This paper was presented as part of the main technical program at IEEE INFOCOM 2011

Distributed Power Control in the SINR Model

Zvi Lotker [1] Merav Parter[2] David Peleg[2] Yvonne Anne Pignolet[3]
zvilo@bgu.ac.il, merav.parter@weizmann.ac.il, david.peleg@weizmann.ac.il, yvonneanne.pignolet@ch.abb.com
[1]Ben Gurion University of the Negev, Beer-Sheva, Israel, [2]Weizmann Institute of Science, Rehovot, Israel,
[3]ABB Corporate Research, Dättwil, Switzerland

***Abstract*—The power control problem for wireless networks in the SINR model requires determining the optimal power assignment for a set of communication requests such that the SINR threshold is met for all receivers. If the network topology is known to all participants, then it is possible to compute an optimal power assignment in polynomial time. In realistic environments, however, such global knowledge is usually not available to every node. In addition, protocols that are based on global computation cannot support mobility and hardly adapt when participants dynamically join or leave the system. In this paper we present and analyze a fully distributed power control protocol that is based on *local* information. For a set of communication pairs, each consisting of a sender node and a designated receiver node, the algorithm enables the nodes to converge to the optimal power assignment (if there is one under the given constraints) quickly with high probability. Two types of bounded resources are considered, namely, the maximal transmission energy and the maximum distance between any sender and receiver. It is shown that the restriction to local computation increases the convergence rate by only a multiplicative factor of $O(\log n + \log\log \Psi_{\max})$, where $\Psi_{\max}$ is the maximal power constraint of the network. If the diameter of the network is bounded by $L_{\max}$ then the increase in convergence rate is given by $O(\log n + \log\log L_{\max})$.**

I. Introduction

A wireless device can decode a signal correctly if the received signal strength exceeds interference from simultaneous transmissions and noise by a device-dependent factor β (or equivalently, its Signal-to-Interference-plus-Noise-Ratio, SINR, exceeds β). This implies that if a device transmits with higher power and no other device changes its transmission power, then the chances that the designated receiver is able to process the message successfully increase.

The *Power Control with Power Constraints* problem involves a set $\mathbf{L} \in \mathbf{S} \times \mathbf{R}$ of n requests (given as sender-receiver pairs), and constraints bounding the maximal and minimal power ($\Psi_{\max}$ and $\Psi_{\min}$). For each request $l_i \in \mathbf{L}$, we need to select a power level $\psi(i)$ in the range $\Psi_{\min} \le \psi(i) \le \Psi_{\max}$ such that the SINR of each l_i exceeds β when all senders are active simultaneously.

The *Power Control with Diameter Constraint* problem again consists of a set of n requests $\mathbf{L} \in \mathbf{S} \times \mathbf{R}$, and in addition, the nodes are given an upper bound $L_{\max}$ on the diameter, i.e., the maximum distance between any sender $s_i \in \mathbf{S}$ and receiver $r_i \in \mathbf{R}$. For each request $l_i \in \mathbf{L}$, we need to select a power level $\psi(i)$ such that all SINR constraints are again satisfied simultaneously.

It is typically tedious, and sometimes even impossible, to recharge the wireless devices' energy source, hence energy conservation is a central goal in device management. In addition, using more power than necessary is not only wasteful but also detrimental to exploiting the full capacity of the medium, as it causes more interference. Thus, we strive to design algorithms that transmit with the lowest transmission energy necessary.

If all participants of a network are known, then it is possible to compute an optimal power assignment for all communication partners in polynomial time [21]. Starting with [1], simple algorithms that iteratively adjust the transmission power level based on the received signal and the interference at the receiver were suggested and their convergence analyzed. All these studies rely on the assumption that there is a *feedback channel* that lets a sender determine the SINR at the receiver. However, this assumption is not realistic for wireless sensor networks or ad hoc networks, where the global knowledge available to every node is typically limited. To cope with the common "partial" knowledge situation, a wireless node must be able to base its decisions for power control on *local* information only and converge to a good solution as fast as possible. An additional motivation for designing local power control algorithms is the fact that they facilitate a network's adaption to changing participants and mobility of nodes.

In this paper, we show how these goals can be achieved for a general scenario of a number of nodes distributed arbitrarily in a metric space. In Section III, we consider the problem of *Power Control with Power Constraints*, dealing with the case where the maximal transmission power is bounded by $\Psi_{\max}$. We propose a

Part of this research was conducted while Y.A. Pignolet was a postdoctoral fellow at the IBM Research Zurich Laboratory.
Z. Lotker, M. Parter and D. Peleg were supported in part by a grant from the Israel Science Foundation.

fully distributed randomized algorithm for this problem, enabling nodes to find the optimal power assignment with high probability within a number of rounds that is logarithmic in the number of nodes and in $\Psi_{\max}$. More precisely, we state the following.

Theorem 1.1: Given n requests, there is a fully distributed randomized algorithm that converges, with high probability, to the optimal power assignment (up to a predefined accuracy level ϵ) in $O(\log(\Psi_{\max}/\epsilon)(\log n + \log\log(\Psi_{\max}/\epsilon)))$ rounds, for $\epsilon > 0$.

The second part of this paper (Section IV) considers the *Power Control with Diameter Constraints* problem, dealing with the case where the network diameter is bounded by $L_{\max}$. For this case we state the following.

Theorem 1.2: Given n requests, there is a fully distributed randomized algorithm that converges, with high probability, to the optimal power assignment (up to a predefined accuracy level ϵ) in $O(\log(L_{\max}/\epsilon)(\log n + \log\log(L_{\max}/\epsilon)))$ rounds, for $\epsilon > 0$.

Interestingly, the second result is derived by a reduction from the *Power Control with Diameter Constraints* problem to the *Power Control with Power Constraints* problem. In particular, we show that if the network diameter is bounded, then so is the maximal transmission energy. A direct consequence of such a relation is that any algorithm designated for the bounded power case naturally extends to the bounded diameter case. We believe this relation may have further implications in the design and analysis of wireless networks protocols.

Related Work: Two main research lines underlie our current work. The first concerns distributed power control algorithms and the second is the recent [11] on scheduling communication requests. We start by providing a short review of iterative power control methods.

Non-uniform power assignments can clearly outperform a uniform assignment [15], [14] and increase the capacity of a network. Iterative distributed power control has been studied for various cellular settings, e.g., joint power control and base station selection [8], asynchronous protocols and their convergence [13], point-to-point, broadcasting and multicasting scenarios [10], multi-hop networks [20], power control under maximum constraints [5] and unifying the above approaches [19]. In [9] it is shown that the algorithms for these problems converge geometrically to the optimal solution. As typically the power cannot be adjusted to arbitrary levels, [2] proves that even for a set of discrete power levels there exists an algorithm that "converges" in a weaker sense. The problem of minimizing the total power transmitted using discrete power levels is considered in [17], where the problem is generalized to an integer programming problem, establishing an iterative power control method that converges to the optimal solution within a polynomial number of rounds (in the number of stations and power classes). Independent from the issue of power control, scheduling a set of communication requests is an important problem. See [4] for an overview of scheduling in the physical model with and without power control. Recently, a randomized distributed algorithm for scheduling a set of requests with *prespecified* power assignment has been proposed in [11], and shown to yield an $O(\log^2 n)$ approximation compared to the optimal schedule for this assignment. One ingredient of this algorithm is a randomized mechanism in which receivers send acknowledgement packets to their intended senders. Since our main challenge is to establish a feedback channel to support the iterative power control process (in a distributed manner), we adapt this acknowledgment mechanism to our settings, where the power is bounded and the assignment is unknown in advance.

