Winning the Lottery: Learning Perfect Coordination with Minimal Feedback

William Zame, Jie Xu, and Mihaela van der Schaar

***Abstract*—Coordination is a central problem whenever stations (or nodes or users) share resources across a network. In the absence of coordination, there will be collision, congestion or interference, with concomitant loss of performance. This paper proposes new protocols, which we call *perfect coordination (PC) protocols*, that solve the coordination problem. PC protocols are completely distributed (requiring neither central control nor the exchange of any control messages), fast (with speeds comparable to those of any existing protocols), fully efficient (achieving perfect coordination, with no collisions and no gaps) and require minimal feedback. PC protocols rely heavily on *learning*, exploiting the possibility to use both actions and silence as messages and the ability of stations to learn from their own histories while simultaneously enabling the learning of other stations. PC protocols can be formulated as finite automata and implemented using currently existing technology (e.g., wireless cards). Simulations show that, in a variety of deployment scenarios, PC protocols outperform existing state-of-the-art protocols – despite requiring much less feedback.**

***Index Terms*—Perfect Coordination, Learning in Networks, Multi-user communication, MAC protocols, Slotted Aloha**

I. Introduction

Multi-access communications form the basis for numerous networks: wireless, cognitive radio, satellite and many others. In these networks, a number of stations share common network resources (e.g., wireless spectrum). If more than one station attempts to access the same resource in the same slot/period, collision, congestion or interference, with concomitant loss of performance, will typically occur. Coordination among stations has the potential to avoid conflict and improve utilization of network resources – especially to increase throughput and decrease delay. Such coordination can be obtained if the network is centrally controlled or if the stations can exchange control messages – but central control and the exchange of control messages between stations are typically wasteful of resources and often impossible, so the first *desideratum* for a coordination protocol is that it be *distributed*. Because gathering information (sensing) also uses resources, a second *desideratum* is that the protocol use *minimal feedback*. A third *desideratum* is that it achieve *maximal goodput*. The search for such protocols is by now quite old. The most familiar distributed MAC protocols are undoubtedly slotted Aloha [1] [2], [3] which was proposed more than 40 years ago, and the distributed coordination function (DCF) used in the 802.11 standard, which is 15 years old [4], [5]. These protocols are widely studied and used, but both fail the last two *desiderata*: they require substantial feedback and, because they do not eliminate collisions and empty slots, do not achieve maximal goodput.

This paper proposes and analyzes new protocols that meet all three of these *desiderata*; we call them *Perfect Coordination protocols*. For concreteness of discussion, we focus on MAC protocols and use the example of a wireless network, but our results apply to multi-access networks quite generally. The protocols we describe are perfectly distributed and require no central control, no exchange of control messages between stations and minimal feedback: stations that transmit learn whether or not their transmission is successful; stations that are idle can not/do not sense the channel and hence learn *nothing*. The assumption of minimal feedback – which distinguishes the present work from all of the literature of which we are aware – makes coordination much more difficult but is important: sensing activity on the channel when not transmitting requires the expenditure of energy and is prone to serious errors (because of the difficulty of distinguishing the traffic of other stations from ambient noise) which increase the fragility of protocols. The protocols we introduce converge as rapidly and with probability as high as previous protocols, and achieve perfect coordination, not merely zero collision – hence greater goodput and smaller delay – *despite requiring much less feedback*.

We emphasize that our protocols prescribe behavior – actions conditional on histories – for *individual* stations. This is necessary because stations are anonymous and *ex ante* indistinguishable and because stations observe only their own histories and not the histories of other stations.[1] However, our protocols prescribe randomized behavior for stations, so that even though stations follow the same instructions they may experience different realizations and hence not take the same actions or experience the same histories. Because they prescribe randomized behavior, our protocols do not guarantee perfect coordination in any given finite length of time because such a guarantee is *impossible*: we prove that, even if we allow for complete sensing, *there do not exist anonymous protocols that guarantee perfect coordination in any given*

W. Zame: Distinguished Professor, Departments of Economics and Mathematics, UCLA, Los Angeles, CA, 90095

J. Xu: Department of Electrical Engineering, UCLA, Los Angeles, CA, 90095; Corresponding Author: jiexu@ucla.edu

M. van der Schaar: Professor, Department of Electrical Engineering, UCLA, Los Angeles, CA, 90095; Fellow, IEEE

Financial support was provided by the National Science Foundation under grants 0617027 and 0830556, by the UCLA Academic Senate Committee on Research and by the Einaudi Institute for Economics and Finance (Rome). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of any funding agency.

[1] If stations were neither anonymous nor indistinguishable, but could be uniquely identified *ex ante* – perhaps by station ID – coordination would be trivial; the setting we consider seems much more common.

finite length of time (Theorem 1).[2] Indeed we prove that the probability of achieving even zero collisions (which is much weaker than perfect coordination) in a given finite length of time is bounded away from one *independently of the protocol or the information structure*.

Our protocols lead stations to learn about the evolving state of the system, to condition their pattern of actions on what is learned, and to enable the learning of other stations. In comparison with previous protocols, stations learn more and use more of what they have learned – especially about the pattern of actions of *other* stations. Remarkably, stations can learn all of this *solely on the basis of their own histories of successful and unsuccessful transmissions*. (We reiterate that stations do not sense the channel, and hence learn nothing, when they are not transmitting.) This is possible because stations learn *cooperatively*: simultaneously learning about themselves and about other stations while also promoting the learning of other stations. This learning exploits two opposite facets of the environment: first, that actions of stations can be used as implicit messages, and second that silence can also be a message.[3] In addition to cooperative learning, actions as signals and silence as a signal, our protocols make extensive use of the idea that stations that are indistinguishable *ex ante* and randomize in an identical fashion may still experience different realizations of that randomization and hence become distinguishable *ex post*. We believe this paper is the first to exploit the combined power of these forces.

We emphasize that the protocols we propose require only finite memory, can be formulated as finite automata, and can be implemented using current wireless cards without requiring the development or deployment of any new hardware; see [7] for instance. We do require that local clocks of the stations be synchronized; this could be accomplished by time stamping from the access point (as in [14]), by GPS coordination (as in [13]) or by various other methods.

We propose two protocols, the first designed for settings in which the number of stations is known, the second designed for settings in which the number of stations is not known (but an upper bound is known); each of these settings represents a realistic set of environments. Not surprisingly, the second protocol runs more slowly than the first: because the number of stations is not known, in order to achieve perfect coordination with neither collisions nor gaps, the second protocol must incorporate checks to see that all stations have been sequenced and that the sequencing of stations has not left gaps, and it must do so in a way that maintains coordination across stations. (We will later identify another, more subtle, reason why the second protocol runs more slowly than the first.) Both our protocols handle smoothly the setting in which some stations have more traffic, and hence require more than one slot, via the simple device of viewing a station that needs k slots as k stations that each need a single slot. Our second protocol, which is designed for the setting in which the number of stations is unknown, is easily adapted to cover a dynamic environment in which stations enter and exit – provided they do not enter or exit too frequently. In practice, the administrator will choose and configure a protocol based on the setting, employing the first when the number of stations is known and fixed and the second when (as is often the case in practice) the number of stations is not known and/or entry and exit are concerns.

We emphasize that our protocols prescribe behavior – actions conditional on histories – for *individual* stations. This is of course necessary because stations are anonymous (indistinguishable to the designer) and hence must be given the same instructions, and because stations observe only their own histories and not the histories of other stations. However, as we have pointed out above, stations that follow the same instructions but experience different realization of their own randomizations need not experience the same histories nor take the same actions.

The idea of utilizing collisions as a coordination device can be found in various existing network protocols, most notably in various versions of slotted Aloha protocols and in exponential backoff protocols. By and large, these protocols, are designed in an *ad hoc* manner and utilize available past information in a limited way, yielding only limited performance improvements. For instance, in the slotted Aloha protocol, the transmission probability used by a station in a given slot/period is independent of the station's experience (waiting, collision) in the previous slot/period and on the time at which the slot/period occurs. As our protocols demonstrate, enormous performance improvements can be achieved by more information.

II. Related Work

Coordination protocols [8]-[20] have received substantial attention in the literature; [8], [9], [10], [11], [12] are closest to the present work.[4] (We provide summaries of some of these protocols in Appendix C.) Table I below provides some comparisons between the current work and those papers. We defer more detailed comparisons until we present simulations in Section V, but for now we would like to highlight a few points. First, all of these other papers assume that stations can sense the state of the channel when they are not transmitting; we do not. Second, all of these papers propose protocols that *do not* learn the number of stations and *do not* achieve perfect coordination, hence yield sub-optimal (possibly highly-suboptimal) goodput; our Perfect Coordination protocols *do* learn the number of stations and *do* achieve perfect coordination, hence yield optimal goodput. Third, the speed of our Perfect Coordination protocols is comparable to that of these

[2] We note that previous protocols also prescribe randomized behavior and that analyses of these protocols do not claim that they guarantee perfect coordination in any given finite length of time.

