On the Feasibility of Unilateral Interference Cancellation in MIMO Networks

Douglas M. Blough, Paolo Santi, and Ramya Srinivasan

***Abstract*—The problem of MIMO feasiblity refers to whether it is possible to support specified numbers of streams allocated to the links of a MIMO network while cancelling all interference. In unilateral interference cancellation, nodes account only for interfering links that they have been assigned to cancel and ignore other interfering links. We present several different formulations of the unilateral MIMO feasibility problem and use these formulations to analyze the problem's complexity and develop heuristic feasibility algorithms. We first prove that the general unilateral feasibility problem is NP-complete. We then identify several special cases where the problem is solvable in polynomial time. These include when only receiver-side interference cancellation is performed, when all nodes have two antenna elements, and when the maximum degree of the network's interference graph is two. Finally, we present several heuristic feasibility algorithms derived from different problem formulations and we evaluate their accuracies on randomly generated MIMO networks.**

***Index Terms*—MIMO Networks, Feasibility, Interference Cancellation, Boolean Satisfiability**

I. Introduction

Multiple-input, multiple-output, or MIMO, technology has been one of the most significant advances in wireless communications in recent years. MIMO technology makes use of antenna arrays, containing multiple antenna elements, at both ends of a communication link. On a single MIMO link, diversity and array gains can be exploited in order to significantly increase the link's capacity. When multiple MIMO links are used concurrently on the same wireless channel, there is also the possibility to cancel interference between links. Interference cancellation provides increased performance benefits on top of diversity and multiplexing gains. For example, in [30], it was shown that, with interference cancellation, the number of concurrent streams that can be supported when every link interferes with every other link is twice the number that can be supported without interference cancellation.

The problem of how to optimally allocate MIMO resources across an arbitrary network configuration is extremely challenging. While interference cancellation improves spatial reuse, it also reduces the resources available to maximize the data rate of each individual link through spatial multiplexing and diversity exploitation. Thus, when approaching the MIMO resource allocation problem at the network level, there is a fundamental trade-off between boosting individual link performance and reducing interference. This tradeoff is called the diversity-multiplexing-interference cancellation trade-off, and achieving optimal performance within this trade-off space is one of the key open problems in the field [10].

At the core of network-level MIMO resource allocation is the *feasibility* problem. Informally, the feasibility problem is defined as follows (see Section IV for a formal definition): "given a specific interference network, the available MIMO resources (antenna array sizes), and a stream allocation vector (an allocation of streams to network links), can the links in the network spatially multiplex the allocated streams while cancelling interference between every pair of interfering links?". Given a stream allocation vector for a set of interfering links, calculating a high-performing set of MIMO beamforming and combining weights is typically done with iterative numerical algorithms that are computationally intensive [8], [21], [36]. If feasibility can be determined before calculating weights, one can avoid running the computation-intensive calculation unnecessarily on infeasible vectors. This approach was used successfully in [8] to substantially reduce the overall time necessary to optimize MIMO networks with up to 10 links. Feasibility testing can also be useful in MIMO stream-controlled MAC layers [32] to ensure that poor stream choices are not made and in joint scheduling and stream assignment algorithms [37] to validate scheduling assignments.

Feasibility has been considered previously as an algebraic problem [11], [26], [38]. The algebraic specification of feasibility permits solutions that make use of bilateral interference cancellation, in which both the transmitter of an interfering link and the receiver of an interfered-with link consider the interference when choosing their beamforming or combining weights. Most MIMO networking research has, instead, assumed that cancellation of interference from one link to another is specifically assigned to either the receiver of the interfered-with link or the transmitter of the interfering link, but not both. This approach has been referred to as *unilateral interference cancellation* [8].

Herein, we consider the problem of feasibility restricted to unilateral interference cancellation solutions. We thoroughly investigate, for the first time, the computational complexity of this unilateral feasibility problem. We begin by specifying the problem in matrix form. We then show that the matrix formulation can be recast as a Boolean satisfiability problem with a specific structure. We also present a graph formulation of the problem for a special case. We first prove that the unilateral feasibility problem is NP-complete. We then proceed to show that the problem can be solved in polynomial time for several special cases, including 1) when interference cancellation is performed only at receivers but not transmitters (or vice versa), 2) when each link interferes with and is interfered

D.M. Blough is with the School of Electrical and Computer Engineering at the Georgia Institute of Technology. P. Santi is with the CNR Istituto di Informatica e Telematica. R. Srinivasan is with Qualcomm.

by at most two other links, and 3) when the antenna array size is two on every node. Finally, we present several polynomial-time feasibility heuristics that arise from our different problem formulations and we evaluate their performances on randomly-generated MIMO networks.

