (9)) in Theorem 3 and 4, we have

$$\begin{aligned}
\mathbb{E}\max(X_1, X_2) &= 2\mathbb{E}X_1 - \mathbb{E}\min(X_1, X_2) \\
&= 2\frac{\sum_{k=0}^{\infty} m_{k-1}}{1 - \sum_{k=0}^{\infty}(m_{k-1} - m_k)^2} - \frac{\sum_{k=0}^{\infty} m_{k-1}^2}{1 - \sum_{k=0}^{\infty}(m_{k-1} - m_k)^2} \\
&= \frac{2\sum_{k=0}^{\infty} m_{k-1} - \sum_{k=0}^{\infty} m_{k-1}^2}{1 - \sum_{k=0}^{\infty}(m_{k-1} - m_k)^2},
\end{aligned} \tag{11}$$

where $m_k = \prod_{i=0}^{k}(1 - p_i)$ with $m_{-1} = 1$.

The only requirement on the sequence $(m_k)_{k=-1}^{\infty}$ is that it is strictly decreasing with $m_k \in [0, 1]$. First we observe if $m_k = m_{k+1}$, we must have both of them equal to zero. Otherwise, we can remove m_k which will leave the denominator unchanged but reduce the numerator. Therefore, the optimal sequence is either a strictly decreasing sequence or a strictly decreasing sequence followed by a tail of zeros. Fix any $v \geq 0$ for which $m_v > 0$ we have, by the optimality of $(m_k)_{k=-1}^{\infty}$,

$$\frac{\partial \mathbb{E}\max(X_1, X_2)}{\partial m_v} = \frac{(2 - 2m_v)B - 2A(m_{v-1} - m_v - (m_v - m_{v+1}))}{D} = 0, \tag{12}$$

where $B = 1 - \sum_{k=0}^{\infty}(m_{k-1} - m_k)^2$ and $A = 2\sum_{k=0}^{\infty} m_{k-1} - \sum_{k=0}^{\infty} m_{k-1}^2$. Let $\gamma = \frac{B}{A} = \frac{1}{\mathbb{E}\max(X_1, X_2)}$, then we have, from Eqn. (12)

$$\begin{aligned}
m_{v+1} &= (2-\gamma)m_v - m_{v-1} + \gamma & (13)\\
\iff (m_{v+1} - 1) &= (2-\gamma)(m_v - 1) - (m_{v-1} - 1) & (14)
\end{aligned}$$

Eqn. (14) defines a linear homogeneous recurrence relation for the sequence $(m_{v+1} - 1)$, whose characteristic roots are $x_1, x_2 = \frac{2-\gamma \pm \sqrt{\gamma^2 - 4\gamma}}{2}$. One may verify that they satisfy the following identities.

$$\begin{aligned}
x_1 + x_2 &= 2 - \gamma & (15)\\
x_1 x_2 &= 1 & (16)
\end{aligned}$$

From Lemma 1 we know $\gamma \in [\frac{1}{4}, \frac{1}{3}]$. Thus we have $\gamma^2 - 4\gamma < 0$ which implies x_1 and x_2 are distinct conjugate numbers and of the same norm $\sqrt{x_1 x_2} = 1$. Then $m_k - 1 = C_1 x_1^k + C_2 x_2^k$ for all k for which at least one of m_{k-1}, m_k or m_{k+1} is greater than zero.

If it were the case that $m_k > 0$ for all k, then by summing (13) up for all $v \in \mathbb{N}$, we have

$$\sum_{k=0}^{\infty} m_{k+1} = (2-\gamma)\sum_{k=0}^{\infty} m_k - \sum_{k=0}^{\infty} m_{k-1} + \gamma \cdot \infty$$

which implies $\sum_{k=0}^{\infty} m_k = \infty$. This is impossible since, by the upper bound of Lemma 1,

$$4 \geq \mathbb{E}\max(X_1, X_2) = \frac{1 + 2\sum_{k=0}^{\infty} m_k - \sum_{k=0}^{\infty} m_k^2}{1 - \sum_{k=0}^{\infty}(m_{k-1} - m_k)^2} \geq \frac{1 + \sum_{k=0}^{\infty} m_k}{1}.$$

Therefore the optimal sequence must be of the form

$$m_k = (C_1 x_1^k + C_2 x_2^k + 1)\mathbb{1}_{k \leq N}$$

for some integer $N \geq 0$, where $C_1 x_1^N + C_2 x_2^N + 1 = 0$ and $C_1 x_1^{-1} + C_2 x_2^{-1} + 1 = 1$. Writing $p_k = 1 - \frac{m_k}{m_{k-1}}$ gives the statement of the theorem. □

The proof of Corollary 2 is in Appendix A.5.

Corollary 2. **MAX-Contention Resolution** *is an optimal protocol for $n = 2$ devices under the* MAX *objective. The expected maximum latency is* $1/\gamma \approx 3.33641$*, where* γ *is the unique root of* $3x^3 - 12x^2 + 10x - 2$ *in the interval* $[1/4, 1/3]$.