[3] That the apparent absence of information can in fact *be* information has been well-known at least since Sherlock Holmes. Recall the exchange between Inspector Gregory of Scotland Yard and Sherlock Holmes in Arthur Conan Doyle's story 'Silver Blaze' [6]: Gregory: "Is there any other point to which you would wish to draw my attention?" Holmes: "To the curious incident of the dog in the night-time." Gregory: "The dog did nothing in the night-time." Holmes: "That was the curious incident."

[4] [13] approach the problem somewhat differently and assume a very different information structure: stations can observe whether or not there was a successful transmission but cannot distinguish between a collision and ambient noise on an idle channel. This seems an unrealistic assumption, especially because such observation would almost surely entail a large error rate, which would be incompatible with the protocols they propose.

	Lee+Walrand (ZC)	*Fang et al (L-ZC)*	*Park + van der Schaar*	*He et al (SRB)*	*This paper (PC)*
Channel Sensing When Not Transmitting	Yes	Yes	Yes	Yes	No
Fully Coordinated	No	No	No	No	Yes
Learn # of Stations	No	No	No	No	Yes
Number of Stations	Unknown	Unknown	Known/ Unknown	Unknown	Known/ Unknown
Convergence (Small Networks)***	Very Fast	Very Fast	Fast	Medium	Very Fast
Convergence (Large Networks)***	Fast	Fast	Medium	Slow	Fast

* Convergence speed is the number of slots required such that the system enters the steady state (perfect coordination or zero collision) with probability 0.999. (Very Fast: < 100 slots; Fast : 100 - 500 slots; Medium: 500 – 5000 slots; Slow: > 5000 slots.)
** Small Networks: number of stations from 1 to 10; Large Networks: number of stations from 20 to 50.

TABLE I
COMPARISON OF ASSUMPTIONS, RESULTS

other protocols, despite the fact that we assume much less feedback.[5]

In what follows, Section III describes the environments in which we work and the framework we use. Section IV demonstrates that deterministic perfect coordination in a given finite time is impossible. Section V presents our protocols, establishes analytic estimates of convergence probabilities and speeds, discusses entry/exit and robustness. Section VI uses simulations to provide performance analyses of our protocols and comparisons with other protocols. Section VII concludes. Appendices A and B present flow charts for our protocols.

III. FRAMEWORK

We consider the interaction among N stations over a (potentially) infinite time horizon. Time is divided into discrete *slots* or *periods* indexed by $t = 0, 1, \ldots$. At each time, each station must choose whether to transmit or not (remain silent); we use $1, 0$ (respectively) for these decisions. We write $A = \{0, 1\}$ for the set of possible actions of each station, $A^N = \{0, 1\}^N$ for the set of possible action profiles and $a^t = (a_1^t, \ldots, a_N^t)$ for the vector of actions of the N stations in slot t. The set of *channel states* is $S = \{s_0, s_1, s_2, \ldots s_N\}$. Given the vector a^t of transmissions at time t the resulting channel state is

$$\sigma(a^t) = s_m \iff \sum_{i=1}^{N} a_i^t = m \tag{1}$$

The channel can support exactly one transmission at a given time: if more than one station transmits, the result is a collision and no transmissions are successful. Thus, we refer to state 0 as *idle* and state 1 as *success*; in other states, there is a collision.

A. Information

We formalize the information available to a station in terms of information partitions of the set S of channel states. By definition an *information partition* of S is a collection $\mathbf{X} = \{X_i\}$ of pairwise disjoint, non-empty subsets of S whose union is S. Because we allow for the possibility that a station that transmits observes more than a station that does not transmit, we prescribe two information partitions $\mathbf{X}^0, \mathbf{X}^1$: if the station chooses action $a = 0, 1$ and the channel state is s the agent observes its own action a and the unique element $\mathbf{X}^a(s) \in \mathbf{X}^a$ that contains s. We consider three information conditions: in the *No Silent Sensing* condition, stations that transmit observe whether or not their transmission is successful but stations that are silent observe nothing; in the *Silent Sensing* condition, stations that transmit observe whether or not their transmission is successful but stations that are silent observe only whether the channel is idle or busy; in the *Complete Sensing* condition, both stations that transmit and stations that are silent observe the number of stations that transmitted. The corresponding information partitions are:

(i) *No Silent Sensing*: $\mathbf{X}^0 = \{S\}$;
$\mathbf{X}^1 = \{\{s_0\}, \{s_1\}, \{s_2, \ldots, s_N\}\}$
(ii) *Silent Sensing*: $\mathbf{X}^0 = \{\{s_0\}, \{s_1, \ldots, s_N\}\}$;
$\mathbf{X}^1 = \{\{s_0\}, \{s_1\}, \{s_2, \ldots, s_N\}\}$
(iii) *Complete Sensing*: $\mathbf{X}^0 = \mathbf{X}^1 = \{\{s_0\}, \{s_1\}, \ldots, \{s_N\}\}$

Our focus is on No Silent Sensing; much of the literature focuses on Silent Sensing; Complete Sensing represents the most information that could possibly be observed (assuming stations are anonymous to each other).

B. Protocols

At each time, each station must choose an action 0 (remain silent) or 1 (transmit); a *protocol* is a set of instructions for making such choices. We allow choices to be random and to depend on the personal history of the station. To make this precise requires some definitions. A *channel event* is a pair $(\mathbf{a}, \sigma(\mathbf{a}))$, specifying an action profile and the corresponding channel state; write $\mathcal{C} \subset A^N \times S$ for the set of channel events. A *channel history of length* $T \geq 0$ is an element of $\mathcal{C}^T = \mathcal{H}_c(T)$, specifying the actions taken and the resulting channel state at each time $t = 0, \ldots, T$; we write $\mathcal{H}_c = \bigcup \mathcal{H}_c(T)$ for the set of all finite channel histories. An *infinite channel history* is an element of $\mathcal{C}^\infty = \mathcal{H}_c(\infty)$, specifying the actions taken and the resulting channel state at each time $t = 0, 1, \ldots$.[6] A *personal event* is a pair (a, X) consisting of an action a and a set X in the corresponding partition $\mathbf{X}^a$, and specifying the action taken and the observation made. The set of personal events can be described as

$$\mathcal{P} = \{(a, X) \in A \times (\mathbf{X}^0 \cup \mathbf{X}^1) : X \in \mathbf{X}^a\}$$

A *personal history of length* $T \geq 0$ is an element of $\mathcal{H}_p(T) = \mathcal{P}^T$, specifying the action taken and the observation made at times $t = 0, \ldots, T$. Note that the empty history is the unique channel history or personal history of length 0. For each n we define a map $\Pi_n : \mathcal{C} \to \mathcal{P}$ from channel events to personal events sby

$$\Pi_n((a_1, \ldots, a_N), s) = (a_n, \mathcal{X}^{a_n}(s))$$

[5] In fact, if we consider duration rather than number of slots, Table I underestimates the speed of our protocols because our protocols do not require full slots. We will discuss this is more detail later.

[6] Since the channel state depends deterministically on the actions, specifying both the vector of actions and the channel state is redundant – but convenient.

Thus Π_n specifies the action taken and the observation made by station n in the given channel event. We abuse notation and continue to write $\Pi_n : \mathcal{H}_c(T) \to \mathcal{H}_p(T)$ for the mapping from channel histories to personal histories that applies Π_n separately to each channel event in the history.

A *protocol* is a set of instructions for a station as a function of the personal history of that station. We allow for randomization, so a protocol is a mapping $f : \mathcal{H}_c \to \Delta(\{0,1\})$ (the set of probability distributions on the set $\{0,1\}$). We write $f(a|h)$ for the probability that action a is chosen following the history h. Because stations are *anonymous*, we do not allow the protocol to depend on the name of the station and we require that all stations use the same protocol. However, because *stations randomize independently*, stations that are *ex ante* identical may experience different realizations and so may become distinguished *ex post*; in particular, stations may choose different actions and experience different histories, even though they are following the same protocol. A protocol f determines, for each $T \geq 0$, a probability measure μ_T on $\mathcal{H}_c(T)$; $\mu_T(h)$ is the probability that the channel history $h \in \mathcal{H}_c(T)$ occur when that the stations follow f.[7]

IV. Impossibility of Deterministic Perfect Coordination

As noted in the Introduction, the protocols we propose (and others used previously) all rely on randomization and so have the property that the probability of convergence to perfect coordination in any given finite number of slots is strictly less than one. This is not an accident: *there are no protocols for which the probability of convergence to perfect coordination in any given finite number of slots is one*. Indeed, given a number $N \geq 2$ of stations and a finite horizon T^*, we find a bound $b < 1$ (independent of the information condition – even Complete Sensing) with the probability that no protocol can achieve perfect coordination in T^* slots or less with probability greater than b. Indeed, no protocol can achieve successful transmission by each station with zero collisions in T^* slots or less with probability greater than b.