II. Model and Preliminaries

The communication system considered here consists of two finite sets $\mathbf{S} = \{s_1, \ldots, s_n\}$, $\mathbf{R} = \{r_1, \ldots, r_n\}$ of *sender and receiver nodes*, forming a set $\mathbf{L}$ of n sender-receiver pairs $\mathbf{L}$=$\{l_1, \ldots, l_n\}$, where $l_i = \langle s_i, r_i\rangle$. Let η denote the ambient noise power level.

In addition, for the problem of *Power Control with Power Constraints*, we are given power constraints $\Psi_{\min}$ and $\Psi_{\max}$ s.t. each device can adjust its transmission power level in the range $\Psi_{\min} \leq \psi(i) \leq \Psi_{\max}$.

Finally, denote by $D(\mathbf{L})$ the *diameter* of the network, i.e., the maximal distance between any sender $s_i \in \mathbf{S}$ and receiver $r_j \in \mathbf{R}$. For the problem of *Power Control with Diameter Constraint* we are also given as input a parameter $L_{\max}$ such that $D(\mathbf{L}) \leq L_{\max}$.

As described in [18], a broad class of reception requirements can be described by a vector inequality of interference constraints of the form

$$\mathbf{\Psi} \ \geq \ \theta(\mathbf{\Psi}, \mathbf{L}, \eta) \ , \qquad (1)$$

where $\mathbf{\Psi} = \langle \psi(1), \ldots, \psi(n)\rangle$ is the *power assignment* vector (with sender s_i transmitting with power $\psi(i)$) and $\theta(\mathbf{\Psi}, \mathbf{L}, \eta) = \langle \theta_{r_1}(\mathbf{\Psi}, \mathbf{L}, \eta), \ldots, \theta_{r_n}(\mathbf{\Psi}, \mathbf{L}, \eta)\rangle$ denotes the *power threshold* induced by the interference of concurrent senders and ambient noise at each receiver r_i, namely, the minimum power level that will enable the transmitter power to overcome that interference.

This paper follows the physical SINR model [6], whereby the energy $E_{r_i}(\psi(i), s_i)$ received at r_i of a signal transmitted by s_i with power $\psi(i)$ is

$$E_{r_i}(\psi(i), s_i) \ = \ \psi(i) \cdot g(s_i, r_i) \ ,$$

where $g(u,v)$ is the *propagation attenuation* (a.k.a. *link gain*) between two nodes u, v, modeled as $g(u,v) = d(u,v)^{-\alpha}$, where $d(u,v)$ denotes the Euclidean distance between u and v. The *path-loss exponent* $\alpha \geq 1$ is a constant typically between 1.6 and 6. Measurements for indoor and outdoor path-loss exponents can be found in [16]. Given a pair $l_i \in \mathbf{L}$, we use the following notation; the length of the link is denoted by $d_i = d(s_i, r_i)$. Without loss of generality we assume $d_i \geq 1$ for every $i \in [1, n]$. We denote $\mathcal{I}_{r_i}(\psi(j), s_j) = E_{r_i}(\psi(j), s_j)$ for every other sender s_j concurrent to s_i, in order to emphasize that the signal transmitted by s_j is perceived at r_i as interference. The *total interference* $\mathcal{I}_{r_i}(\mathbf{\Psi}, \mathbf{L})$ experienced by a receiver r_i is now simply the sum of the interferences created by the $\mathbf{S}$ nodes transmitting simultaneously (except the intended sender s_i), i.e.,

$$\mathcal{I}_{r_i}(\mathbf{\Psi}, \mathbf{L}) = \sum_{s_j \in \mathbf{S} \setminus \{s_i\}} \mathcal{I}_{r_i}(\psi(j), s_j) \ .$$

Finally, define

$$\begin{aligned} \mathrm{SINR}_{r_i}(\mathbf{\Psi}, \mathbf{L}, \eta) &= \frac{E_{r_i}(\psi(i), s_i)}{\eta + \mathcal{I}_{r_i}(\mathbf{\Psi}, \mathbf{S})} \\ &= \frac{\frac{\psi(i)}{d(s_i, r_i)^\alpha}}{\eta + \sum_{j \neq i} \frac{\psi(j)}{d(s_j, r_i)^\alpha}} \end{aligned} \qquad (2)$$

where r_i correctly receives s_i's transmission (i.e., l_i has an *acceptable connection*) if and only if

$$\mathrm{SINR}_{r_i}(\mathbf{\Psi}, \mathbf{L}, \eta) \ \geq \ \beta \ , \qquad (3)$$

where $\beta \geq 1$ is the minimum ratio required for a successful message reception.

The total interference (due to competing transmitters and ambient noise) induces a *power threshold*

$$\begin{aligned} \rho &= \beta d_i^\alpha \left(\eta + \sum_{j \neq i} \psi(j) / d(s_j, r_i)^\alpha \right) \ , \\ \theta_{r_i}(\mathbf{\Psi}, \mathbf{L}, \eta) &= \min \{\Psi_{max}, \max \{\Psi_{min}, \rho\}\} \ . \end{aligned} \qquad (4)$$

A power assignment $\mathbf{\Psi}$ is called *feasible* if $\Psi_{\min} \leq \psi(i) \leq \Psi_{\max}$ for every $1 \leq i \leq n$ and Eq. (3) is satisfied. A set of requests $\mathbf{L}$ is *feasible* if there exists at least one feasible assignment $\mathbf{\Psi}$ for it. Denote by $\max(\mathbf{\Psi})$ the ratio between the largest and the smallest entries in the power assignment vector $\mathbf{\Psi}$ and by $\mathbf{\Psi}^*$ the optimal feasible assignment, such that $\max(\mathbf{\Psi}^*) \leq \max(\mathbf{\Psi})$ for every feasible assignment $\mathbf{\Psi}$. We adopt the convention that the vector inequality $\mathbf{\Psi} > \mathbf{\Psi'}$ indicates a strict inequality on all components. Let $\mathbf{\Psi}_{\max}$ be the vector with all entries equal to $\Psi_{\max}$. $\Psi_{\max}$ and $\Psi_{\min}$ can be normalized by fixing $\Psi_{\min} = 1$. Our distributed algorithm requires the satisfying power assignment, $\mathbf{\Psi}^*$, to obey two natural conditions. First, it is required that the received energy at each $r_i \in \mathbf{R}$ under $\mathbf{\Psi}^*$ is at least β/ω, for constant $\omega > 1$, that is

$$E_{r_i}(\psi^*(i), s_i) \ = \ \psi^*(i)/d_i^\alpha \ \geq \ \beta/\omega \ . \qquad (5)$$

(This should not be confused with the SINR criterion of Eq. (3).) Second, the powers are required to be sufficiently large so that the ambient noise, η, plays a minor role compared to interference. Formally, it is required that under $\mathbf{\Psi}^*$, the energy $E_{r_i}(\psi^*(i), s_i)$ is higher by at least some constant factor than the minimum power needed to deal with noise, for every $r_i \in \mathbf{R}$. As in [11], we require this constant to be 2. Combining this with Eq. (5), we have that for every $l_i \in \mathbf{L}$,

$$E_{r_i}(\psi^*(i), s_i) = \psi^*(i)/d_i^\alpha \geq \beta \cdot \max\{2\eta, 1/\omega\} \ . \quad (6)$$

Our problem can now be stated as follows. Given a configuration $\langle \mathbf{L}, \beta, \eta \rangle$ consisting of a feasible set $\mathbf{L}$, an SINR threshold β and an ambient noise level η, it is required to compute an optimal power assignment $\mathbf{\Psi}^*$.