II. Related Work

The general feasibility problem in MIMO networks has been shown to be equivalent to finding whether a system of multivariate polynomial (non-linear) equations is solvable [11], [26], [38]. To be specific, if the set of links is $\{l_1, l_2, \ldots, l_L\}$ and the number of streams allocated to an arbitrary link l_j is denoted by s_j, then the stream allocation is feasible if and only if the following conditions can be simultaneously satisfied for every receiver j:

$$\mathbf{U}_j^T \mathbf{H}_{ji} \mathbf{V}_i = 0, \ \ \forall i \neq j \quad (1)$$

and

$$\operatorname{rank}\left(\mathbf{U}_j^T \mathbf{H}_{jj} \mathbf{V}_j\right) = s_j$$

In these conditions, $\mathbf{U}_j$ represents the combining weights at receiver j, $\mathbf{V}_i$ represents the beamforming, or precoding, weights at transmitter i, and $\mathbf{H}_{ji}$ represents the channel coefficients matrix between transmitter i and receiver j. The first of these conditions states that all interference on receiver j is in the subspace orthogonal to $\mathbf{U}_j$ (meaning all interference is cancelled at receiver j) and the second condition ensures that the resulting system matrix on link j has sufficient rank to multiplex s_j streams. Finding an exact or approximate solution to these conditions has been the subject of extensive research [11], [14], [16], [18], [19], [24], [25], [27], [28]. Practical solutions often minimize (rather than completely eliminate) interference in order to try to maximize the sum data rate across all links. The problem of determining when these conditions are simultaneously satisfiable (the general feasibility problem) has been solved only for the case where every link carries one stream (the single-beam case) [22], [38] and for specific small networks, e.g. networks with three links [5].

The general feasibility problem allows bilateral interference cancellation. In this paper, we restrict solutions to unilateral interference cancellation. In this case, each transmitter or receiver has a local system of equations to solve. In one node's local system of equations, links for which the node has not been assigned to perform interference cancellation are ignored. A related topic is interference alignment [5], [11], [25] in which multiple transmitters align their interfering signals at a particular receiver so that the receiver can be made orthogonal to all interfering signals. In fact, any solution to Equation 1, by definition, ends up with interference being aligned. Thus, we do not consider interference alignment as a separate technique but rather as an end goal that can be arrived at via different techniques (unilateral vs. bilateral cancellation, for example).

With unilateral interference cancellation, each local system of equations is linear, assuming that the weights at the other side are fixed. As long as the number of streams multiplexed on a node's link plus the number of streams on interfering links that the node is assigned to cancel does not exceed the antenna array size of the node, and assuming a rich scattering environment, this local system of equations is solvable at every node [12]. In fact, in general, there are many solutions to each of these local systems.

Since the local systems solved with unilateral cancellation are dependent, a relevant question is whether compatible local solutions exist that simultaneously solve each of these systems. One approach to this, called order-based interference cancellation (OBIC) has been studied in [17], [29]. In order-based interference cancellation, nodes are assigned an order for local solution and each node must cancel interference with all nodes that precede it in the order. In this way, the weights of the preceding nodes are fixed and known at the time of local solution and each node's local solution is forced to work with prior ones. The OBIC approach restricts the feasible stream allocation space somewhat, because it does not permit cycles of interference cancellation assignments. However, it guarantees that local solutions can be pieced together to form a valid global solution and it is extremely fast because it is a one-step (non-iterative) approach. In [8], it was shown that, using iterative solution techniques, piecing local solutions into a global solution can be done successfully even when cycles exist in the interference cancellation assignments. Therefore, we consider a unilateral interference cancellation assignment to be feasible whenever the local systems of equations are all solvable. This is, in fact, an implicit assumption that has been commonly made in work on MIMO by the networking community, e.g. [4], [12], [23], [30], [31], [35] all use the unilateral interference cancellation model considered herein.

To our knowledge, the only existing works on computational complexity of the feasibility problem are [26], [31]. In [26], the authors prove two complexity results. First, they show that finding the maximum number of degrees of freedom is NP-complete. Second, they prove the stronger result that the simpler problem of determining whether a given degrees of freedom allocation is achievable via linear schemes is NP-complete. It is interesting to observe that, while the first technical result of [26] can be readily applied to the unilateral interference cancellation model considered herein[1], the second result of [26] cannot be extended to unilateral interference cancellation. A major technical contribution of this paper is proving that the feasibility problem remains NP-complete also in the unilateral interference cancellation model, which is a strictly simpler model than the linear scheme model considered in [26]. We remark that the extension to the results of [26] presented in this paper is highly non-trivial, as it is based on a completely different construction.