Theorem 1: (Impossibility) Given the number of stations $N \geq 2$, a finite horizon T^* and a protocol f, the probability that successful transmission occurs in slot T^* is no greater than $1 - (2^{-N+1})^{T^*+1}$. In particular, the probability that f converges to perfect coordination or achieves successful transmission by each station in T^* slots or less is no greater than $1 - (2^{-N+1})^{T^*+1}$.

Proof: Set

$$A' = \{(a_1, \ldots, a_n) \in A^N : a_i = a_j \text{ for all } i, j\}$$

For each $T \geq 0$ write $\mathcal{H}_c{}'(T) = (A' \times S)^T$; A' is the set of action profiles in which all stations take the same action and $\mathcal{H}_c{}'(T)$ is the set of channel histories of length T for which all stations take the same action in each slot. Fix a protocol f and as above let μ_T be the probability measure on $\mathcal{H}_c(T)$ induced by f. We first establish the following CLAIM:

$$\mu_T\left(\mathcal{H}_c{}'(T)\right) \geq (2^{-N+1})^T \text{ for each } T \geq 0$$

To prove this, we proceed by induction on T. The assertion is true for $T = 0$ since the unique history of length 0 is the empty history, which belongs to $\mathcal{H}_c{}'(0)$. Suppose that the assertion is true for $T = T_0$; we must show it is true for $T = T_0 + 1$. If $h \in \mathcal{H}_c{}'(T_0)$ then $\Pi_m(h) = \Pi_n(h)$ for all stations m, n, so all stations follow the same probability distribution over actions in slot T_0+1; say p is the common probability of transmitting. Since there are N stations, the probability that they all transmit is p^N and the probability that they are all silent is $(1-p)^N$, so the probability that they all choose the *same* action is $p^N + (1-p)^N$. This probability is minimized if $p = 1/2$; the minimum is 2^{-N+1}. Hence, conditional on the history up to slot T_0 belonging to $\mathcal{H}_c{}'(T_0)$, the probability that the history through slot $T_0 + 1$ belongs to $\mathcal{H}_c{}'(T_0 + 1)$ is at least 2^{-N+1}. The inductive assumption is that $\mu_{T_0}\left(\mathcal{H}_c{}'(T_0)\right) \geq (2^{-N+1})^{T_0}$ so it follows that $\mu_{T_0+1}\left(\mathcal{H}_c{}'(T_0+1)\right) \geq (2^{-N+1})^{T_0+1}$ which is the CLAIM. To derive the asserted inequality note that if $h \in \mathcal{H}_c{}'(T^* + 1)$ then all stations take the same action in slot $T^* + 1$ so either they are all silent or there is a collision; in either case, successful transmission does not occur. This completes the proof. ■

V. Perfect Coordination Protocols

In this section we describe our proposed protocols. For clarity of exposition, we present descriptions in words; formal flow charts are in the Appendix.

A. N is known

For a known number of stations N, we propose a family $\Phi(N, K, p)$ of Perfect Coordination protocols, depending on a length parameter K and a vector p of probabilities (strictly between $0, 1$) to be chosen by the administrator. (Different choices of K, p lead to different probabilities and speeds of convergence.) Each protocol is divided into *phases*, each phase is divided into *cycles*, each cycle is divided into *slots*. In the *Learning* phase, stations learn their place in an endogenously determined sequencing of all the stations; in the *Transmission* phase, stations transmit in the sequence determined in the first stage.

By definition a protocol specifies random actions conditional on personal history; however $\Phi(N, K, p)$ uses only some of the information contained in those histories (in addition to the parameters K, p):

- the current phase;
- the location of the current cycle within the current phase;
- the location of the current slot within the current cycle;
- personal information about the previous slot;
- a summary statistic of personal history: the station's *index*, an integer between 0 and N.

A station's index is a convenient way for the station to keep track of whether it has at some point won a lottery (as

[7] To define μ_T rigorously, recall first that $\mathcal{H}_c(0)$ consists of the single empty history; require μ_0 to assign probability 1 to the empty history. Now proceed inductively. Suppose $\mu_0, \ldots, \mu_{T-1}$ have been defined. For $h \in \mathcal{H}_c(T)$ write $h = (h', ((a_1, \ldots, a_N), s))$ where $h' \in \mathcal{H}_c(T-1)$ and $((a_1, \ldots, a_N), s)) \in \mathcal{C}$; define

$$\mu_T(h) = \mu_{T-1}(h') \times f(a_1|\Pi_1(h')) \times \cdots \times f(a_N|\Pi_N(h'))$$

.

described below) and, if so, when. Each station knows only its *own* index. At any slot, a station's index will depend on its index in the previous slot and the station's personal observation about the previous slot. We initialize so that *index*$(z) = 0$ for all stations z.

Fix a positive integer K and an N-vector $p = (p_1, \ldots, p_N)$ of probabilities strictly between $0, 1$. The protocol $\Phi(N, K, p)$ begins in the Learning phase.

- In slot 1, all stations randomize: they transmit with probability p_1 and remain silent with the complementary probability $(1 - p_1)$. This randomization creates an *endogenous lottery*.
- In slot 2 each station conditions on what it observes about slot 1 (its personal history). A station z that transmitted in slot 1 and was successful – *won the lottery* – sets its *index*$(z) = 1$ and transmits with probability 1 in the second slot (and in every succeeding slot in the current cycle); a station that did transmit but was unsuccessful – did not win the lottery – or did not transmit (and hence observed nothing) randomizes again in slot 2, with the same probabilities $p_1, 1 - p_1$.
- The process continues through slot K (one cycle of the Learning phase).
- At the end of the first through N-th cycles of the Learning phase, the protocol repeats the process above, with three changes:
  - stations that won a lottery (those with *index*$(z) > 0$) remain silent (inactive) throughout the current cycle;
  - stations that have not won a lottery (those with *index*$(z) = 0$) randomize with probabilities $p_n, 1 - p_n$;
  - a station z that wins the lottery in cycle n sets *index*$(z) = n$[8]

  At any point in time, *index*$(z) > 0$ if and only if station z won the lottery in cycle *index*(z) of the Learning phase; if *index*$(z) = 0$ then station z has not won a lottery in any cycle.
- At the end of N cycles of the Learning phase, stations enter the Transmission phase. If *index*$(z) = n > 0$ then station z transmits in slot n and only in slot n; if *index*$(z) = 0$, then station z transmits in *every* slot of the Transmission phase.
- After N slots (one cycle) of the Transmission phase *all* stations will have had the *same* experience: either at least one collision or success whenever transmitting. If all stations experienced at least one collision, then stations are not perfectly coordinated; in that case, all stations set *index* $= 0$ and return to the beginning of the protocol. If all stations experienced success whenever transmitting, then stations are perfectly coordination; in that case, all stations return to the beginning of the Transmission phase and repeat the Transmission phase indefinitely.

(For flow charts, see Figures 5 and 6 in the Appendix.)

[8]A station that wins the lottery in some slot during cycle n transmits in every succeeding slot in that cycle, so at most one station can transmit successfully in any given cycle and so no two stations can have the same index.

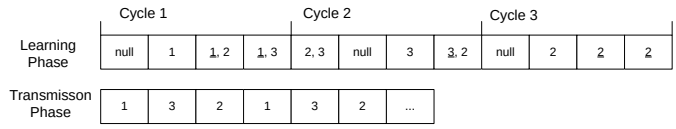

Fig. 1. Illustration of PC protocol when N is known. ($N = 3, K = 4$)

Both individual and cooperative learning play important roles in this protocol. In the Learning phase, stations are creating an endogenous lottery: creating the lottery is cooperative (although it requires no communication, only conformity to the protocol); winning the lottery is individualistic. A station that wins the lottery in the current cycle transmits in every remaining slot of the current cycle, and stations that have won lotteries in previous cycles do not participate in the current cycle; both of these cooperative activities avoid interference with the learning of other stations. Learning of a different sort occurs in the Transmission phase: all stations learn either that perfect coordination has been achieved, so that they should remain in the Transmission phase indefinitely, or that perfect coordination has not been achieved, so that they should re-initialize and begin the Learning phase again.

Figure 1 illustrates the operation of this protocol with $N = 3$, $K = 4$. Station 1 wins the lottery in Cycle 1, Slot 2; Station 3 wins the lottery in Cycle 2, Slot 3; Station 2 wins the lottery in Cycle 3, Slot 2.

1) The Probability of Convergence: It is straightforward to calculate the probability that perfect coordination is achieved in one round (N cycles) of the Learning phase. From this we can calculate the probability that the protocol converges to perfect coordination within a given number of rounds and hence in a given number of slots. (Translation from rounds into slots is straightforward: Every station begins in the Learning Phase, remains there for N cycles – one *round* – exits to the Transmission Phase and remains there for N slots (one cycle), then either remains in the Transmission Phase indefinitely or starts the Learning Phase over again; one round of the Learning Phase requires NK slots and and one round of the Transmission phase requires N slots so if perfect coordination is achieved after R rounds of the Learning Phase then $RNK + (R - 1)N$ slots will have elapsed.) As a consequence, it follows that, for every N and every choice of K, p, convergence to perfect coordination in finite time occurs with probability 1; i.e., for each $\varepsilon > 0$ there is some T so that the probability of convergence to perfect coordination in T slots or less is at least $1 - \varepsilon$ (Corollary 1).