III. Algorithm for Power-Restricted Configurations

This section establishes Theorem 1.1. We first briefly review the extensively studied *iterative power control* method [5], [13], [18], [10], [22] and some of its attributes and implications, aiming towards transforming it into a fully distributed algorithm in later sections. Throughout this section, we assume that the powers of the stations are in the range $[1, \Psi_{\max}]$.

A. Iterative Power Control Algorithm

For a given requests set $\mathbf{L}$, let $\mathbf{G} = [g(s_j, r_i)]$ be the link-gain matrix and let $\mathbf{Z} = [g(s_j, r_i)/g(s_i, r_i)]$ be its corresponding normalized matrix. Zander [21] showed that the *maximum achievable SINR* (denoted SINR^*) can be computed efficiently by solving an eigenvector problem of $\mathbf{Z}$. Given a system with generic power threshold constraints $\tilde{\theta}$ as in (1), consider the iterative power control algorithm where each sender s_i adjusts its transmission power for the next time slot $t+1$ to

$$\psi(i)^{t+1} \leftarrow \tilde{\theta}_{r_i}(\mathbf{\Psi}^t, \mathbf{L}, \eta) \ . \qquad (7)$$

Yates proved that this algorithm converges to the optimal solution if for all non-negative power assignment vectors $\mathbf{\Psi} \geq 0$, the power threshold function $\tilde{\theta}(\mathbf{\Psi}, \mathbf{L}, \eta)$ satisfies the following three requirements: (R1) positivity: $\tilde{\theta}(\mathbf{\Psi}, \mathbf{L}, \eta) > 0$, (R2) monotonicity: $\mathbf{\Psi} \geq \mathbf{\Psi'} \Rightarrow \tilde{\theta}(\mathbf{\Psi}, \mathbf{L}, \eta) \geq \tilde{\theta}(\mathbf{\Psi}, \mathbf{L}, \eta)$, and (R3) scalability: $a \cdot \tilde{\theta}(\mathbf{\Psi}, \mathbf{L}, \eta) > \tilde{\theta}(a \cdot \mathbf{\Psi}, \mathbf{L}, \eta)$ for any constant $a > 1$ [18]. Our interest is in the specific power threshold function given by Eq. (4). By Eq. (2) and (4), the rule of (7) and the corresponding iterative algorithm translate into the following procedure and algorithm.

Procedure PowerUpdateRule (for sender s_i)
$\psi' \leftarrow \psi(i)^t \cdot \frac{\beta}{\mathbf{SINR}_{r_i}(\mathbf{\Psi^t}, \mathbf{L}, \eta)}$;
$\psi(i)^{t+1} \leftarrow \min\{\Psi_{\max}\ ,\ \max\{1\ ,\ \psi'\}\}$;

Algorithm PowerUpdate
Set $\mathbf{\Psi^0} \leftarrow \mathbf{\Psi}_{\max}$;
For $t = 1$ to T_s do:

- Invoke Procedure PowerUpdateRule;

Throughout, it is assumed that $\beta < SINR^*$. It was shown that the function θ_{r_i} of Eq. (4) is positive, monotone and scalable [18], and this holds also under maximum and minimum power constraints [5]. Hence Alg. PowerUpdate converges to a fixed point. This convergence is guaranteed regardless of the initial power assignment $\mathbf{\Psi}^0$ and also holds in asynchronous systems, where some senders perform power adjustments faster than others, as well as in settings where some nodes base their decision on outdated information.

Let T_s denote the number of adjustment iterations necessary until the distance to the optimum solution $\mathbf{\Psi}^*$ is smaller than ϵ for some constant $\epsilon > 0$. To explicitly express T_s, we restrict ourselves to the case where PowerUpdate operates in a synchronous manner and where the initial power assignment $\mathbf{\Psi}^0$ is given by $\mathbf{\Psi_{max}}$ (as in Alg. PowerUpdate). By applying techniques from [9], T_s can be bounded as follows. Consider the norm $||\cdot||$ defined by $||\mathbf{\Psi}|| = \max_i |\psi(i)/\psi^*(i)|$ [5].

Theorem 3.1: $||\mathbf{\Psi}^t - \mathbf{\Psi}^*||$ is at most ϵ after $T_s = \lceil\log(\Psi_{\max}/\epsilon)/\log(SINR^*/\beta)\rceil$ iterations.

Proof: [9] shows that the sequence $\mathbf{\Psi}^t$ converges geometrically to $\mathbf{\Psi}^*$, s.t. $||\mathbf{\Psi}^t - \mathbf{\Psi}^*|| \le ||\mathbf{\Psi}^0 - \mathbf{\Psi}^*|| \cdot a^t$, for non-negative constant $a < 1$, for every $t \ge 0$. By Lemma 5 of [5], $a = \beta/SINR^*$, where $SINR^*$ is the maximum achievable SINR for L. As $||\mathbf{\Psi}^0 - \mathbf{\Psi}^*|| \le \Psi_{\max}$, it follows that $||\mathbf{\Psi}^t - \mathbf{\Psi}^*|| \le \epsilon$ is guaranteed after $t \ge T_s = \lceil\log(\Psi_{\max}/\epsilon)/\log(SINR^*/\beta)\rceil$ iterations. ■

As mentioned earlier, the sender cannot know the interference $\theta_{r_i}(\mathbf{\Psi}^t, \mathbf{L}, \eta)$ at the receiver, i.e., the receiver needs to inform the sender of the measured channel activity. However, sending these measurements back is problematic, as it is necessary to determine for each receiver the point in time at which it should transmit its measurements and the power level it should use. Hence it seems necessary to solve the joint problem of power control and scheduling for the dual set of feedback requests (where the roles of sender and receiver nodes are swapped) in order to be able to solve the power control problem. The key observation that allows us to tackle this circularity is in the following lemma, which rephrases Lemma 1 from [18] and follows by the monotonicity of the standard interference function.

Lemma 3.2: [18] If $\mathbf{\Psi}^0 \ge \theta(\mathbf{\Psi^0}, \mathbf{L}, \eta)$, and $\mathbf{\Psi}^{t+1}$ is obtained by applying Alg. PowerUpdate, then $\mathbf{\Psi^t}$ is a monotone decreasing sequence of power assignments that converges to the unique fixed point $\mathbf{\Psi}^*$.

Corollary 3.3: Procedure PowerUpdateRule changes the power level of a sender s_i in the next power vector (i.e., adjusts the corresponding entry $\mathbf{\Psi^{t+1}}(i)$) if and only if its link l_i has an acceptable connection when all senders transmit according to $\mathbf{\Psi^t}$.

Proof: Recall that the initial power vector is $\mathbf{\Psi^0} = \mathbf{\Psi}_{\max}$. By Lemma 3.2, in each application of Proc. PowerUpdateRule, the power of each sender s_i can only decrease. One can verify that such a decrease is only possible if the recent SINR measurement of s_i transmission was above β, i.e., l_i is satisfied when the senders transmit according to $\mathbf{\Psi^t}$. ■

Cor. 3.3 allows us to reduce the scheduling of *feedback requests* into a simpler problem for which techniques from [11] can be applied as discussed later. By a procedure named Feedback, a feedback channel is implemented that enables the receivers to transmit their measurements to the senders in $T_r = O(\log n + \log\log(\Psi_{\max}/\epsilon))$ steps with high probability over all T_s *feedback phases* until convergence. This allows us to apply Procedure PowerUpdateRule within an iterative process in our setting. All senders start with power level $\Psi_{\max}$. Each phase proceeds as follows. After the senders transmit, the receivers measure their observed SINR and use Procedure Feedback to inform their respective senders in $T_r = O(\log n + \log\log(\Psi_{\max}/\epsilon))$ steps. Upon getting this feedback, each sender acknowledges its receiver, adapts its transmission power and the process proceeds to the next phase.