4 Conclusion

In this paper we established the existence of optimal contention resolution policies for any *reasonable* cost metric, and derived the first optimal protocols for resolving conflicts between $n = 2$ parties under the AVG, MIN, and MAX objectives.

Generalizing our results to $n \geq 3$ or to more complicated cost metrics (e.g., the ℓ_2 norm) is a challenging problem. Unlike the $n = 2$ case, it is not clear, for example, whether the optimal protocols for $n = 3$ select their transmission probabilities from a finite set of reals. It is also unclear whether the optimal protocols for $n = 3$ satisfy some analogue of Theorem 2, i.e., that they are "recurrent" in some way.

References

[1] Norman Abramson. The ALOHA system: another alternative for computer communications. In *Proceedings of the November 17-19, 1970, fall joint computer conference*, pages 281–285. ACM, 1970.

[2] D. J. Aldous. Ultimate instability of exponential back-off protocol for acknowledgment-based transmission control of random access communication channels. *IEEE Trans. Information Theory*, 33(2):219–223, 1987.

[3] B. Awerbuch, A. W. Richa, and C. Scheideler. A jamming-resistant MAC protocol for single-hop wireless networks. In *Proceedings of the Twenty-Seventh Annual ACM Symposium on Principles of Distributed Computing (PODC)*, pages 45–54, 2008.

[4] R. Bar-Yehuda, O. Goldreich, and A. Itai. On the time-complexity of broadcast in multi-hop radio networks: An exponential gap between determinism and randomization. *Journal of Computer and System Sciences*, 45(1):104–126, 1992.

[5] M. Bender, T. Kopelowitz, S. Pettie, and M. Young. Contention resolution with constant throughput and log-logstar channel accesses. *SIAM Journal on Computing*, 47(5):1735–1754, 2018.

[6] M. A. Bender, M. Farach-Colton, S. He, B. C. Kuszmaul, and C. E. Leiserson. Adversarial analyses of window backoff strategies. In *Proceedings 18th International Parallel and Distributed Processing Symposium (IPDPS)*, 2004.

[7] M. A. Bender, M. Farach-Colton, S. He, B. C. Kuszmaul, and C. E. Leiserson. Adversarial contention resolution for simple channels. In *Proceedings of the 17th Annual ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 325–332, 2005.

[8] M. A. Bender, J. T. Fineman, S. Gilbert, and M. Young. Scaling exponential backoff: Constant throughput, polylogarithmic channel-access attempts, and robustness. *J. ACM*, 66(1):6:1–6:33, 2019.

[9] J. Capetanakis. Tree algorithms for packet broadcast channels. *IEEE Trans. Information Theory*, 25(5):505–515, 1979.

[10] Y.-J. Chang, W. Jin, and S. Pettie. Simple contention resolution via multiplicative weight updates. In *Proceedings 2nd Symposium on Simplicity in Algorithms (SOSA)*, pages 16:1–16:16, 2019.

[11] Y.-J. Chang, T. Kopelowitz, S. Pettie, R. Wang, and W. Zhan. Exponential separations in the energy complexity of leader election. In *Proceedings 49th Annual ACM Symposium on Theory of Computing (STOC)*, pages 771–783, 2017.

[12] R. G. Gallager. Conflict resolution in random access broadcast networks. In *Proceedings AFOSR Workshop on Communications Theory Applications, Provincetown, MA, Sept 17–20*, pages 74–76, 1978.

[13] Van Jacobson. Congestion avoidance and control. In *ACM SIGCOMM computer communication review*, volume 18, pages 314–329. ACM, 1988.

[14] B.-J. Kwak, N.-O. Song, and L. E. Miller. Performance analysis of exponential backoff. *IEEE/ACM Transactions on Networking (TON)*, 13(2):343–355, 2005.

[15] R. M. Metcalfe and D. R. Boggs. Ethernet: Distributed packet switching for local computer networks. *Communications of the ACM*, 19(7):395–404, 1976.

[16] Robert M Metcalfe and David R Boggs. Ethernet: Distributed packet switching for local computer networks. *Communications of the ACM*, 19(7):395–404, 1976.

[17] V. A. Mikhailov and B. S. Tsybakov. Upper bound for the capacity of a random multiple access system. *Problemy Peredachi Informatsii*, 17(1):90–95, 1981.

[18] J. Mosely and P. A. Humblet. A class of efficient contention resolution algorithms for multiple access channels. *IEEE Trans. Communications*, 33(2):145–151, 1985.

[19] K. Nakano and S. Olariu. Uniform leader election protocols for radio networks. *IEEE Trans. Parallel Distrib. Syst.*, 13(5):516–526, 2002.

[20] C. Newport. Radio network lower bounds made easy. In *Proceedings of the 28th International Symposium on Distributed Computing (DISC)*, pages 258–272, 2014.