Also for the same problem, [26] proves that feasibility is solvable in polynomial time when every node has 2 antenna elements. We proved this same result under the unilateral interference cancellation model in [31].

In this paper, we significantly extend our preliminary work in [31] by adding:

- several new formulations of the unilateral MIMO feasibil-

[1]This is due to the fact that, in the construction used to prove the first complexity result of [26], each node has one antenna element, and under this condition the unilateral interference cancellation becomes equivalent to bilateral interference cancellation.

ity problem, including formal specification as a Boolean satisfiability problem and specification as a graph problem,
- proof of NP-completeness for unilateral MIMO feasibility,
- proof that unilateral MIMO feasibility is solvable in polynomial time when the maximum degree of the conflict graph is two (based on the Boolean assignment formulation), and
- a new unilateral feasibility heuristic for the case where every node has at most three antenna elements (based on the graph formulation).

Other works on MIMO networks have considered different problems, e.g. MAC protocols and/or scheduling [1], [6], [32], [34], [39], throughput optimization [3], [9], [12], [30], [35], [37], and routing [13], [33]. While these works did not explicitly consider the problem of feasibility, most assume there exists a way to evaluate whether a given stream allocation is feasible.

III. Background

A. Interference Cancellation with MIMO

The availability of channel state information at both transmitters and receivers allows both types of nodes to participate in interference cancellation. For a given transmitter interfering with a given receiver, the cancellation is done by setting the transmitter's beamforming weights and/or the receiver's combining weights in such a way as to make the interfering signal orthogonal to the receiving array [2]. In the communications literature, it is usually assumed that the interfering transmitter and interfered-with receiver both account for the interference while calculating their weights [11], [14], [24], [28]. However, in the networking literature, it is more commonly assumed that interference cancellation is assigned to either the transmitter or the receiver, but not both [3], [4], [17], [23], [29], [30], [31], [35]. In [8], Cortés–Peña, et al., compared these two approaches, referring to the former as *bilateral* interference cancellation and the latter as *unilateral* interference cancellation. In this paper, we assume unilateral interference cancellation. As we will show in the next section, one of the advantages of unilateral cancellation is that feasibility can be viewed as a combinatorial problem, in contrast with the classical algebraic formulation [38].

The capability of a node to cancel interference is determined by the number of antenna elements it possesses and how many streams are multiplexed on an interfering node's communication link. Let the number of streams spatially multiplexed on any link $l_j = (t_j, r_j)$ be denoted by s_j. A transmitter (or receiver) node i with k antenna elements can spatially multiplex s_i streams on its link and cancel interference at the receivers (or from the transmitters) of a set of links denoted by $\mathcal{L}_i$ if and only if:

$$s_i + \sum_{j \in \mathcal{L}_i} s_j \leq k \tag{2}$$

Here, we assume a rich scattering environment, where the full capabilities of the MIMO antenna array can be exploited.

Equation 2 shows that there is a trade-off between the number of streams a node can multiplex on its own link versus the number of interfering streams it can cancel, which is determined by the size of the node's antenna array. Thus, the antenna elements are degrees of freedom that the node can use either for spatial multiplexing or interference cancellation. In the networking literature, this model is sometimes referred to as the DOF (Degrees of Freedom) Model [3], [12], [29], [31]. In the communications literature, however, "degrees of freedom" is typically used to mean the total number of streams that can be simultaneously transmitted across the entire network [5], [15]. So as not to produce confusion between these concepts, we try to avoid referring to Equation 2 as the DOF Model. When unavoidable, we use the term "antenna DOFs" to refer to the degrees of freedom associated with an individual node, in order to distinguish the term from the total DOFs of the network.

B. Feasibility Examples

To discuss some concrete examples, we will adopt notation similar to that from [38]: an $(M \times N, S)^L$ network is one where every transmitter has M antenna elements, every receiver has N antenna elements, there are L links, and every link carries S streams.

Three link networks have been well studied. Two examples are a $(4\times 4, 2)^3$ network with 3 links, 4 antenna elements per node, and two streams per link, and a $(2\times 2, 1)^3$ network with 3 links, 2 antenna elements per node, and one stream per link, which were shown to be feasible in [5] using a cooperating transmitter solution. Both networks are also feasible with unilateral interference cancellation. Due to the symmetry of these networks, each transmitter can cancel its interference on exactly one receiver and each receiver can cancel interference from exactly one transmitter. Thus, an interference cancellation assignment where link 1 cancels all interference with link 2,[2] link 2 cancels all interference with link 3, and link 3 cancels all interference with link 1 satisfies Inequality 2 at every node and cancels all interference, for both of these networks. Using numerical solution techniques, $(4\times 4, 2)^3$ networks were also empirically shown to be feasible with both unilateral and bilateral interference cancellation in [8].