By Theorem 3.1, the convergence rate of the iterative Algorithm PowerUpdate based on Procedure PowerUpdateRule is $T_s = O(\log(\Psi_{\max}/\epsilon))$. Using the algorithm for the feedback channel slows down the entire process by a factor of $T_r = O(\log n + \log\log(\Psi_{\max}/\epsilon))$. In what follows, we provide a detailed description of the protocol and analyze its behavior.

B. Feedback Channel Implementation

Consider a feasible set $\mathbf{L}$ of n requests. We begin by providing some notation adopted from [11]. Let $\widehat{l_i} = \langle r_i, s_i\rangle$ denote the *dual* request of l_i. Clearly, the dual of the dual link is the link itself, that is $\widehat{(\widehat{l_i})} = l_i$. Finally, the dual of the requests set $\mathbf{L}$ is denoted by $\widehat{\mathbf{L}} = \{\widehat{l_1}, \ldots, \widehat{l_n}\}$. For two requests l_i, l_j with power assignment $\mathbf{\Psi}$, the *affectance* of link l_i on link l_j is given by

$$a_{\mathbf{\Psi}}(l_i, l_j) = \min\left\{1, \beta\frac{\psi(i)}{d(s_i, r_j)^\alpha} / \left(\frac{\psi(j)}{d_j^\alpha} - \beta\eta\right)\right\} . \quad (8)$$

Let the *affectance* of link l_i on itself be $a_{\mathbf{\Psi}}(l_i, l_i) = 0$. The *affectance* of set of links $\mathbf{M}$ on link l_j is given by

$$a_{\mathbf{\Psi}}(\mathbf{M}, l_j) = \sum_{l_i \in \mathbf{M}} a_{\mathbf{\Psi}}(l_i, l_j) .$$

Let the *maximal average affectance* of a requests set $\mathbf{L}$ with power vector $\mathbf{\Psi}$ be given by

$$\overline{A}(\mathbf{L}, \mathbf{\Psi}) = \max_{\mathbf{M} \subseteq \mathbf{L}} \frac{1}{|\mathbf{M}|} \sum_{l_j \in \mathbf{M}} a_{\mathbf{\Psi}}(\mathbf{M}, l_j) . \quad (9)$$

Claim 3.4: If $\mathbf{L}$ is a feasible set of requests under assignment $\mathbf{\Psi}$, then $\overline{A}(\mathbf{L}, \mathbf{\Psi}) \leq 1$.

In our setting the nodes do not know the locations of other nodes or the power assignment that will allow them to communicate concurrently. Each node in $\mathbf{S} \cup \mathbf{R}$ is assumed to have the following knowledge: SINR parameters β and (for receivers only) α, some lower estimate, $SINR^-$ for $SINR^*$, the maximum achievable SINR, where $\beta < SINR^-$, link length d_i, maximal power level $\Psi_{\max}$, number of total requests n, and (for receivers only) an estimate for the minimal level of received energy, β/ω, and an estimate for the ambient noise level η, as will be elaborated later.

In this section we propose a randomized distributed algorithm that enables the receivers to send messages to their respective senders (namely, perform the dual request set $\widehat{\mathbf{L}}$) with high probability. These messages contain information on the SINR measured at the receiver, thus allowing the senders to adjust their transmission power and converge to the optimal power assignment. The key observation is that, when starting with power assignment $\mathbf{\Psi}^0 = \mathbf{\Psi}_{\max}$, it is sufficient for only a subset of requests in $\mathbf{L}$ to fulfil the dual request set $\widehat{\mathbf{L}}$. Interestingly, this subset is a feasible requests set under the current power assignment vector. Although our setting, as well as the goal, are different from those of the recent work of [11], the latter observation allows us to apply their techniques, as will be elaborated later on. The pseudocode of the procedures to be run by each sender and receiver in the system is also given.

1) Protocol Description: Our *Distributed Power Control* algorithm DPC is composed of two alternating phases, namely the *senders phase* and the *feedback phase* conducted by the receivers. Let T_s be the number of iterations required for the convergence of the power control algorithm with accuracy ϵ, as given by Theorem 3.1. Let T_r be the duration of each *feedback phase* (e.g., number of transmission rounds). Overall, our randomized algorithm consists of T_s *senders phases*, each of one round, and T_s *feedback phases*, each of T_r rounds.

Procedure PowerControl (for sender s_i)
$\psi(i) \leftarrow \Psi_{\max}$;
$p \leftarrow 1/16$;
$T_s \leftarrow \lceil \log(\Psi_{\max}/\epsilon) / \log(SINR^-/\beta) \rceil$;
$\Delta = \lceil \log \omega \rceil$;
$T_r \leftarrow \lceil 4/p((c+1)\ln n + \ln T_s) \rceil \cdot \Delta$;
On round $1 \leq t \leq T_r \cdot T_s$ do:

1) If $t \mod (T_r + 1) = 1$ then Transmit protocol message M_{s_i} to r_i using power $\psi(i)$;
2) Else do:
   Listen
   Upon successful reception of $M_{r_i}(\delta)$ from r_i:
   a) Send acknowledgment message $M_{s_i}(ack)$ to r_i using power $\psi(i)$.
   b) Set $\psi' \leftarrow \max\{1, \frac{\psi(i) \cdot \beta}{\delta}\}$;
   c) Set $\psi(i) \leftarrow \min\{\Psi_{\max}, \psi'\}$;

We begin by describing the senders protocol, formally given by Procedure PowerControl. Each senders phase consists of a single transmission round, in which all senders transmit simultaneously. In the first senders phase, all senders transmit with power $\Psi_{\max}$. In the feedback phase following each senders phase, each sender s_i whose power level needs to be updated, successfully receives (with high probability) from the corresponding receiver r_i the SINR measurement of its previously sent message. The sender s_i then sends an acknowledgment message to its receiver r_i and uses the information to recalculate its next power assignment according to Procedure PowerUpdateRule. Subsequently, the power level vector $\mathbf{\Psi}$ monotonically decreases until it converges to its optimum value $\mathbf{\Psi}^*$ with accuracy ϵ.