[21] Ravi Rajwar and James R Goodman. Speculative lock elision: Enabling highly concurrent multithreaded execution. In *Proceedings of the 34th annual ACM/IEEE international symposium on Microarchitecture*, pages 294–305. IEEE Computer Society, 2001.

[22] H. L. Royden and P. Fitzpatrick. *Real analysis*, volume 32. Macmillan, New York, 1988.

[23] B. S. Tsybakov and V. A. Mikhailov. Slotted multiaccess packet broadcasting feedback channel. *Problemy Peredachi Informatsii*, 14(4):32–59, 1978.

[24] Dingyu Wang. Optimal protocols for 2-party contention resolution. In *International Symposium on Stabilization, Safety, and Security of Distributed Systems*, 2021.

[25] D. E. Willard. Log-logarithmic selection resolution protocols in a multiple access channel. *SIAM J. Comput.*, 15(2):468–477, 1986.

Appendix A Proof of Theorems

A.1 Deduction on Random Board Model

In order to implement a contention resolution policy we need to generate biased random bits, e.g., in order to transmit with probability 1/3. However, the bias of each time step could be different and depend on the outcome of previous time steps. It will be convenient if we can generate all randomness in advance and dynamically set the biases as we go. To that end we define the Random Board model, which is just an infinite number of uniform and i.i.d. random reals in $[0, 1]$. If device k in time step t wants to generate a biased random bit b with probability 1/3 of 1, it sets $b = \mathbb{1}(U_{k,t} < 1/3)$, where $(U_{k,t})$ is the random board.

Definition 2 (Random Board)**.** A random board U with n rows are a set of i.i.d. uniformly distributed random variables $(U_{k,t})_{k\in[n],t\in\mathbb{N}}$ with range $[0, 1]$.

Definition 3 (Deduction on Random Board)**.** Given a policy $f \in \mathcal{F}$ and an outcome $(u_{k,t})_{k\in[n],t\in\mathbb{N}}$ of an n-row random board where $u_{k,t} \in [0, 1]$, we deduce $(\bar{d}_{k,t}, \bar{r}_{k,t})_{k\in[n],t\in\mathbb{N}}$ iteratively by the following rule.

$$\bar{d}_{k,t} = \mathbb{1}(u_{k,t} < f((\bar{r}_{k,0}\bar{r}_{k,1} \ldots \bar{r}_{k,t-1}))) \qquad \text{(1 iff device } k \text{ transmits at time } t\text{)}$$

$$\bar{r}_{k,t} = \begin{cases} 0, & \bar{d}_{k,t} = 0 \\ 1, & (\bar{d}_{k,t} = 1) \wedge (\forall j \neq k, \bar{d}_{j,t} = 0) \\ 2_+, & (\bar{d}_{k,t} = 1) \wedge (\exists j \neq k, \bar{d}_{j,t} = 1) \end{cases} \qquad \text{(the response at time } t + 0.5\text{)}$$

We define $\bar{D}_{k,t}$ and $\bar{R}_{k,t}$ be the random variables that maps outcomes to the deduced $\bar{d}_{k,t}$ and $\bar{r}_{k,t}$ respectively.

We demonstrate the deduction on random board by the following example.

Example 1. Let Table 1 be an outcome w_0 of a 3-row random board.

0.23371	0.281399	0.375409	0.927202	0.0824814	0.0473227	...
0.216321	0.4534	0.377702	0.573771	0.704855	0.497943	...
0.888769	0.939998	0.261829	0.343283	0.830001	0.43118	...

Table 1: An outcome w_0 of a 3-row random board

Let f_1 be the policy with constant sending probability $\frac{1}{2}$ until the message gets successfully transmitted. Let f_2 be the policy with constant sending probability $\frac{1}{3}$ before success. Then the results of deduction are shown in the following two tables $(\bar{d}_{k,t}, \bar{r}_{k,t})_{k\in[n],t\in\mathbb{N}}$ (successful transmissions are marked by a star $*$).

(send,2^+)	(send,2^+)	(send,2^+)	(idle,0)	(send*,1)	(idle,0)	...
(send,2^+)	(send,2^+)	(send,2^+)	(idle,0)	(idle,0)	(send*,1)	...
(idle,0)	(idle,0)	(send,2^+)	(send*,1)	(idle,0)	(idle,0)	...

Table 2: Deduction on w_0 using policy f_1 (constant sending probability $\frac{1}{2}$)

Lemma 2. *Given an n-row random board U, for any policy $f \in \mathcal{F}$, the deduced set of random variables $(\bar{D}_{k,t}, \bar{R}_{k,t})_{k\in[n],t\in\mathbb{N}}$ satisfies both equations (1) and (2).*

(send,2^+)	(send*,1)	(idle,0)	(idle,0)	(idle,0)	(idle,0)	...
(send,2^+)	(idle,0)	(idle,0)	(idle,0)	(idle,0)	(idle,0)	...
(idle,0)	(idle,0)	(send*,1)	(idle,0)	(idle,0)	(idle,0)	...