Theorem 2: Fix $N, K, p = (p_1, \ldots, p_N)$ (with $0 < p_n < 1$ for all n). The probability that the protocol $\Phi(N, K, p)$ achieves perfect coordination in no more than R rounds (no more than $RNK + (R - 1)N$ slots) is exactly

$$1 - \left(1 - \prod_{n=1}^{N} \left\{1 - [1 - (N - n + 1)p_n(1 - p_n)^{N-n}]^K\right\}\right)^R \quad (2)$$

Proof: By construction, the protocol $\Phi(N, K, p)$ achieves perfect coordination in a given round of the Learning phase and only if one station (necessarily unique and distinct) wins the lottery in each of the N cycles in that round, so we first calculate the probability π that this occurs. Fix a cycle n

of the Learning phase; consider the k-th slot in this cycle. Assume that there has been a winner in each of the previous cycles $1, \ldots, (n-1)$ but that there has not been a winner in previous slots of the current cycle n. By definition, there will be a winner in the current slot if and only if exactly one station transmits. Since $n-1$ stations were winners in previous cycles, exactly $N-n+1$ stations are randomizing in the current cycle n. By construction, each of these stations transmits with probability p_n so the probability that exactly one of them transmits, and hence wins the lottery in this particular slot, is $(N-n+1)p_n(1-p_n)^{N-n}$ and the complementary probability, that no station wins the lottery in this slot, is $1-(N-n+1)p_n(1-p_n)^{N-n}$. Hence the probability that *no station* wins the lottery in cycle n is $[1-(N-n+1)p_n(1-p_n)^{N-n}]^K$ and the probability that *some station* wins the lottery in cycle n is the complementary probability $1-[1-(N-n+1)p_n(1-p_n)^{N-n}]^K$. (Recall that if some station transmits successfully in some slot during a cycle, then the same station will transmit for the remainder of that cycle, so at most one station will be assigned an index in a given cycle.) Thus, the probability that some station wins the lottery in *every cycle*, and hence that perfect coordination is achieved after this particular round of the Learning phase is

$$\pi = \prod_{n=1}^{N} \left\{ 1 - [1-(N-n+1)p_n(1-p_n)^{N-n}]^K \right\} \quad (3)$$

Perfect coordination will *fail* to be achieved within the first R rounds of the Learning Phase if it fails to be achieved in *each* of the first R rounds; the probability that this occurs is $(1-\pi)^R$, and the probability that perfect coordination will be achieved in no more than R rounds is the complementary probability $1-(1-\pi)^R$. Substituting into equation (3) yields the desired expression (2). As we have noted, if perfect coordination is achieved after R rounds of the Learning Phase, then $R-1$ rounds of the Transmission Phase will have occurred; the R-th time the protocol enters the Transmission Phase, perfect coordination is already achieved, and the protocol remains in the Transmission Phase forever. This completes the proof. ■

Corollary 1: For every N and every choice of $K, p = (p_1, \ldots, p_N)$ (with $0 < p_n < 1$ each n), the protocol $\Phi(N, K, p)$ converges to perfect coordination in finite time with probability 1.

Proof: For fixed N, K, p the probability in (2) tends to ∞ with R. ■

2) A Numerical Estimate: The expression (2) is somewhat intractable and does not provide a simple rule for an an administrator who wishes to guarantee a given probability of convergence. However a judicious choice of the probability vector p and some simple manipulation yields (Theorem 3) both a tractable estimate and a useful procedure for choosing K, R to guarantee a given probability of convergence. We first isolate two simple lemmas.

Lemma 1: For $x > 1$, $1-\left(\frac{x-1}{x}\right)^{x-1}$ is a decreasing function of x and

$$\lim_{x\to\infty} \left(\frac{x-1}{x}\right)^{x-1} = \frac{1}{e}$$

Proof: Take logarithms and differentiate; then use the fact that log is a concave function to derive the first assertion and L'Hopital's rule to derive the second assertion. ■

Lemma 2: For ℓ a strictly positive integer

$$\max_{0<p<1} \ell p(1-p)^{\ell-1} = \left(\frac{\ell-1}{\ell}\right)^{\ell-1} = \left(1-\frac{1}{\ell}\right)^{\ell-1}$$

Moreover, the maximum is attained when $p = 1/\ell$.

Proof: Apply the the first-order condition. ■

Theorem 3: Fix N and set $p^* = (p_1^*, \ldots, p_N^*)$, $p_n^* = \frac{1}{N-n+1}$.

(i) The probability that the protocol $\Phi(N, K, p)$ converges to perfect coordination in at most R rounds of the learning phase ($RNK + (R-1)N$ slots) is at least

$$1 - \left(1 - \left\{1 - \left[1 - \left(\frac{1}{e}\right)\right]^K\right\}^N\right)^R$$

(ii) Fix $\varepsilon > 0$. In order that $\Phi(N, K, p)$ converge to perfect coordination in at most R rounds of the learning phase with probability at least $1-\varepsilon$ it is sufficient that

$$K \geq -\left(\frac{\log N - (1/R)\log\varepsilon}{\log[1-(1/e)]}\right)$$

Proof: (i) Substituting the specific probability vector p^* into (2), using the estimate from Lemma 2 and simplifying, shows that the probability of convergence in one round of the Learning Phase is at least $\beta = \left\{1 - [1-(1/e)]^K\right\}^N$. Hence the probability of *failure* of convergence in one round of the Learning Phase is no more than $\gamma = 1-\beta$, the probability of failure in each of R rounds is no more than γ^R and the probability of success within at most R rounds is at least $1-\gamma^R = 1-(1-\beta)^R$, which is the assertion.

(ii) Fix $\varepsilon > 0$. We want to choose K, R so that the probability of failure in each of R rounds is at most ε. Maintaining the notation of (i), this means we want to choose K, R so that $\gamma^R \leq \varepsilon$, which is to say $\gamma \leq \varepsilon^{1/R}$. Observe that for $-1 < x < 1$, the function $(1-x)^N$ is convex and so is everywhere above its tangent line at $x = 0$; i.e., $(1-x)^N \geq 1 - Nx$. Applying this observation to $\gamma = 1-\beta$ yields

$$\gamma = 1-\beta \leq 1 - \left(1 - N[1-(1/e)]^K\right) = N[1-(1/e)]^K$$

Hence, it suffices to choose K, R so that $N[1-(1/e)]^K \leq \varepsilon^{1/R}$. Taking logarithms and doing the requisite algebra yields the assertion. ■

To make Theorem 3 more concrete, take $N = 10$ and $\varepsilon = 10^{-4}$ so that we seek convergence with probability at least 0.9999. In view of (ii), it is sufficient to take $R = 1, K = 25$ or $R = 2, K = 15$. The former choice would seem to be superior to the latter (250 slots vs. 310 slots) but note that if $K = 15$ then convergence actually occurs in one round (150 slots) with probability greater than 0.99 – so using the latter choice leads to a a much smaller *expected* time to convergence.

As the simulations in Section VI will make clear, these analytic estimates are not sharp; but they do have the virtue of simplicity, and the rate of convergence they provide is already quite good (especially when interpreted in terms of time rather than number of slots).

B. N is Unknown

Now we assume that the true number N of stations is unknown but that an upper bound on the number of stations $N^{\max}$ is known. We could apply any of the protocols $\Phi(N^{\max}, K, p)$ in this setting, but the result might not be very satisfactory: If $N < N^{\max}$ then $\Phi(N^{\max}, K, p)$ will eventually reach the situation in which the N stations transmit in turn but in a cycle of length $N^{\max}$, so that there will be gaps and throughput will be less than optimal; moreover since $\Phi(N^{\max}, K, p)$ executes at least $N^{\max}$ cycles in the Learning phase, the protocol might take much longer than is necessary – much longer than necessary if N is much smaller than $N^{\max}$. We avoid both of these problems by modifying the Learning phase slightly and adding two intermediate phases so that the protocol moves into the Transmission phase as soon as all stations have learned their indices and closes gaps along the way.

As before, our protocols depend on $N^{\max}$ and on parameters that can be chosen by the administrator: a length parameter K (a positive integer) and a vector of probabilities parameters $q = (q_1, \ldots, q_{N^{\max}})$, each q_m strictly between 0, 1. The protocol $\Psi(N^{\max}, K, q)$ consists of four phases: *Learning-to-Win*, *Rectifying-the-Count*, *Learning-the-Losers*, *Coordinated Transmission*.