Next, we describe the receivers protocol, formally given by Procedure Feedback above. Let f_t denote the t^{th} *feedback phase*, for $1 \leq t \leq T_s$. Each receiver becomes active when it receives a signal from its sender. As all senders transmit simultaneously in the *senders phase* round, the receivers wake up synchronously and feedback phase f_t is initiated. When a receiver r_i successfully receives a signal from its sender s_i, it measures the observed SINR δ of the message just received, $\delta = \text{SINR}_{r_i}(\mathbf{\Psi}, \mathbf{L}, \eta)$. If it equals SINR threshold (i.e., $\delta = \beta$), then r_i does nothing and becomes inactive until the arrival of the next message. Otherwise ($\delta > \beta$),

Procedure Feedback (for receiver r_i)
$\psi(i) \leftarrow \Psi_{\max}$;
$\Delta = \lceil \log \omega \rceil$;
$T_s \leftarrow \lceil \log(\Psi_{\max}/\epsilon) / \log(SINR^{-}/\beta) \rceil$;
$p \leftarrow 1/16$;
$T_r \leftarrow \lceil 4/p((c+1)\ln n + \ln T_s) \rceil \cdot \Delta$;
Upon successful reception of a protocol message M_{s_i} from s_i do:

1) $success \leftarrow false$;
2) $\delta \leftarrow measured\ SINR\ of\ M_{s_i}$;
3) If $\delta > \beta$ then do:
   a) $\widetilde{\psi}(i) \leftarrow \Psi_{\max} \cdot \beta \cdot \max\{2\eta, 1\} \cdot d_i^{\alpha} / \psi(i)$;
   b) $\widehat{\psi}(i) \leftarrow \min\{\Psi_{\max}, \widetilde{\psi}(i)\}$;
   c) $\psi(i) \leftarrow \min\{\Psi_{\max}, \max\{1, \frac{\psi(i)\cdot\beta}{\delta}\}\}$;
   d) $j = \lfloor \log(\widehat{\psi}(i) \cdot \omega / \widetilde{\psi}(i)) \rfloor$;
   e) **Repeat** for $t = 1$ to T_r rounds:
      If $j = (t \mod (\Delta + 1)) - 1$ and $success \neq true$ then:
      - With probability p transmit $M_{r_i}(\delta)$ to s_i using power $\widehat{\psi}(i)$;
      - Listen
      - If acknowledgment message $M_{s_i}(ack)$ has been received then $success \leftarrow true$;

r_i attempts to transmit δ to its sender s_i using energy $\widehat{\psi}(i)$. This is repeated until successfully receiving an acknowledgment from the sender s_i or after T_r trials. Inspired by [11], $\widehat{\psi}(i)$ is given by

$$\widehat{\psi}(i) = \min\{\widetilde{\psi}(i), \Psi_{\max}\}\ , \qquad (10)$$

$$\widetilde{\psi}(i) = \Psi_{\max} \cdot \max\{2\widetilde{\eta}, 1\} \cdot \beta \cdot \frac{d_i^{\alpha}}{\psi(i)}\ , \qquad (11)$$

where $\widetilde{\eta}$ is an upper bound on the ambient noise that satisfies Ineq. (6). Note that in contrast with the solution of [11], the power levels are not assumed to be given in advance, nor are they expected to follow specific conditions. Moreover, in [11], the power level is not constrained as in our setting. Hence our dual power level function is distinct from [11], due to differences in setting and requirements. Yet, the two definitions share a common desired property, as will be seen later.

To be able to compute the dual power level, the receiver should have the knowledge of the current power level of its sender. This knowledge can be obtained either by requiring the sender to include this information in its message (this suffices since feedback is required only on successful transmissions). Alternately, knowing the senders protocol, the receiver can keep calculating the power level used by the sender on its own, by applying Procedure PowerUpdateRule in each phase (initially on $\Psi_{\max}$) and recording the obtained value for the next feedback phase if necessary. The set of dual links is partitioned into $\Delta = \lceil \log \omega \rceil$ classes as will be elaborated later. Let $\Delta(i) \in [0, \Delta - 1]$ denote the *link class* of link l_i. Each successful link l_i of class $\Delta(i)$ repeatedly attempts to send δ to s_i at the k^{th} round, for $\Delta(i) = (k \bmod (\Delta + 1)) - 1, \quad k \in [1, T_r]$. This is done by making a random choice, transmitting δ with probability p and remaining silent with probability $1-p$.

2) Correctness Analysis: Let $\mathbf{F} = \{f_1, \ldots, f_{T_s}\}$ be the set of feedback phases. Let $\widehat{\mathbf{L}}_\mathbf{t}$ be the set of dual requests that have to be fulfilled in feedback phase f_t. We say that f_t is successful if every dual request $\widehat{l_i} \in \widehat{\mathbf{L}}_\mathbf{t}$ is satisfied (e.g., successfully received at its target s_i) at least once during the phase f_t.

Lemma 3.5: Fix a constant $c > 1$, and let $T_r = \lceil \log \omega \rceil \cdot \lceil 64 \cdot (c+1) \cdot \ln n + \ln \log T_s \rceil$ be the number of rounds in every feedback phase f_t, for $1 \leq t \leq T_s$. Then $\mathbb{P}(\text{all feedback phases are successful}) \geq 1 - 1/n^c$.

The remainder of this section is dedicated to proving Lemma 3.5. For ease of analysis we focus on a certain f_t but the same applies to any $f_t \in \mathbf{F}$.

Given that each sender successfully receives the feedback regarding the last measured SINR activity at its receiver, Procedure PowerControl is guaranteed to converge with accuracy of ϵ within $T_s = \lceil \log(\Psi_{\max}/\epsilon) / \log(SINR^{-}/\beta) \rceil$ iterations (by Theorem 3.1 and Lemma 5 in [5]). Consequently, it remains to show that with probability greater than $1 - 1/n^c$ all T_s feedback phases of Procedure Feedback are successful when setting T_r as in the lemma. The analysis starts at the point in time where the receivers have just received the t^{th} message sent by their respective senders, using powers according to $\mathbf{\Psi^t}$.

By Cor. 3.3 it follows that the *feedback transmission* task for the set $\mathbf{L}$ is reduced to a *scheduling* task of only a subset $\widehat{\mathbf{L}}_\mathbf{t} \in \widehat{\mathbf{L}}$, whose dual $\mathbf{L_t}$ is feasible when the senders transmit according to $\mathbf{\Psi^t}$.

The next lemma shows that the power assigned to each dual link $\widehat{l_i} \in \widehat{\mathbf{L}}_\mathbf{t}$ follows two general conditions with respect to power constraints and ambient noise level. In particular, condition (b) of the lemma implies that noise has a relatively minor effect in feedback transmission.

Lemma 3.6: Let $\psi(i)$ be the power of s_i in the t^{th} message transmission. Let $\widehat{\psi}(i)$ be the corresponding energy of r_i in the feedback transmission. Then the dual energy $\widehat{\psi}(i)$ satisfies the following two conditions:
(a) $1 \leq \widehat{\psi}(i) \leq \Psi_{\max}$ for every $1 \leq i \leq n$,
(b) The energy received at any sender s_i satisfies $E_{s_i}(\widehat{\psi}(i), r_i) \geq \beta \cdot \max\{2\eta, 1/\omega\}$.

Proof: We begin with property (a). If $\widehat{\psi}(i) = \Psi_{\max}$, then property (a) trivially holds. Else, $\widehat{\psi}(i) < \Psi_{\max}$ and $\widehat{\psi}(i) = \widetilde{\psi}(i)$, so it remains to show that $\widetilde{\psi}(i) \geq 1$. This holds as $d_i^\alpha \geq 1$, $\psi(i) \leq \Psi_{\max}$ and $\beta \geq 1$.

Next, we prove property (b). Again we first consider the case where $\widehat{\psi}(i) = \Psi_{\max}$. As $\widehat{\psi}(i) = \Psi_{\max} \geq \psi^*(i)$,

$$E_{s_i}(\widehat{\psi}(i), r_i) \;\geq\; E_{r_i}(\psi^*(i), s_i) \;\geq\; \beta \cdot \max\{2\eta, 1/\omega\} \;,$$

where the last inequality follows by Requirement (6). Next, assume the complementary case where $\widehat{\psi}(i) = \widetilde{\psi}(i)$. Observe that the received energy at s_i when r_i transmits with power $\widehat{\psi}(i)$ is given by

$$E_{s_i}(\widehat{\psi}(i), r_i) \;=\; \frac{\widehat{\psi}(i)}{d_i^\alpha} \;=\; \frac{\Psi_{\max}}{\psi(i)} \cdot \beta \cdot \max\{2\eta, 1\} \;.$$

Since $\Psi_{\max}/\psi(i) \geq 1$ we also have that $E_{s_i}(\widehat{\psi}(i), r_i) \geq \beta \cdot \max\{2\eta, 1/\omega\}$, as required. ■

Note that condition (a) in Lemma 3.6 is required by the constraints of the system. Condition (b) is required for the sake of the analysis, where it allows a simpler treatment of ambient noise compared to [11]. In particular, it is necessary for proving that the *maximal average affectance* of the dual configuration is bounded from above by a constant.