Table 3: Deduction on w_0 using policy f_2 (constant sending probability $\frac{1}{3}$)

Proof. Equation (2) follows directly by the definition of deduction. For equation (1), we have

$$\begin{aligned}\mathbb{P}(\bar{D}_{k,t} = 1 \mid \bar{R}_{k,1}\bar{R}_{k,2}\ldots\bar{R}_{k,t} = h) &= \mathbb{P}(U_{k,t} < f(h)) \\ &= f(h).\end{aligned}$$

□

From Lemma 2 we know, using the same f, the random processes $(D_{k,t}, R_{k,t})_{k\in[n],t\in\mathbb{N}}$ and $(\bar{D}_{k,t}, \bar{R}_{k,t})_{k\in[n],t\in\mathbb{N}}$ are of identical distributions and thus $(X_i)_{i\in[n]}$ and $(\bar{X}_i)_{i\in[n]}$ are also identically distributed. Having this, we now can consider all the policies working on a same sample space, i.e. the sample space of the random board.

A.2 Existence Issues

For notation, we denote any random variable X (e.g. X_i, $D_{k,t}$ and $R_{k,t}$) induced by policy f as X^f.

Having the random board model set up, we now can go on proving the existence theorem. We first formally define the class of *reasonable* objective functions.

Definition 4. We say an objective function $T : \mathbb{R}^n \to \mathbb{R}^+$ is *reasonable* if for any $s > 0$, there exists some $N_s > 0$ such that $\mathbb{1}(T(X_1, X_2, \ldots, X_n) < s)$ only depends on the deduction of the first N_s time slots. In other words, there exists a function h such that $\mathbb{1}(T(X_1, X_2, \ldots, X_n) < s) = h(\{D_{k,t}, R_{k,t}\}_{k\in[n],t\in[N_s]})$.

Proof of Theorem 1. Let $g : \mathcal{F} \to \mathbb{R}_+ \cup \{\infty\}$ be the function that maps policy f to $\mathbb{E}T(X_1^f, X_2^f, \ldots, X_n^f)$. For simplicity, we denote $(X_1^f, X_2^f, \ldots, X_n^f)$ by X^f.

Define J as $\inf\{g(f) \mid f \in \mathcal{F}\}$ which exists since g is bounded below (g is non-negative). Since the set of finite strings is countable, we can identify $\mathcal{F}$ by $[0,1]^{\mathbb{N}}$. Note that $[0,1]$ is compact and thus, by diagonalization argument, we can find a sequence of $(f_i)_{i\in\mathbb{N}}$ where $f_i \in \mathcal{F}$ converges to f^* point-wisely with $\lim_{i\to\infty} g(f_i) = J$. We have to show $g(f^*) = J$. By definition, we have

$$\begin{aligned}g(f^*) &= \mathbb{E}T(X^{f^*}) \\ &= \int_0^\infty \mathbb{P}(T(X^{f^*}) \geq s)ds \\ &= \int_0^\infty (1 - \mathbb{P}(T(X^{f^*}) < s))ds\end{aligned}$$

For a fixed $s > 0$, by the definition of a reasonbale objective function, there exists a constant $N_s > 0$ such that $\mathbb{1}(T(X^f) < s)$ only depends on the deduction of the first N_s time slots. Since f_j converges to f^* point-wisely and the set of strings $\{w \mid |w| < N_s\}$ is finite[6], we can choose ϵ and N so that for any $j > N$ and w with $|w| < N_s$, we have $|f_j(w) - f^*(w)| < \epsilon$. Then for the two policies f_j and f^*, the probability that the deduction of the first N_s time slots will not change is larger than $(1-\epsilon)^{nN_s}$.

[6] $|w|$ is equal to the length of the string w.

In fact, let w be the history of a user with $|w| < N_s$ and U be the cell on the random board that will be used by the user, we have $\mathbb{P}(\mathbb{1}(U < f(w)) = \mathbb{1}(U < f^*(w))) = 1 - |f(w) - f^*(w)| \geq (1-\epsilon)$. Therefore, $\mathbb{P}(E_k \mid E_{k-1}) \geq (1-\epsilon)^n$, where E_k is the event that the deductions of the first k columns are identical. We get the estimation we want by writing out

$$\begin{aligned}\mathbb{P}(E_{N_s-1}) &= \mathbb{P}(E_{N_s-1} \mid E_{N_s-2})\mathbb{P}(E_{N_s-2} \mid E_{N_s-3}) \cdots \mathbb{P}(E_1 \mid E_0)\mathbb{P}(E_0) \\ &\geq (1-\epsilon)^{nN_s}\end{aligned}$$