As before, $\Psi(N^{\max}, K, q)$ uses only some of the information contained in personal histories (in addition to K, q):

- the current phase;
- the location of the current cycle within the current phase;
- the location of the current slot within the current cycle;
- the station's personal information about the previous slot;
- two summary statistics of the station's personal history: *stationcount* (an integer between 0 and $N^{\max}$), which will tell the station when to move to the Coordinated Transmission phase; and *index* (either $*$ or an integer between 0 and $N^{\max}$), which will tell the station in which slot of the Coordinated Transmission phase to transmit).

Each station knows its *own* statistics – *stationcount* and *index*. At any slot, a station's *stationcount* and *index* will depend on the values in the previous slot and the station's observation of the previous slot. It is convenient to refer to a station z as a *winner* if *index*$(z) > 0$, a *loser* if *index*$(z) = 0$ and as a *waiter* if *index*$(z) = *$. Winners are stations that have won a lottery *and* learned their indexed, waiters are stations that have won a lottery but have not yet learned their index, losers are stations that have not yet won a lottery. At various times, some of these categories will be empty (e.g., at the beginning of the protocol all stations are losers); at any moment in time there is at most one waiter. Because stations always know their own statistics, they know to which category they belong. The protocol begins in the Learning-to-Win phase.

- In the Learning-to-Win phase, waiters and winners remain silent throughout.[9] Losers randomize: if this is the M-th time the protocol has entered the Learning-to-Win phase, Losers transmit with probability q_m and remain silent with the complementary probability $(1-q_m)$, where $m = \min\{N^{\max}, M\}$. (Stations can compute m without computing M because they only need to count as high as $m = N^{\max}$ – and then stop counting.) As before, this creates an endogenous lottery.
- A station z that wins the lottery in some slot in the current Learning-to-Win cycle sets *index*$(z) = *$ and transmits in every subsequent slot; losers continue to randomize.
- The process continues through slot K (one cycle of Learning-to-Win).
- At the end of each cycle of Learning-to-Win, the protocol enters the Rectifying-the-Count phase which consists of m slots. In each slot: losers remain silent; the waiter (if one exists) transmits in every slot; winners transmit in the slot corresponding to their index. An inductive argument shows that if there are $W < m$ winners at this stage then their indices are $1, \ldots, W$. Hence if there is a waiter, all winners experience a collision and the waiter experiences collisions in slots $1, \ldots, W$ and success in slot $W + 1$. The waiter sets its *index* equal to $W + 1$; the waiter and all winners set their *stationcount* equal to $W + 1$. Note that at this point $W + 1$ is indeed the *correct* number of winners.
- At the end of one cycle of Rectifying-the-Count, the protocol enters the Learning-the-Losers phase which again consists of m slots. In every slot: losers transmit in every slot; winners transmit in the slot corresponding to their index. (There are no waiters at this point.) If there are no losers, the winners all experience success; if there are losers then both winners and losers experience collision. Hence after this cycle, all stations know whether or not losers remain.
- If no losers remain after Rectifying-the-Count and Learning-the-Losers, then every station moves to the Coordinated Transmission Phase and all stations transmit in sequence according to their indices, in a cycle of length *stationcount* (w) for any station w. This is perfect coordination and repeats indefinitely. If some loser or losers remain after Rectifying-the-Count and Learning-the-Losers then all stations return to Learning-to-Win and proceed as above.

(For formal flow charts, see Figures 7, 8, 9 and 10 in the Appendix.)

Again, both individual and cooperative learning play important roles here. In the Learning-to-Win phase there is little to add to what we have already said about the Learning phase in the protocol $\Phi(N, K, p)$. In the Rectifying-the-Count phase, winners have learned the current total number of winners and their own place in the sequence; now winners and the waiter (if there is one) cooperate so that all can learn whether there is a waiter, in which case the waiter becomes a winner, all winners learn the new total number of winners, and all winners learn

[9] In fact there will never be a waiter at this point.

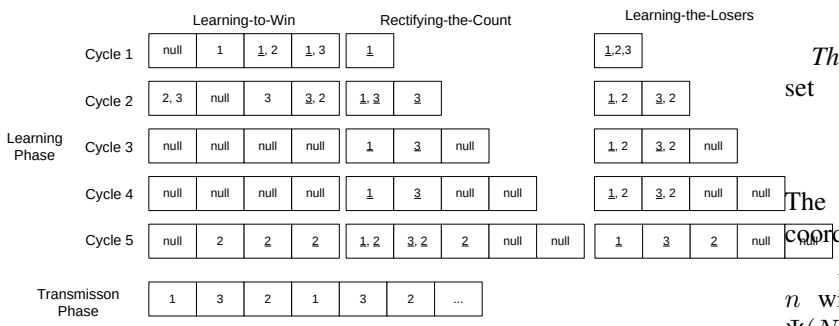

Fig. 2. Illustration of PC protocol when N is known. ($N = 3, K = 4, N^{max} = 4$)

their places in the sequence. Note that waiter can count the number of current winners because each of them transmits in turn, so that the waiter experiences collisions, but then the current winners are silent, so that the waiter experiences success. In the Learning-the-Losers phase, all stations learn whether there are remaining losers; if so they return to the Learning-to-Win phase, if not they transit to the Transmission phase.

Figure 2 provides an illustration of the operation of the protocol, with $N = 3$ and $N^{\max} = 4$ (so that each cycle of the Learning-to-Win phase consists of $K = 4$ slots). Station 1 wins the lottery in Cycle 1 Slot 2; Station 3 wins the lottery in Cycle 2 Slot 3; there are no winners in Cycles 3, 4; Station 2 wins the lottery in Cycle 3 Slot 2. The winners transmit in their corresponding slots in the Rectifying-the-Count and Learning-the-Losers phases.

1) The Probability of Convergence: At this point it is natural to ask for a parallel to Theorem 2 that provides, for each K, q, the precise probability of convergence to perfect coordination in a specified number of slots. However, a closed form expression for the precise probability seems difficult to provide. Instead, Theorem 3 provides convergence estimates for the particular vector of probabilities $q^* = (q_1^*, \ldots, q_{N^{\max}}^*)$ defined by

$$q_1^* = \frac{1}{N^{\max}}, q_2^* = \frac{1}{(N^{\max} - 1)}, \ldots, q_{N^{\max}-1}^* = q_{N^{\max}}^* = \frac{1}{2}$$

As our simulations will show these estimates are not at all sharp (much worse than the estimates from Theorem 2 when N is known), but they are good enough to establish (Corollary 2) convergence in finite time with probability 1. To state the estimate formally, write $B(N, R; \zeta)$ for the probability of getting at least N successes in $R \geq N$ independent trials each of which has probability of success ζ. (As before, we can we can translate rounds into slots although the translation is more complicated. Learning-to-Win requires K slots in every round; if the current round is r then Rectifying-the-Count and Learning-the-Losers each require $\min\{r, N^{\max}\}$ slots. Hence if $R \leq N^{\max}$ the total number of slots that have occurred in R rounds of the first three phases is $\sum_{r=1}^{R}(K + r + r) = RK + R(R+1)$; if $R > N^{\max}$ then – because there are $R - N^{\max}$ rounds in which Rectifying-the-Count and Learning-the-Losers require $N^{\max}$ slots – the total number of slots is $N^{\max}K + N^{\max}(N^{\max} + 1) + (R - N^{\max})(K + 2N^{\max})$.)

Theorem 4: If N is the true number of stations and $R \geq N$; set

$$\zeta = 1 - \left[1 - \left(\frac{1}{N^{\max} - N + 1}\right)\left(\frac{1}{e}\right)\right]^K$$

The probability that $\Psi(N^{\max}, K, q^*)$ converges to perfect coordination in no more than R rounds is at least $B(N, R; \zeta)$.

Proof: Fix $r \geq 0$ and $n < N$, suppose that exactly n winners have been identified in the first r rounds of $\Psi(N^{\max}, K, q^*)$, and consider what happens in the Learning-to-Win cycle of round $r + 1$. By construction,there are $N - n$ losers randomizing in this cycle, each transmitting with probability $q^*_{(N^{\max}-n)}$ and remaining silent with the complementary probability. The probability that there is a winner in one particular slot in this cycle is exactly

$$\pi(n+1) = (N - n)\left(\frac{1}{N^{\max} - n}\right)\left(1 - \frac{1}{N^{\max} - n}\right)^{N-n-1}$$

(Notice that this expression depends on n but not on r.) As in the proof of Theorem 2, this yields the inequality

$$\pi(n + 1) \geq \left(\frac{1}{N^{\max} - N + 1}\right)\left(\frac{1}{e}\right) = \eta$$

Hence the probability that there will be a winner in the current cycle is at least $1 - (1 - \eta)^K = \zeta$; note that this estimate is independent of r and n.