We begin by providing two preliminary observations.

Observation 3.7: $\widetilde{\psi}(i)/\omega \leq \widehat{\psi}(i) \leq \widetilde{\psi}(i)$, for every $i \in [1, n]$.

Proof: The second inequality follows directly from Eq. (10). To prove the first inequality, let $\kappa = \max\{2\eta, 1\} \cdot \beta d_i^\alpha / \psi(i)$. By Eq. (11) it holds that $\widetilde{\psi}(i) = \kappa \cdot \Psi_{\max}$. If $\widetilde{\psi}(i) \leq \Psi_{\max}$, then $\widehat{\psi}(i) = \widetilde{\psi}(i)$, and the observation holds. We now consider the case where $\widetilde{\psi}(i) > \Psi_{\max}$. In this situation $\widehat{\psi}(i) = \Psi_{\max}$, so it remains to show that $\widetilde{\psi}(i) < \omega \cdot \Psi_{\max}$, or equivalently, that $\kappa \leq \omega$. As Ψ is monotonically decreasing, $\psi(i) \geq \psi^*(i)$, and we are guaranteed by Requirement (6) that

$$\kappa \leq \max\{2\eta, 1\} / \max\{2\eta, 1/\omega\}.$$

Thus there are three cases to consider. Case 1: $2\eta \geq 1$. In this case, $2\eta \geq 1/\omega$ and $\kappa \leq 1$. Case 2: $1 > 2\eta > 1/\omega$. In this case $\kappa \leq 1/(2\eta) < \omega$. Case 3: $2\eta \leq 1/\omega$. In this case $\kappa \leq 1/(1/\omega) = \omega$. As a consequence the claim follows. ■

Hereafter, we restrict our attention to sets of dual links in f_t that correspond to link classes defined as follows. Let $\widehat{\mathbf{L}}^{\mathbf{j}}$ be the set of dual links belonging to link class j in phase f_t of Procedure Feedback. That is, $\widehat{\mathbf{L}}^{\mathbf{j}} = \{\widehat{l}_i \mid \lfloor \log(\omega \cdot \widehat{\psi}(i)/\widetilde{\psi}(i)) \rfloor = j\}$. The set of dual links that attempt to transmit in step f_t and belong to link class j is given by $\widehat{\mathbf{L}}^{\mathbf{j}}_{\mathbf{t}} = \widehat{\mathbf{L}}_{\mathbf{t}} \cap \widehat{\mathbf{L}}^{\mathbf{j}}$. There are at most $\Delta = \lceil \log \omega \rceil$ classes as a direct consequence of Obs. 3.7.

Corollary 3.8: $\bigcup_{j \in [0, \lceil \log \omega \rceil - 1]} \widehat{\mathbf{L}}^{\mathbf{j}} = \widehat{\mathbf{L}}$.

Proof: By Obs. 3.7, $\omega \cdot \widehat{\psi}(i)/\widetilde{\psi}(i) \geq 1$. Therefore, link $\widehat{l}_i$ belongs to class j if

$$2^j \cdot \widetilde{\psi}(i)/\omega \;\leq\; \widehat{\psi}(i) \;\leq\; 2^{j+1} \cdot \widetilde{\psi}(i)/\omega \;. \qquad (12)$$

Since $j \in [0, \lceil \log \omega \rceil - 1]$, for any link $\widehat{l}_i$ where

$$\widetilde{\psi}(i)/\omega \;\leq\; \widehat{\psi}(i) \;\leq\; \widetilde{\psi}(i) \qquad (13)$$

there exists some $j \in [0, \Delta - 1]$ such that $\widehat{l}_i \in \widehat{\mathbf{L}}^{\mathbf{j}}$. By Obs. 3.7, it follows that Eq. (13) holds for any $\widehat{l}_i \in \widehat{\mathbf{L}}$ and the corollary is established. ■

Observation 3.9: Let $l_{i_1}, l_{i_2} \in \mathbf{L}^{\mathbf{j}}$ for any $j \in [0, \Delta - 1]$. Then $a_{\widehat{\mathbf{\Psi}}}(\widehat{l}_{i_1}, \widehat{l}_{i_2}) \;\leq\; 4 \cdot a_{\mathbf{\Psi}}(l_{i_2}, l_{i_1})$.

Proof: By Eq. (8) and Lemma 3.6(b),

$$a_{\widehat{\mathbf{\Psi}}}(\widehat{l}_{i_1}, \widehat{l}_{i_2}) \;\leq\; 2 \cdot \min\left\{1 \;,\; \beta \frac{\widehat{\psi}(i_1)}{d(r_{i_1}, s_{i_2})^\alpha} \Big/ \left(\frac{\widehat{\psi}(i_2)}{d_{i_2}^\alpha}\right)\right\} \;.$$

Recall that $\widehat{l}_{i_1}, \widehat{l}_{i_2} \in \widehat{\mathbf{L}}^{\mathbf{j}}$ and therefore by integrating Eq. (12) we have that $2^j \cdot \widetilde{\psi}(i_2)/\omega \leq \widehat{\psi}(i_2)$ and $2^{j+1} \cdot \widetilde{\psi}(i_1)/\omega \geq \widehat{\psi}(i_1)$, so

$$a_{\widehat{\mathbf{\Psi}}}(\widehat{l}_{i_1}, \widehat{l}_{i_2}) \;\leq\; 4 \cdot \min\left\{1, \beta \frac{\widetilde{\psi}(i_1)}{d(r_{i_1}, s_{i_2})^\alpha} \Big/ \left(\frac{\widetilde{\psi}(i_2)}{d_{i_2}^\alpha}\right)\right\} \;.$$

Next, by plugging in Eq. (11), we get that

$$\begin{aligned} & a_{\widehat{\mathbf{\Psi}}}(\widehat{l}_{i_1}, \widehat{l}_{i_2}) \\ \leq\;\; & 4 \cdot \min\left\{1, \beta \frac{\psi(i_2)}{d(s_{i_2}, r_{i_1})^\alpha} \Big/ \left(\frac{\psi(i_1)}{d_{i_1}^\alpha}\right)\right\} \\ \leq\;\; & 4 \cdot \min\left\{1, \beta \frac{\psi(i_2)}{d(s_{i_2}, r_{i_1})^\alpha} \Big/ \left(\frac{\psi(i_1)}{d_{i_1}^\alpha} - \beta \cdot \eta\right)\right\} \\ =\;\; & 4 \cdot a_{\mathbf{\Psi}}(l_{i_2}, l_{i_1}) \;, \end{aligned}$$

which establishes our claim. ■

We now bound the affectance for each class.

Lemma 3.10: The maximal average affectance satisfies $\overline{A}(\widehat{\mathbf{L}}^{\mathbf{j}}_{\mathbf{t}}, \widehat{\mathbf{\Psi}}^{\mathbf{t}}) \leq 4$ for every $j \in [0, \Delta - 1]$.