Thus, let Ω be the sample space of the random board,

$$\begin{aligned}|\mathbb{P}(T(X^{f_j}) < s) - \mathbb{P}(T(X^{f^*}) < s)| &= \left| \int_\Omega \mathbb{1}(T(X^{f_j}(w)) < s) - \mathbb{1}(T(X^{f^*}(w)) < s) dw \right| \\ &\leq \int_\Omega \left| \mathbb{1}(T(X^{f_j}(w)) < s) - \mathbb{1}(T(X^{f^*}(w)) < s) \right| dw \\ &= \mathbb{P}\left(\mathbb{1}\left((T(X^{f_j}) < s\right) \neq \mathbb{1}\left(T(X^{f^*}) < s)\right)\right) \\ &= 1 - \mathbb{P}\left(\mathbb{1}\left((T(X^{f_j}) < s\right) = \mathbb{1}\left(T(X^{f^*}) < s)\right)\right) \\ &\leq 1 - \mathbb{P}(E_{N_s}) \\ &\leq 1 - (1-\epsilon)^{nN_s}\end{aligned}$$

which can be made arbitrarily small by choosing ϵ small enough. Thus we have $\mathbb{P}(T(X^{\lim_{j\to\infty} f_j}) < s) = \lim_{j\to\infty} \mathbb{P}(T(X^{f_j}) < s)$. Therefore we have

$$\begin{aligned}g(f^*) &= \int_0^\infty (1 - \mathbb{P}(T(X^{f^*}) < s))ds \\ &= \int_0^\infty (1 - \lim_{j\to\infty} \mathbb{P}(T(X^{f_j}) < s))ds \\ &= \int_0^\infty \lim_{j\to\infty} \mathbb{P}(T(X^{f_j}) \geq s)ds \\ &\leq \lim_{j\to\infty} \int_0^\infty \mathbb{P}(T(X^{f_j}) \geq s)ds \\ &= J\end{aligned}$$

Note that we have used the Fatou's Lemma [22] to change the order of integration and the limit. We have $g(f^*) \leq J$ but J is the infimum. Thus f^* is an optimal policy.

□

A.3 Optimal Recurrent Policies — Proof of Theorem 2

Proof of Theorem 2. Let C be the random variable indicating the time of the first collision. We define $C = -1$ if there is no collision. By the definition of f^*, we know $C^{f^*} = C^f$ since $f^*(w) = f(w)$ if $2_+ \notin w$. Also, if there is no collision, the deduction of the two policies should be identical. Thus we have $\mathbb{E}(T(X_1^f, X_2^f) \mid C^f = -1) = \mathbb{E}(T(X_1^{f^*}, X_2^{f^*}) \mid C^{f^*} = -1) =: M$. Let $\mathbb{P}(C^f = k) = \mathbb{P}(C^{f^*} = k) = q_k$ and $f_k(w) = f(0^k 2_+ w)$ for any $w \in \mathcal{R}^*$. Note that both users must keep idle before the first collision (if there is a collision). Therefore if the first collision happens at time k, after time k, the two users would behave as if they have restarted the process with a new policy f_k. Then we have,

$$\begin{aligned}
\mathbb{E}T(X_1^f, X_2^f) &= Mq_{-1} + \sum_{k=0}^{\infty} \mathbb{E}(T(X_1^f, X_2^f) \mid C^f = k) q_k \\
&= Mq_{-1} + \sum_{k=0}^{\infty} \mathbb{E}T(k+1+X_1^{f_k}, k+1+X_2^{f_k}) q_k \\
&= Mq_{-1} + \sum_{k=0}^{\infty} (k+1) q_k + \sum_{k=0}^{\infty} \mathbb{E}T(X_1^{f_k}, X_2^{f_k}) q_k \\
&\geq Mq_{-1} + \sum_{k=0}^{\infty} (k+1) q_k + \mathbb{E}T(X_1^f, X_2^f) \sum_{k=0}^{\infty} q_k
\end{aligned}$$

Note that if $\sum_{k=0}^{\infty} q_k = 1$, then $q_{-1} = 0$ and we would have $\sum_{k=0}^{\infty}(k+1)q_k \leq 0$, which is impossible. Thus we must have $\sum_{k=0}^{\infty} q_k < 1$.

Now by the definition of f^*, we have $f^*(w) = f^*(0^k 2_+ w)$. Similarly, we have

$$\begin{aligned}
\mathbb{E}T(X_1^{f^*}, X_2^{f^*}) &= Mq_{-1} + \sum_{k=0}^{\infty} \mathbb{E}(T(X_1^{f^*}, X_2^{f^*}) \mid C^{f^*} = k) q_k \\
&= Mq_{-1} + \sum_{k=0}^{\infty} \mathbb{E}T(k+1+X_1^{f^*}, k+1+X_2^{f^*}) q_k \\
&= Mq_{-1} + \sum_{k=0}^{\infty} (k+1) q_k + \sum_{k=0}^{\infty} \mathbb{E}T(X_1^{f^*}, X_2^{f^*}) q_k
\end{aligned}$$

We conclude that $\mathbb{E}T(X_1^f, X_2^f) \geq \frac{Mq_{-1} + \sum_{k=0}^{\infty}(k+1)q_k}{1 - \sum_{k=0}^{\infty} q_k} = \mathbb{E}T(X_1^{f^*}, X_2^{f^*})$. Since f is optimal, then f^* is also optimal. □