Perfect coordination will be achieved as soon as N cycles of Learning-to-Win have produced a winner (and the succeeding Rectifying-the-Count and Learning-the-Losers phases have concluded). If the probability that a winner would be produced in a given round were independent of history and equal to ζ, the probability that N winners would be produced in at most R cycles would be at least $B(N, R; \zeta)$ – the probability of at least N successes in R trials with success probability ζ. In fact, the probability of a winner occurring in a given cycle is *not* independent of history, but in view of the construction above, the probability that a winner occurs in any given cycle is *at least* ζ, so the probability that N winners are produced in at most R cycles is at least $B(N, R; \zeta)$; this completes the proof. ■

Corollary 2: For every $N, N^{\max}$ and every choice of K, the protocol $\Psi(N, K, q^*)$ converges to perfect coordination in finite time with probability 1.

Proof: For all N and all $\zeta > 0$ the probability $B(N, R; \zeta)$ tends to ∞ with R. ■

We caution the reader that the estimate in Theorem 4 is very crude and that simulations show that the actual rate and probability of convergence is much much better. For instance, suppose $N = 5, N^{\max} = 10, K = 20$. Theorem 4 estimates that the probability of convergence to perfect coordination in 10 rounds is at least .96 – but simulations show that the true probability of convergence to perfect coordination is greater than .9999.

2) Entry and exit: To this point we have assumed that the number of stations, although unknown, is fixed – no new stations enter the network and no old stations leave the network. In many settings this is not realistic, but the protocols $\Psi(N^{\max}, K, q*)$ are easily adapted to allow for entry and exit,

with little loss of performance, so long as entry and exit are not too frequent. To do this we use the familiar idea of a *superframe*.

At the beginning of each superframe, active stations are coordinated by following the PC protocol $\Psi(N^{\max}, K, q^*)$. After perfect coordination is achieved, stations transmit in a perfectly coordinated fashion for the remainder of the superframe. During the superframe, stations can exit but new stations cannot enter; at the end of the superframe, stations may exit and other stations may enter. Following the end of the superframe, the newly active stations (which may include some or all of the previously active stations) are again coordinated by following $\Psi(N^{\max}, K, q^*)$. (All superframes have the same duration so stations know when a new superframe starts and they can enter.) Because the PC protocol achieves perfect coordination, this results in transmissions without collisions; if stations exit only at the end of a superframe this also results in transmissions without gaps. Even if stations exit during the superframe, the PC protocol leaves fewer gaps than [8] or [9] and hence is more efficient.

C. Robustness

All protocols have some sensitivity to errors. The most obvious errors in the context we consider are the failure to detect an idle channel (not distinguishing a channel that is busy from a channel that is idle but noisy), and the failure to recognize a successful transmission (loss of acknowledgement). Because our protocols do not require stations that do not transmit to observe anything, our protocols are completely immune to the first of these errors – unlike other protocols that depend heavily on detection of idle channels. Failure to recognize a successful transmission in a given slot during the learning phases of either Φ or Ψ would have exactly the same effect as an unsuccessful transmission in that slot, so a small probability of loss of acknowledgement simply raises the probability of an unsuccessful transmission in a single slot by the same small amount. Failure to recognize a successful transmission in a given slot during the Transmission phase of Φ could result in confusion: some stations would experience success and believe perfect coordination had been achieved, hence repeat the transmission phase, while at least one station would experience collision and believe perfect coordination had not been achieved, hence return to the learning phase. However, the probability that confusion occurs could be made arbitrarily small simply by repeating the Transmission phase several times; confusion would occur only if some station experienced errors in each repetition. If loss of acknowledgement occurs with probability δ then the probability that confusion would occur in a single cycle of the Transmission phase would be $1-(1-\delta)^N$ but k repetitions would reduce the probability to $1-(1-\delta^k)^N$ – and k repetitions require only kN slots. For instance, if $\delta = 10^{-2}$ and $N = 10$, $k = 3$ repetitions would reduce the probability of confusion below 10^{-5} at the cost of only 30 slots – less than 7ms. Repetition would be similarly effective with similarly low cost in the Rectifying-the-Count and Learning-the-Losers phases of the protocol Ψ.

VI. Simulations

In this section, we provide simulation results to evaluate the performance of the proposed Perfect Coordination protocols and provide comparisons with existing protocols. For these simulations, we adopt the parameters specified by IEEE 802.11a, as in Table II.

TABLE II
Simulation configuration

Parameters	Values
Packet payload	1024 octets
MAC header	28 octets
ACK frame size	14 octets
Data rate	54 Mbps
PHY header time	20 μs
SIFS	16 μs
DIFS	34 μs

A. Impact of knowing the number of stations

Table III documents the estimated and simulated speed of convergence of our protocols for $N = 4, 8, 16, 24, 32$: the number of slots required for the stations to achieve perfect coordination with pre-specified probabilities $0.99, 0.999, 0.9999$. For N known, we record the estimate implied by Theorem 3; for N unknown we assume $N^{\max} = 32$ and use the estimate implied by Theorem 4. In both cases the simulation results are generated from 10^4 Monte Carlo tests.

There are two reasons why convergence is slower when the number of stations N is unknown. As we have noted earlier, when N is unknown, the protocol must go through the additional phases of *Rectifying-the-Count* and *Learning-the-Losers* at least N times; this requires at least $N \times (N+1)$ slots, which is substantial when N is large. More subtly, and more importantly, when N is not known we can tailor the probability parameter to the upper bound $N^{\max}$ – but we cannot tailor it to N itself.

However, we note that actual convergence as measured by *time* – rather than number of slots – is very fast even if the number of slots is substantial, especially since during the coordination phase(s) we can use small slots in which stations only send a small payload (e.g. 100 bytes) in each slot. Allowing for packet overhead and signaling intervals, these small slots can be as short as 90 μs. (After perfect coordination is achieved, stations can revert to regular slots, typically 230 μs.) If we use these small slots, the actual simulated time to convergence in the worst case illustrated in Table III is only a few tenths of a second.

B. The path to convergence

By construction, the protocol $\Psi(N^{\max}, K, q^*)$ runs through the first three phases over and over until all stations have won a lottery. If the true number of stations is N, this necessarily takes at least N rounds but typically takes more. Figure 3 illustrates this point in a simple way by showing sample paths of the number of winners as a function of the number of rounds in an environment with $N = 8$ and $N^{\max} = 32$. (For example, Sample Path 1 shows that in each of the first 6 rounds, there was a winner, for the the next 4 rounds there was no winner, etc.)

N is Known						
Prob.		N = 4	N = 8	N = 16	N = 24	N = 32
0.99	Sim.	40	104	240	384	544
	Est.	56	120	272	408	576
0.999	Sim	56	128	320	456	704
	Est.	72	160	336	528	736
0.9999	Sim.	64	160	368	552	768
	Est.	92	192	416	648	864
N is Unknown (N^{max} = 32)						
Prob.		N = 4	N = 8	N = 16	N = 24	N = 32
0.99	Sim.	288	442	740	1053	1483
	Est.	1224	1976	2850	2704	1516
0.999	Sim.	380	532	836	1120	1582
	Est.	1656	2511	3302	2998	1592
0.9999	Sim.	456	588	874	1176	1592
	Est.	2106	2998	3682	3226	1626

TABLE III
NUMBER OF SLOTS FOR PERFECT COORDINATION WITH PROBABILITIES .99, .999, .9999

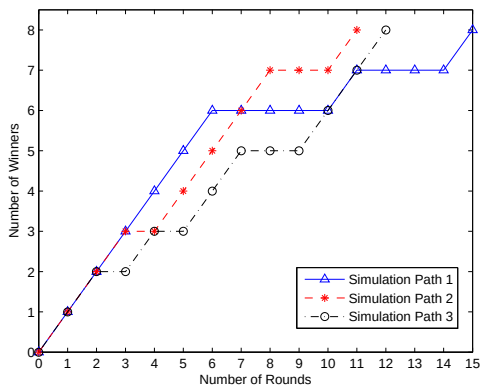

Fig. 3. Sample paths of the PC protocol

C. Goodput comparison

Assuming $N^{\max}$ is known but N is unknown, we compare the goodput of our proposed Perfect Coordination protocol against the Zero Collision protocol (ZC) proposed in [8]. As above, we assume that our Perfect Coordination protocol uses small slots of 90 μs in the coordination phases and regular slots of 230 μs (including a 1024-Byte Payload and overhead) in the Transmission phase. To be generous to ZC, we assume it uses idle slots of 34 μs (i.e. a DIFS duration with a 9μs empty slot) when no stations are transmitting and regular slots of 230 μs when some station is transmitting. In both cases we assume the upper bound on the number of stations is $N^{\max} = 32$, which means that in ZC the contention window size (the number of slots in a cycle) is also $M = 32$.

In this setting we carry out two simulation exercises (using 10^3 Monte Carlo tests).