Proof: As $\mathbf{L}^{\mathbf{j}}_{\mathbf{t}}$ is a feasible requests set under $\mathbf{\Psi}^{\mathbf{t}}$, by Claim 3.4, $\overline{A}(\mathbf{L}^{\mathbf{j}}_{\mathbf{t}}, \mathbf{\Psi}^{\mathbf{t}}) \leq 1$. Combining Eq. (9) and Obs. 3.9 yields that for every $\widehat{\mathbf{M}} \subseteq \widehat{\mathbf{L}}^{\mathbf{j}}_{\mathbf{t}}$,

$$\begin{aligned} \overline{A}(\widehat{\mathbf{L}}^{\mathbf{j}}_{\mathbf{t}}, \widehat{\mathbf{\Psi}}^{\mathbf{t}}) \;&=\; \frac{1}{|\widehat{\mathbf{M}}|} \sum_{\widehat{l}_{i_2} \in \widehat{\mathbf{M}}} \sum_{\widehat{l}_{i_1} \in \widehat{\mathbf{M}}} a_{\widehat{\mathbf{\Psi}}^{\mathbf{t}}}(\widehat{l}_{i_1}, \widehat{l}_{i_2}) \\ &\leq\; \frac{4}{|\mathbf{M}|} \sum_{l_{i_2} \in \mathbf{M}} \sum_{l_{i_1} \in \mathbf{M}} a_{\mathbf{\Psi}^{\mathbf{t}}}(l_{i_1}, l_{i_2}) \\ &=\; 4 \cdot \overline{A}(\mathbf{L}^{\mathbf{j}}_{\mathbf{t}}, \mathbf{\Psi}^{\mathbf{t}}) \;\leq\; 4 \;, \end{aligned}$$

and the claim follows. ■

We are now ready to prove Lemma 3.5. Let $p = 1/16$ be the transmission probability of each $\widehat{l_i} \in \widehat{\mathbf{L}}^{\mathbf{j}}_{\mathbf{t}}$. By Lemma 3.10 it follows that

$$p \ \leq \ 1/(4 \cdot \overline{A}(\widehat{\mathbf{L}}^{\mathbf{j}}_{\mathbf{t}}, \widehat{\mathbf{\Psi}}^{\mathbf{t}})) \ .$$

Let n_t be a random variable indicating the number of dual requests that have not been successfully scheduled after $\Delta \cdot t$ rounds. That is, n_t is the sum of unsatisfied requests after each class executed t rounds of feedback attempts (where in each round, r_i that has not been acknowledged transmits with probability p). Then following the analysis of Lemma 1 and Theorem 2 in [11],

$$\mathbb{P}(n_t \neq 0) \leq \mathbb{E}(n_t) \leq (1/e)^{t \cdot p/4} \cdot n \ .$$

We want to determine t that guarantees that all T_s phases of Procedure Feedback, $\mathbf{F}$, are successful with probability of at least $1 - 1/n^c$. To do so, we apply the union bound over all phases such that

$$(1/e)^{t \cdot p/4} \cdot T_s \cdot n \leq 1/n^c \ .$$

Solving for t yields that

$$t \ \geq \ (4/p)((c+1)\ln n + \ln T_s) \ .$$

Recall that this corresponds to the number of rounds required for each link class j. As there are at most $\Delta = \lceil \log \omega \rceil$ classes, it follows that it suffices to use $T_r = \Delta \cdot (4/p)((c+1)\ln n + \ln T_s)$, and Lemma 3.5 is established.

Proof of Theorem 1.1: Correctness of Procedure PowerControl is established in Section III-A. The number of feedback iterations is bounded by $T_s \cdot T_r$. The claim now follows by Thm. 3.1 and Lemma 3.5. ■

Note that there are two degrees of freedom in our algorithm, namely, the parameters ϵ and c. Both affect the runtime and success guarantees, and reflect a time-success trade-off. Specifically, the lower ϵ is, the closer is the final power control vector to the optimal Ψ^*. Also, the higher c is, the greater is the chance that the algorithm will converge to the optimal vector up to ϵ.

IV. Algorithm for Diameter-Restricted Configurations

We now turn to prove Theorem 1.2. Recall that in the *Power Control with Diameter Constraints* problem, the network diameter is bounded by $L_{\max}$. As no power constraints are imposed, $\Psi_{\max} = \infty$, so Alg. DPC (of Section III-B) cannot be immediately applied. In what follows we show a reduction from the *Power Control with Diameter Constraints* problem to the *Power Control with Power Constraints* problem. That is, we show how the diameter constraint can be converted into power constraints that eventually allow us to apply Alg. DPC.

A. From constrained diameter to constrained power

Consider a feasible requests set $\mathbf{L}$ with ambient noise η, and without loss of generality assume that the distances are normalized so that $\min_{i,j} d(s_i, r_j) = 1$. Recall that $L_{\max} \geq \max_{i,j}\{d(s_i, r_j)\}$. Let $\mathbf{Z}$ be the normalized link-gain matrix for $\mathbf{L}$, with maximal eigenvalue λ and corresponding eigenvector $\mathbf{\Gamma} = \langle \gamma_1, \ldots, \gamma_n \rangle$. Recall that the elements of $\mathbf{Z}$ are $z(i,j) = d(s_i, r_i)^\alpha / d(s_j, r_i)^\alpha$. We next summarize some useful properties of $\mathbf{Z}$.

Lemma 4.1: For a given requests set $\mathbf{L}$, the normalized link-gain matrix $\mathbf{Z}$ satisfies the following properties.

1) $\mathbf{Z}$ consists of positive values only.
2) $\mathbf{Z}$ has only one real positive eigenvalue λ, i.e., the maximal one for which the corresponding eigenvector $\mathbf{\Gamma} = \langle \gamma(1), \ldots, \gamma(n) \rangle$ is positive.
3) $\mathbf{\Gamma}$ equals the optimal $\mathbf{\Psi}^*$ up to a common constant scaling factor c, i.e., $\mathbf{\Psi}^* = c \cdot \mathbf{\Gamma}$.
4) A requests set $\mathbf{L}$ is feasible iff $\lambda \ \leq \ 2$.

Without loss of generality let i_{max} (respectively i_{min}) be the index of the maximal (resp., minimal) entry in $\mathbf{\Gamma}$. Since $\mathbf{L}$ is a feasible set of links with ambient noise η, it is also feasible when there is no noise in the system, $\eta = 0$. Then by properties (2) and (4) in Lemma 4.1,

$$\mathbf{Z} \cdot \mathbf{\Gamma} = \lambda \cdot \mathbf{\Gamma} \leq \ 2\ \mathbf{\Gamma} \ ,$$

and hence, combined with the fact that the entries of $\mathbf{Z}$ and $\mathbf{\Gamma}$ are strictly positive (by properties (1) and (2) in Lemma 4.1), we have for every $1 \leq i,j \leq n$,

$$z(j,i) \cdot \gamma(i) \ < \ \sum_{k=1}^{n} z(j,k) \cdot \gamma(k) \ \leq \ 2\gamma(j) \ . \quad (14)$$

Observation 4.2: $1/L^\alpha_{\max} \leq z(i,j) \leq L^\alpha_{\max}$ for every $1 \leq i,j \leq n$.