A.4 Proof of Corollary 1 — The Avg Objective

Proof of Corollary 1. By Theorem 3, the optimal sequence $(m_k)_{k=-1}^{\infty}$ is given by $m_k = (a_0 + a_1 k + a_2 k^2)\mathbb{1}_{k \leq N}$, for some integer $N > 0$, where

$$a_0 - a_1 + a_2 = 1 \tag{17}$$

$$a_0 + a_1 N + a_2 N^2 = 0 \tag{18}$$

Note that by the definition, we must have $m_{k-1} \leq m_k$ for all $k = 0, 1, \ldots, N$. We necessarily have

$$m_{-1} \geq m_0 \implies a_0 - a_1 + a_2 \geq a_0 \tag{19}$$

$$m_{N-1} \geq m_N \implies a_0 + a_1(N-1) + a_2(N-1)^2 \geq a_0 + a_1 N + a_2 N^2 \tag{20}$$

From Eqn. (17) and (18), we have $a_0 = \frac{N - a_2(N^2+N)}{N+1}$ and $a_1 = \frac{-a_2(N^2-1)-1}{N+1}$. Insert both expressions into Eqn. (19) and (20), we get the constraints on a_2:

$$a_2 \in \left[-\frac{1}{N+N^2}, \frac{1}{N+N^2}\right]. \tag{21}$$

In fact, the conditions that $m_{-1} \geq m_0$ and $m_{N-1} \geq m_N$ are sufficient for $m_{k-1} \leq m_k$ for all $k = 0, 1, \ldots, N$. To see this, let $g(x) = a_0 + a_1 x + a_2 x^2$ be the quadratic function where $g(k) = m_k$

for all $k = -1, 0, 1, \ldots, N$. Once we know $g(-1) \geq g(0)$ and $g(N-1) \geq g(N)$, we know there exist $x_1 \in [-1, 0]$ and $x_2 \in [N-1, N]$ such that $g'(x_1) \leq 0$ and $g'(x_2) \leq 0$. Since $g'(x) = a_1 + 2a_2 x$ is a linear function, $g'(x_1), g'(x_2) \leq 0$ implies $g'(x) \leq 0$ for any $x \in [x_1, x_2]$. Thus $g'(x) \leq 0$ for any $x \in [0, N-1]$. We conclude that $m_{k-1} = g(k-1) \leq g(k) = m_k$ is also true for any $k = 1, \ldots, N-1$.

Using Eqn. (4) derived in Theorem 3, we have

$$\begin{aligned} \mathbb{E}X_1 &= \frac{\sum_{k=-1}^{N-1}(a_0 + a_1 k + a_2 k^2)}{1 - \sum_{k=0}^{N}(-a_1 + a_2(-2k+1))^2} \\ &= \frac{(N+1)(N+2)(a_2 N(N+1) - 3)}{2N\left(a_2^2(N+1)^2(N+2) - 3\right)} \end{aligned} \tag{22}$$

Fixing N, to minimize $\mathbb{E}X_1$, we have

$$\frac{\partial \mathbb{E}X_1}{\partial a_2} = -\frac{(N+1)^2(N+2)(a_2^2 N(N+2)(N+1)^2 - 6a_2(N+2)(N+1) + 3N)}{2N\left(a_2^2(N+1)^2(N+2) - 3\right)^2} = 0$$

Then $a_2 = \frac{3N^2+9N+6 \pm \sqrt{3}\sqrt{-N^5-N^4+13N^3+37N^2+36N+12}}{N^4+4N^3+5N^2+2N}$. It is easy to check that a_2 becomes complex for all $N \geq 5$, which implies $\frac{\partial \mathbb{E}X_1}{\partial a_2} \leq 0$ for any a_2. That is to say for $N \geq 5$, the larger a_2 is, the smaller $\mathbb{E}X_1$ is. Thus the best a_2 one can pick under the constraint (21) is $\frac{1}{N+N^2}$. In this case, by inserting $a_2 = \frac{1}{N+N^2}$ into Eqn. (22) we have

$$\begin{aligned} \mathbb{E}X_1 &= \frac{(N+1)(N+2)(1-3)}{2N\left(\frac{N+2}{N^2} - 3\right)} \\ &= \frac{N(N+1)(N+2)}{3N^2 - N - 2}. \end{aligned} \tag{23}$$

Then we differentiate Eqn. (23) by N:

$$\frac{\partial \mathbb{E}X_1}{\partial N} = \frac{1}{15}\left(-\frac{8}{(3N+2)^2} - \frac{18}{(N-1)^2} + 5\right) > 0$$

for $N \geq 5$. Therefore, in the case where $N \geq 5$, it it best to choose $N = 5$, yielding $\mathbb{E}X_1 = \frac{N(N+1)(N+2)}{3N^2-N-2} = \frac{105}{34} \approx 3.09$.