- Table IV shows goodput achieved by PC, ZC and L-ZC (the modification of ZC suggested by [9]) as a function of the true number of stations N. For each protocol, we show goodput in Mbps and the fraction of theoretically optimal goodput (which in this case is 35.94 Mbps, after taking overhead into account). PC reaches perfect coordination rapidly and achieves optimal goodput independently of N; ZC and L-ZC both leave gaps, but leave fewer gaps – use the available slots more efficiently, hence achieve greater goodput – when N is large than when N is small. (If the size of the contention window is $M = N^{\max} = 32$ and the true number of stations is N, both ZC and L-ZC leave $M - N$ gaps.)
- Table IV compares goodput after coordination has been achieved, but ignores the time required to coordinate. Table V compares goodput *including the coordination period* for T = 0.1, 0.2, 0.3, 0.4, 0.5, 1.0, 2.0, 5.0, 10.0 seconds (for $N = 8, N^{\max} = 32$). PC requires a slightly longer coordination time than ZC and L-ZC (in the average simulation, roughly 50ms for PC, 30 ms for ZC and L-ZC); this leads to lower goodput for the first 0.1 seconds, but PC outperforms both ZC and L-ZC within 0.2 seconds.

		N = 4	N = 8	N = 16	N = 24	N = 32
ZC/L-ZC	Mbps	17.58	24.82	31.28	34.24	35.94
	Fraction	0.49	0.69	0.87	0.95	1
PC	Mbps	35.94	35.94	35.94	35.94	35.94
	Fraction	1	1	1	1	1

TABLE IV
GOODPUT COMPARISON FOR VARIOUS NUMBER OF STATIONS ($N^{\max} = 32$)

T (sec)		0.1	0.2	0.3	0.4	0.5	1.0	2.0	5.0	10.0
ZC	Mbps	24.69	24.76	24.80	24.81	24.81	24.81	24.82	24.83	24.83
	Fraction	0.68	0.68	0.69	0.69	0.69	0.69	0.69	0.69	0.69
L-ZC	Mbps	24.83	24.83	24.83	24.83	24.83	24.83	24.83	24.83	24.83
	Fraction	0.69	0.69	0.69	0.69	0.69	0.69	0.69	0.69	0.69
PC	Mbps	21.85	28.95	31.28	32.44	33.14	34.54	35.24	35.66	35.80
	Fracion	0.61	0.81	0.87	0.90	0.92	0.96	0.98	0.99	1.00

TABLE V
GOODPUT AND LEARNING COST COMPARISON ($N = 8, N^{\max} = 32$)

D. Entry and exit and the length of the superframe

As discussed above, entry and exit can be accommodated by the use of superframes, in which case the length of the superframes is a design variable: if the superframe is shorter a larger fraction of it will be used up in coordination and goodput will be less; if the superframe is longer potential entrants will be kept waiting longer. Thus, there is a tension between maximizing goodput and minimizing the amount of time stations spend waiting to enter. For a simple illustration, assume the objective of the designer is to maximize $\text{Utility}(L) = \text{Goodput}(L) - \alpha\text{Wait}(L)$, where $\text{Goodput}(L), \text{Wait}(L)$ are aggregate goodput and average waiting time achieved by using a superframe of length L and α is a purely numerical parameter that trades off goodput against waiting time. Figure 4 displays utility as a function of the superframe length for $\alpha = 0.1, 0.3, 0.5$ (using $N = 12, N^{\max} = 32$). As may be seen, for these parameters, there are optimal values of the superframe length L – but a suboptimal choice makes very little difference in overall utility, never more than 5%. Of course, the appropriate value of α (hence the proper resolution of the tension between goodput and waiting time) depends on the objective of the designer which will in turn depend on the environment, including the arrival/departure rates of stations and the nature of the network (home, office, coffee shop, airport, etc.).

VII. CONCLUSION

In this paper, we have proposed a new class of MAC protocols that deploy sophisticated learning techniques to

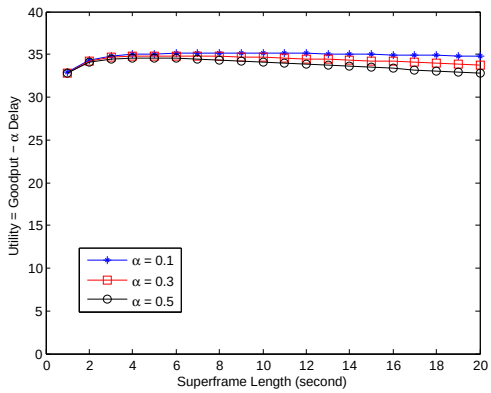

Fig. 4. Optimal superframe length.

achieve perfect coordination. Our proposed protocols are completely distributed, requiring neither any central control nor any exchange of control messages between stations and use minimal feedback: stations nodes that transmit in a given slot/period learn for certain whether or not some other station also transmitted in that slot/period but stations that do not transmit in a given slot/period learn nothing. Our results show that despite this minimal feedback the proposed protocols converge very quickly to perfect coordination and yield optimal throughput. Thus, they outperform existing coordination protocols, despite using much less feedback. Our protocols can easily be extended to provide quality-of-service differentiation across stations and adapted for deployment in cognitive radio networks in which stations can transmit across multiple frequency bands.

Appendix A: Flow Charts

Figures 5 and 6 are flow charts for the two phases of the protocol $\Phi(N, K, p)$; Figures (7), (8), (9) and (10) are flow charts for the four phases of the protocol $\Psi(N^{\max}, K, q)$.

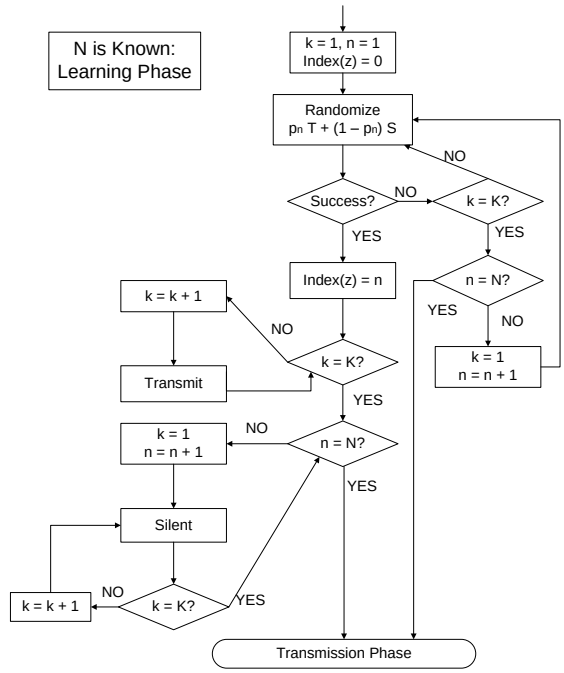

Fig. 5. $\Phi(N, K, p)$: Learning Phase

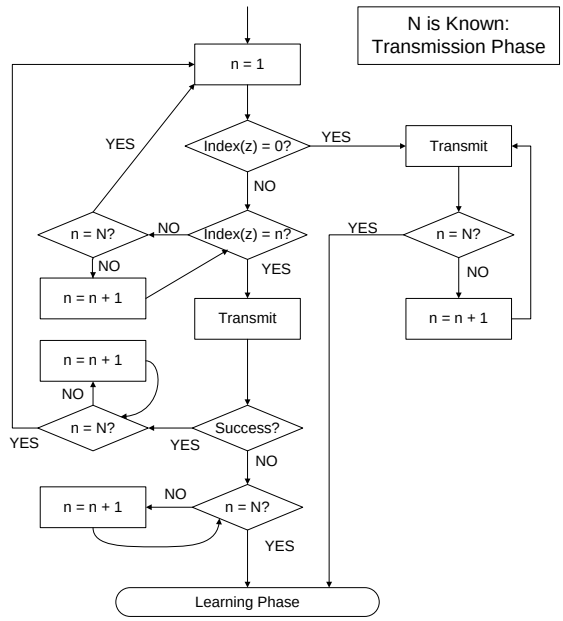

Fig. 6. $\Phi(N, K, p)$: Transmission Phase

Appendix B: Summary of ZC, L-ZC, SRB Protocols

Here we give brief summaries of the main protocols we have used for comparison purposes. In each case there are N users; N is unknown but less than some maximum M.