Proof: $\min_{i,j} d(s_i, r_j) = 1$ by normalization of the distances, and $\max_{i,j} d(s_i, r_j) = L_{\max}$ by the $L_{\max}$ definition. Therefore, as $z(i,j) = d(s_i, r_i)^\alpha / d(s_j, r_i)^\alpha$, it follows that $1/L^\alpha_{\max} \leq z(i,j) \leq L^\alpha_{\max}$. ■

We are now ready to state our main lemma.

Lemma 4.3: If $\mathbf{L}$ is bounded by $L_{\max}$, then $\Psi_{\max} \leq 2L^\alpha_{\max}$.

Proof: By property (3) in Lemma 4.1, $\Psi_{\max} = \gamma(i_{max})/\gamma(i_{min})$. Using Obs. 4.2 and Eq. (14), we get

$$\begin{aligned} 1/L^\alpha_{\max} \cdot \Psi_{\max} \quad &\leq \quad z(i_{min}, i_{max}) \cdot \Psi_{\max} \\ &= \quad z(i_{min}, i_{max}) \cdot \frac{\gamma(i_{max})}{\gamma(i_{min})} \ < \ 2 \end{aligned}$$

and the lemma follows. ■

We are now ready to provide the proof of Theorem 1.2.

Proof of Theorem 1.2: Given a network with diameter bounded by $L_{\max}$, it follows by Lemma 4.3 that $\Psi_{\min} =$

1 and $\Psi_{\max} = 2L_{\max}^{\alpha}$. Applying Theorem 1.1 with the latter value yields our theorem. ■

Finally, we show another utilization of Lemma 4.3 in a somewhat different context. A recent work of [3] considers the capacity gap induced by the ability to adjust the transmission power (a.k.a. *power control*) if the available resources are bounded. We start by stating the following Theorem from [3].

Lemma 4.4 ([3]): Let $\mathbf{L}$ be a feasible requests set with a non-uniform power assignment vector $\mathbf{\Psi}^*$. Then there exists a subset $\mathbf{L}' \subseteq \mathbf{L}$ such that $\mathbf{L}'$ is feasible under a uniform power assignment and

$$|\mathbf{L}'| \geq |\mathbf{L}| \cdot \Omega(1/\log \Psi_{\max}) \ .$$

For the bounded diameter case we state the following.

Lemma 4.5: Let $\mathbf{L}$ be a feasible requests set with a non-uniform power assignment vector $\mathbf{\Psi}^*$. Then there exists a subset $\mathbf{L}' \subseteq \mathbf{L}$ such that $\mathbf{L}'$ is feasible under a uniform power assignment and

$$|\mathbf{L}'| \geq |\mathbf{L}| \cdot \Omega(1/\log L_{\max}) \ .$$

This lemma was proved by [3] only for the *one-dimensional* setting. Indeed, at the end of their paper, the authors pose an open question regarding extending this result to the *two-dimensional* case. We now provide this extension by proving Lemma 4.5.

Proof of Lemma 4.5: The proof follows the exact same line of the proof for Theorem 1.2. By Lemma 4.3, we convert the diameter constraint $L_{\max}$ into power constraints given by $\Psi_{\min} = 1$ and $\Psi_{\max} = 2L_{\max}^{\alpha}$. Applying Lemma 4.4, the result follows. ■

V. Conclusion

The first part of this paper presented a fully distributed power control protocol for a wireless network with power constraints in the physical SINR model. The iterative algorithm we suggest is randomized and converges to the optimal assignment quickly with high probability. It assumes nodes have minimal information on the network and it is based solely on local computation performed at each node. Interestingly, the capability to design such an algorithm for a bounded power network results from the fact that for constrained networks there exist starting points (power assignments) for which the iteratively assigned power changes monotonically. This allows one to apply simplifying reductions combined with existing tools. In the second part of this paper, we consider the case where the network diameter is bounded (for example, for networks embedded in a bounded region). We showed how the diameter constraint can be converted into a power constraint, which allows us to automatically apply the tools designated for the constrained power case. This relation can be applied in another context, namely, to study the capacity gap between uniform and non-uniform power assignments. In fact, we believe this reduction can serve as a useful general tool in the design and the analysis of networks with bounded resources.

References

[1] J. Aein. Power balancing in systems employing frequency reuse. *COMSAT Technical Review*, 3:277–300, 1973.

[2] M. Andersin, Z. Rosberg, and J. Zander. Distributed discrete power control in cellular PCS. *Wireless Personal Communications*, 6:211–231, 1998.

[3] C. Avin, Z. Lotker, and Y.A. Pignolet. On the Power of Uniform Power: Capacity of Wireless Networks with Bounded Resources. *17th ESA*, 2009.

[4] O. Goussevskaia, Y.A. Pignolet, and R. Wattenhofer. Efficiency of Wireless Networks: Approximation Algorithms for the Physical Interference Model. *Foundations and Trends in Networking: Vol. 4: No 3, pp 313-420.*, 2010.

[5] S. Grandhi, J. Zander, and R. Yates. Constrained power control. *Wireless Personal Communications*, 1:257–270, 1994.

[6] P. Gupta and P. R. Kumar. The Capacity of Wireless Networks. *IEEE Tr. Inf. Theory*, 46:388–404, 2000.

[7] M. Halldorsson and R. Wattenhofer. Wireless Communication is in APX. *36th ICALP*, 2009.

[8] S. Hanly. An algorithm for combined cell-site selection and power control to maximize cellular spread spectrum capacity. *IEEE J. Select. Areas Comm.*, 13:1332–1340, 1995.

[9] C.-Y. Huang and R. D. Yates. Rate of convergence for minimum power assignment algorithms in cellular radio systems. *Wirel. Netw.*, 4:223–231, 1998.

[10] S. Jagannathan, A.T. Chronopoulos and S. Ponipireddy. Distributed power control in wireless communication systems. *Proc. ICCCN*, 493–496, 2002.

[11] T. Kesselheim and B. Vöcking. Distributed Contention Resolution in Wireless Networks. 2010.

[12] T.-H. Lee, J.-C. Lin, and Y. T. Su. Downlink Power Control Algorithms for Cellular Radio Systems. *IEEE Tr. Vehic. Technol.*, 44:89–94, 1995.

[13] D. Mitra. An asynchronous distributed algorithm for power control in cellular radio systems. *Wireless and Mobile Comm.*, pages 177–186, 1994.

[14] T. Moscibroda and R. Wattenhofer. The Complexity of Connectivity in Wireless Networks. *25th INFOCOM*, 2006.

[15] T. Moscibroda, R. Wattenhofer, and Y. Weber. Protocol Design Beyond Graph-based Models. *5th HOTNETS*, 2006.

[16] T. Rappaport. *Wireless communications*. Prentice Hall, 2002.

[17] C. Wu and D. Bertsekas. Distributed Power Control Algorithms for Wireless Networks. *IEEE Tr. Vehic. Technol.*, 50:504–514, 1999.

[18] R. Yates. A framework for uplink power control in cellular radio systems. *IEEE J. Select. Areas Comm.*, 13:1341–1347, 1995.

[19] R. Yates and C. Huang. Integrated power control and base station assignment. *IEEE Tr. Vehic. Technol.*, 44:638–644, 1995.

[20] X. Yufang and E.M. Yeh. Throughput Optimal Distributed Power Control of Stochastic Wireless Networks. *IEEE Tr. Networking*, 18:1054–1066,2010.

[21] J. Zander. Performance of optimum transmitter power control in cellular radiosystems. *IEEE Tr. Vehic. Technol.*, 41:57–62, 1992.

[22] J. Zander. Distributed cochannel interference control in cellular radio systems. *IEEE Tr. Vehic. Technol.*, vol. 41, Aug. 1992.