Now we consider the cases where $N \in \{1, 2, 3, 4\}$. We give the optimal results for each case in the Table 4. Note that we have $a_2 \in [-\frac{1}{N+N^2}, \frac{1}{N+N^2}]$ for each N.

Combining all previous results we conclude that a globally optimal algorithm that minimizes the average waiting time of a two-party contention is obtained when $N = 2$ and $a_2 = \frac{1}{2} - \frac{1}{\sqrt{6}}$. In this case, we have $a_0 = \frac{N - a_2(N^2+N)}{N+1} = \frac{1}{3}(\sqrt{6} - 1)$ and $a_1 = \frac{-a_2(N^2-1)-1}{N+1} = \frac{\sqrt{6}-5}{6}$. Then we have $m_0 = a_0 = \frac{1}{3}(\sqrt{6}-1)$, $m_1 = a_0 + a_1 + a_2 = \frac{1}{3}(\sqrt{6}-2)$ and $m_2 = 0$. Finally we get the optimal probability sequence.

$$\begin{aligned} p_0 &= 1 - m_0 = \frac{4 - \sqrt{6}}{3} \approx 0.516837, \\ p_1 &= 1 - \frac{m_1}{m_0} = \frac{1 + \sqrt{6}}{5} \approx 0.689898, \\ p_2 &= 1 - \frac{m_2}{m_1} = 1. \end{aligned}$$

This is precisely the AVG-**Contention Resolution** protocol. □

N	$\mathbb{E}X_1$ (Eqn. (22))	optimal a_2	optimal $\mathbb{E}X_1$
1	$\frac{3-2a_2}{1-4a_2^2}$	$\frac{3}{2}-\sqrt{2}$	$\sqrt{2}+\frac{3}{2}\approx 2.91421$
2	$\frac{3-6a_2}{1-12a_2^2}$	$\frac{1}{2}-\frac{1}{\sqrt{6}}$	$\frac{1}{2}(\sqrt{6}+3)\approx 2.72474$
3	$\frac{10-40a_2}{3-80a_2^2}$	$\frac{1}{12}$	$\frac{30}{11}\approx 2.72727$
4	$\frac{5\,(20a_2-3)}{200a_2^2-4}$	$\frac{1}{20}$	$\frac{20}{7}\approx 2.85714$

Table 4: Solutions for $N\in\{1,2,3,4\}$

A.5 Proof of Corollary 2 — The Max Objective

Proof of Corollary 2. We continue to use all the derived results in Theorem 5. Remember that x_1 and x_2 depend on γ, and satisfy the equalities $x_1+x_2=2-\gamma$ and $x_1x_2=1$. Recall that $1/\gamma=\mathbb{E}\max(X_1,X_2)$. Using the fact that $m_{-1}=1$, it follows from Eqn. (11) that γ can be written as

$$\begin{aligned}\gamma &= \frac{1-\sum_{k=0}^{N+1}(m_{k-1}-m_k)^2}{1+2\sum_{k=0}^{N}m_k-\sum_{k=0}^{N}m_k^2}\\ &= \frac{1-(1-m_0)^2-\sum_{k=0}^{N}\left(C_1x_1^k(1-x_1)+C_2x_2^k(1-x_2)\right)^2}{1+2\sum_{k=0}^{N}\left(C_1x_1^k+C_2x_2^k+1\right)-\sum_{k=0}^{N}\left(C_1x_1^k+C_2x_2^k+1\right)^2}\end{aligned}$$

Since $x_1x_2=1$, this can be written as follows.

$$=\frac{1-(1-m_0)^2-\sum_{k=0}^{N}\left(C_1^2\left(1-x_1\right)^2\left(x_1^2\right)^k+2C_1C_2(1-x_1)(1-x_2)+C_2^2(1-x_2)^2\left(x_2^2\right)^k\right)}{N+2-\sum_{k=0}^{N}\left(C_1^2\left(x_1^2\right)^k+C_2^2\left(x_2^2\right)^k+2C_1C_2\right)}$$

By organizing the terms we have (using identities (15), (16), the polynomial (10) that defines x_1 and x_2 and the definition $m_k=C_1x_1^k+C_2x_2^k+1$)

$$\begin{aligned}0 &= \gamma(N+2)-1+(1-m_0)^2-2(N+1)C_1C_2(-(1-x_1)(1-x_2)+\gamma)\\ &\quad +C_1^2\left((1-x_1)^2-\gamma\right)\sum_{k=0}^{N}\left(x_1^2\right)^k+C_2^2\left((1-x_2)^2-\gamma\right)\sum_{k=0}^{N}\left(x_2^2\right)^k\end{aligned}$$

Using the fact that $(1-x_1)(1-x_2)=2-(x_1+x_2)=\gamma$ and that $-\gamma(1+x_1)=(x_1+x_2-2)(1+x_1)=x_1^2-x_1-1+x_2=(1-x_1)^2-2+x_1+x_2=(1-x_1)^2-\gamma$, we can simplify this as follows.