- **ZC** [8]
  1) Each round consists of M slots; this is a *fixed number*.
  2) At the beginning of the first round, each station chooses a slot uniformly in $\{1, 2, , M\}$ and transmits in the chosen slot.
  3) If the station transmits successfully, then it chooses the same slot again for the next round.
  4) If the station transmits unsuccessfully (collides), then

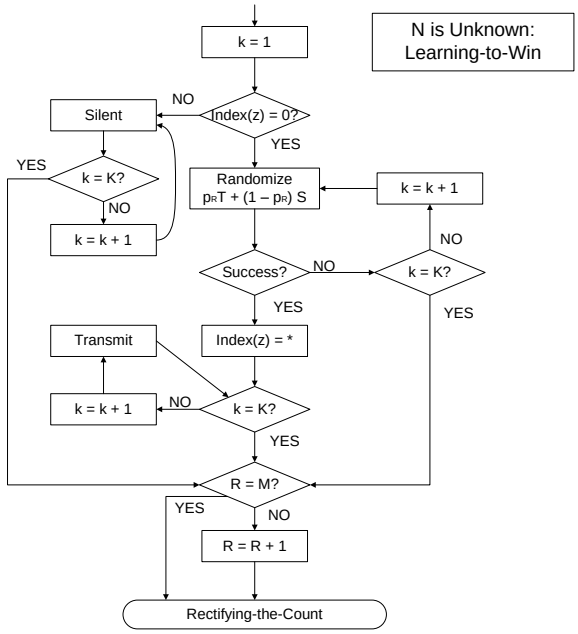

Fig. 7. $\Psi(N^{\max}, K, q)$: Learning-to-Win Phase

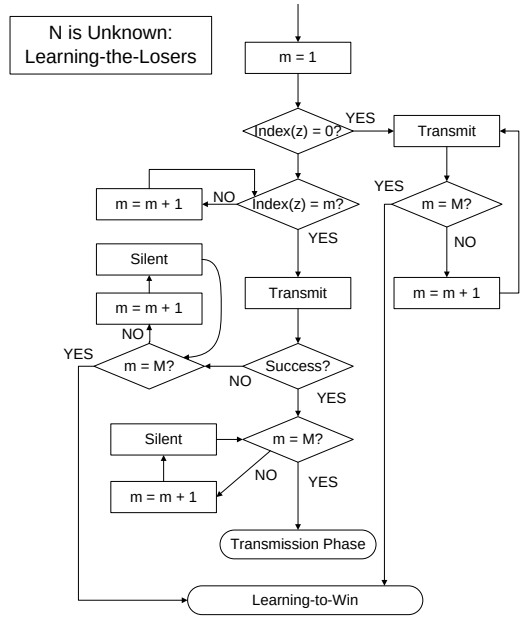

Fig. 9. $\Psi(N^{\max}, K, q)$: Learning-the-Losers Phase

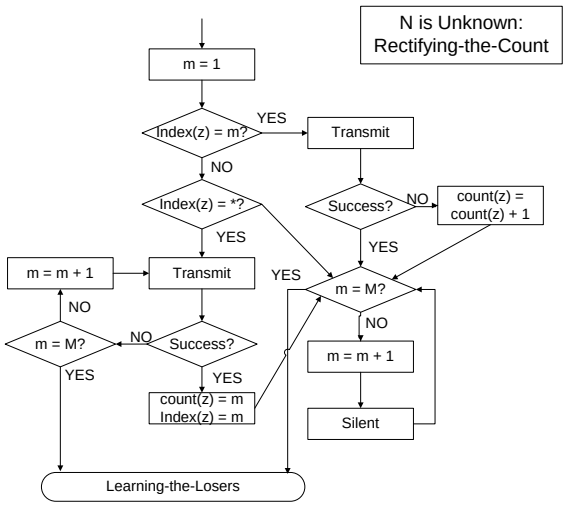

Fig. 8. $\Psi(N^{\max}, K, q)$: Rectifying-the-Count Phase

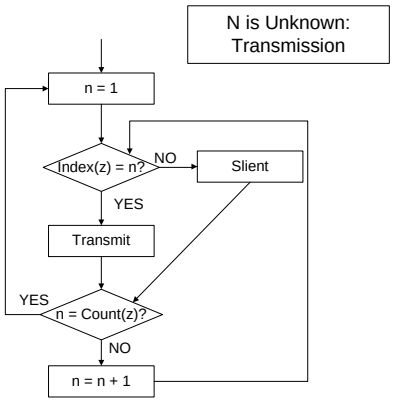

Fig. 10. $\Psi(N^{\max}, K, q)$: Transmission Phase

the station observes which slots were idle in this round; in the next round the station chooses a slot uniformly from among the idle slots.

5) The cycle continues; eventually there are no collisions (with probability 1).

- **L-ZC** [9]
  1) Each round consists of M slots; this is a *fixed number*.
  2) At the beginning of the first round, each station chooses a slot uniformly in $\{1, 2, , M\}$ and transmits in the chosen slot.
  3) If the station transmits successfully, then it chooses the same slot again for the next round.
  4) If the station transmits unsuccessfully (collides), then the station observes which slots were idle in this round; in the next round the station chooses the same slot with probability γ and a slot uniformly from among the idle slots with probability $1-\gamma$. [γ is a parameter of the protocol; [9] suggest the optimal value of γ is $\gamma^* = 1/(MN+2)$.)
  5) The cycle continues; eventually there are no collisions (with probability 1).
- **SRB** [10]
  1) Each station i uses a contention window $CW(i)$ and a cycle length M; $CW(i)$ is variable, M is fixed and constant across stations.
  2) At the beginning of the protocol, each station sets $CW(i) = CW_{\min}$ (a pre-determined integer) and chooses a slot uniformly in $\{1, \ldots, CW(i)\}$.
  3) If the station transmits successfully, then it transmits again M slots later.
  4) If the station transmits unsuccessfully (collides), then it doubles its contention window: $CW(i) = 2 * CW(i)$ and chooses a slot uniformly in $\{1, \ldots, CW(i)\}$.

5) The cycle continues; eventually there are no collisions (with probability 1).

REFERENCES

[1] N. Abramson, "The Aloha system: another alternative for computer communications," in *Proc. the AFIPS Fall Joint Computer Communication*, vol. 37, pp. 281-285, 1970.

[2] L. G. Roberts, "Aloha packet system with and without slots and capture," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 5, no. 2, pp. 28-42, Apr. 1975.

[3] W. Crowther, R. Rettberg, D. Walden, S. Ornstein, and F. Heart, "A system for broadcast communication: reservation-ALOHA," in *Proc. 6th Hawaii Int. Conf. Syst. Sci.*, Jan. 1973, pp. 371-374.

[4] G. Bianchi, "Performance analysis of the IEEE 802.11 distributed coordination function," *IEEE J. Sel. Areas Commun.*, vol. 18, no. 3, pp. 535-547, Mar. 2000.

[5] IEEE 802.11a, *Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications: Highspeed Physical Layer in the 5 GHz Bang*, Supplement to IEEE 802.11 Standard, Sep. 1999.

[6] A.C. Doyle, "Silver Blaze," *Strand Magazine* December 1892.

[7] I. Tinnirello, G. Bianchi, P. Gallo, D. Garlisi, F. Giuliano, F. Gringoli, "Wireless MAC processors: programming MAC protocols on commodity hardware," *IEEE INFOCOM 2012: 1269-1277*.

[8] J. Lee and J. Walrand, "Design and analysis of an asynchronous zero collision MAC protocol", Arxiv preprint 0806.3542, 2008.

[9] M. Fang, D. Malong, K. R. Duffy and D. J. Keith, "Decentralised learning MACs for collision-free access in WLANS," *Wireless Networks*, DOI: 10.1007/s11276-012-0452-1, May 2012.

[10] Y. He, R. Yuan, J. Sun, W. Gong, "Semi-random backoff: towards resource reservation for channel access in wireless LANs," *17th IEEE International Conference on Network Protocols (ICNP)*, 2009.

[11] J. Barcelo, C. Cano, B. Bellalta, M. Oliver, "Learning-BEB: avoiding collisions in WLAN," *Eunice Summer School*, 2008.

[12] J. Park and M. van der Schaar, "Medium access control protocols with memory," *IEEE/ACM Trans. Networking*, vol. 18, no. 6, pp. 1921-1934, Dec. 2010.

[13] K. Nakano and S. Olariu, "Randomized initialization protocols for ad hoc networks," *IEEE Trans. Parallel and Distributed Systems*, vol. 11, no. 7, Jul. 2000.

[14] K. Römer, "Time synchronization in ad hoc networks,"*ACM MobiHoc'01*, Long Beach, CA, Oct. 2001.

[15] R. T. Ma, V. Misra, and D. Rubenstein, "An analysis of generalized slotted-aloha protocols," *IEEE/ACM Trans. Netw.*, vol. 17, no. 3, pp. 936-949, Jun. 2009.

[16] R. L. Rivest, "Network control by Bayesian broadcast," *IEEE Trans. Inf. Theory*, vol. 33, no. 3, pp. 323-328, May 1987.

[17] J. -W. Lee, A. Tang, J. Huang, M. Chiang, and A. R. Calderbank, "Reverse-engineering MAC: a non-cooperative game model," *IEEE J. Sel. Areas Commun.*, vol. 25, no. 6, pp. 1135-1147, Aug. 2007.

[18] J. Capetanakis, "Tree algorithms for packet broadcast channels," *IEEE Trans. Inf. Theory*, vol. 25, no. 5, pp. 505-515, Sep. 1979.

[19] D. J. Goodman, R. A. Valenzuela, K. T. Gayliard, and B. Ramamurthi, "Packet reservation multiple access for local wireless communications," *IEEE Trans. Commun*, 37(8):885-890, 1989.

[20] N. Amitay and S. Nanda, "Resource auction multiple access (rama) for statistical multiplexing of speech in wireless pcs," *IEEE Trans. Veh. Tech.*, 43(3 Part I): 584-596, 1994.