$$\begin{aligned}&= \gamma(N+2)-1+(1-m_0)^2-C_1^2\gamma(1+x_1)\sum_{k=0}^{N}\left(x_1^2\right)^k-C_2^2\gamma(1+x_2)\sum_{k=0}^{N}\left(x_2^2\right)^k\\ &= \gamma(N+2)-1+(1-m_0)^2-C_1^2\gamma\frac{1-\left(x_1^2\right)^{N+1}}{1-x_1}-C_2^2\gamma\frac{1-\left(x_2^2\right)^{N+1}}{1-x_2}\end{aligned}$$

We will now put the last two terms under the common denominator $(1-x_1)(1-x_2)=\gamma$, thereby cancelling the γ factors in both terms.

$$
\begin{aligned}
&= \gamma(N+2) - 1 + (1-m_0)^2 \\
&\quad - \left(C_1^2 + C_2^2 - C_1^2 x_2 - C_2^2 x_1 - C_1^2 x_1^{2(N+1)} - C_2^2 x_2^{2(N+1)} + C_1^2 x_1^{2N+1} + C_2^2 x_2^{2N+1}\right)
\end{aligned}
$$

Since $m_0 = C_0 + C_1 + 1$, we have

$$
\begin{aligned}
&= \gamma(N+2) - 1 + 2C_1C_2 \\
&\quad - (-C_1^2 x_2 - C_2^2 x_1 - C_1^2 x_1^{2(N+1)} - C_2^2 x_2^{2(N+1)} + C_1^2 x_1^{2N+1} + C_2^2 x_2^{2N+1}) \\
&= \gamma(N+2) - 1 + (m_{N+1}-1)^2 - (-C_1^2 x_2 - C_2^2 x_1 + C_1^2 x_1^{2N+1} + C_2^2 x_2^{2N+1}) \\
&= \gamma(N+2) - 1 + (m_{N+1}-1)^2 - (m_N - 1)(m_{N+1}-1) + (m_0-1)(m_{-1}-1)
\end{aligned}
$$

Note that we have $m_{-1} = 1$ and $m_{N+1} = 0$, hence

$$
= \gamma(N+2) + m_N - 1.
$$

Rearranging terms, we have:

$$
N = \frac{1-m_N}{\gamma} - 2 < 2 \tag{24}
$$

The last inequality follows from the fact that $m_N \in (0,1]$ and $\frac{1}{\gamma} \in [3,4]$. Thus the only possible values of N are 0 and 1.

If $N = 0$, we have the following set of equations. [7]

$$
\begin{aligned}
\gamma(N+2) + m_N - 1 = 2\gamma + C_1 + C_2 &= 0 \\
C_1 x_1^{-1} + C_2 x_2^{-1} + 1 &= 1 \qquad (m_{-1} = 1) \\
C_1 x_1 + C_2 x_2 + 1 &= 0 \qquad (m_1 = 0)
\end{aligned}
$$

whose solution is $\gamma = \frac{2-\sqrt{2}}{2}$ where $\frac{1}{\gamma} \approx 3.41421$.

If $N = 1$, we have the following set of equations.

$$
\gamma(N+2) + m_N - 1 = 3\gamma + C_1 x_1 + C_2 x_2 = 0 \tag{25}
$$

$$
C_1 x_1^{-1} + C_2 x_2^{-1} + 1 = 1 \qquad (m_{-1} = 1) \tag{26}
$$

$$
C_1 x_1^2 + C_2 x_2^2 + 1 = 0 \qquad (m_2 = 0) \tag{27}
$$

whose solution gives

$$
\gamma \approx 0.299723, \text{ where } 3\gamma^3 - 12\gamma^2 + 10\gamma - 2 = 0
$$

We see in this case $\frac{1}{\gamma} \approx 3.33641$, which is better than the case where $N = 0$. The corresponding $C_1 \approx -0.264419 - 0.426908i$ and $C_2 \approx -0.264419 + 0.426908i$, which are a pair of conjugate roots of polynomial $76x^6 - 532x^5 + 664x^4 + 3288x^3 + 4680x^2 + 2268x + 729$.

Finally we get

$$
\begin{aligned}
p_0 &= 1 - m_0 = -C_1 - C_2 = \alpha \approx 0.528837, \text{ where } \alpha^3 + 7\alpha^2 - 21\alpha + 9 = 0, \\
p_1 &= 1 - \frac{m_1}{m_0} = 1 - \frac{C_1 x_1 + C_2 x_2 + 1}{C_1 + C_2 + 1} = \beta \approx 0.785997, \text{ where } 4\beta^3 - 8\beta^2 + 3 = 0, \\
p_2 &= 1 - \frac{m_2}{m_1} = 1.
\end{aligned}
$$

This is precisely the MAX-**Contention Resolution** protocol. □

[7] We have used the algebra software *Mathematica* to give the exact form of the solutions for both cases.