Another example is a $(2 \times 3, 1)^4$ network with one stream per link, 4 links, two antenna elements per transmitter, and 3 antenna elements per receiver. This network is shown to be feasible in [38] using algebraic techniques. It is unilaterally feasible in our model, because if links are arranged in a circle, each transmitter can cancel interference to the next receiver and each receiver can cancel interference from the next two transmitters and all interference is cancelled.

It should also be mentioned that all of the above examples ($(4 \times 4, 1)^3$, $(2 \times 2, 1)^3$, and $(2 \times 3, 1)^4$) are not feasible according to the stricter unilateral criterion of [29].

For a final example, consider $(5\times 5, 2)^4$ networks. Algebraic techniques in [38] and numerical solution in [11] suggest these networks are feasible. However, there is no feasible unilateral

[2] I.e., transmitter 1 nulls at receiver 2 and receiver 1 cancels interference from transmitter 2.

interference cancellation assignment for these networks. To see this, note that each node has only enough antenna elements to cancel interference from/to one other node but not enough for two nodes. Thus, the total number of cancellations that can occur is 8. However, a total of 12 cancellations is necessary to eliminate all interference (four receivers, each of which is interfered by 3 transmitters).

C. Network Model

Our results apply to MIMO networks with non-uniform antenna array sizes. In particular, we consider a MIMO network with a set of L links[3] given by $\mathcal{L} = \{l_1 = (t_1, r_1), \ldots l_L = (t_L, r_L)\}$, where vectors $K^t = [k^t_1, \ldots k^t_L]$ and $K^r = [k^r_1, \ldots k^r_L]$ are used to denote the number of antenna elements available at the transmitters and receivers of the links. For our study, we assume there is an interference threshold, below which interference can safely be ignored. Thus, if the received power of an interfering signal is below the interference threshold, we do not consider it. Interference relationships between links can, therefore, be described by a directed conflict graph $G_c = (V_c, E_c)$, where V_c is the set of links in the network and $e = (l_i, l_j) \in E_c$ if and only if the transmission on link l_i interferes with the receiver of link l_j (the received power of t_i's signal at r_j is above the interference threshold). With a slight abuse of notation, we denote both the conflict graph and its adjacency matrix representation by G_c. Therefore, $G_c[i, j] = 1$ if $e = (l_i, l_j) \in E_c$, and $G_c[i, j] = 0$ otherwise. We set the diagonal elements of the adjacency matrix G_c to 1.

Furthermore, we assume that each node is equipped with only a single radio. Therefore, a basic constraint on concurrency of transmissions is that each node can participate in one transmission at a time, either as transmitter or as receiver. A set of links is said to be *primary-interference-free* if and only if it satisfies that condition, i.e. that every node in the network appears as an endpoint of at most one link in the set.

IV. Unilateral Feasibility Problem Definition

A. Matrix Formulation

Consider a multi-hop MIMO network defined as in the previous section. Let the network's link set be denoted by $\mathcal{L}$, its conflict graph by G_c, and let $S = [s_1 \ldots s_L]$ be an $L \times 1$ stream allocation vector containing the number of data streams multiplexed by each link, where $L = |\mathcal{L}|$.

For the stream allocation vector S to be feasible over $\mathcal{L}$, interference between every pair of links must be cancelled. However, in cancelling interference with unilateral interference cancellation, each node is limited by the number of antenna elements it possesses. For a transmitter t_i, this constraint can be expressed in the following way:

$$s_i + \sum_{j=1, j\neq i}^{L} a^t_{ij} s_j \leq k^t_i \;\; \forall i \in 1 \ldots L \tag{3}$$

[3]The communications literature on the MIMO interference channel commonly refers to communicating entities as "users", while we prefer the networking terminology of "links".

where a^t_{ij} is a Boolean variable such that $a^t_{ij} = 1$ if the transmitter of link i cancels the generated interference at the receiver of link j, and $a^t_{ij} = 0$ otherwise. If we let $a^t_{ii} = 1 \; \forall i$, and $a^t_{ij} = 0$ if $G_c[i, j] = 0$. Equation (3), across all transmitters, can be written as:

$$A_t S \leq K^t \;, \tag{4}$$

where A_t is a Boolean matrix containing the a^t_{ij} values. Similarly, for a receiver r_i, we can write:

$$s_i + \sum_{j=1, j\neq i}^{L} a^r_{ij} s_j \leq k^r_i \;\; \forall i \in 1 \ldots L \tag{5}$$

where a^r_{ij} is a Boolean variable such that $a^r_{ij} = 1$ if the receiver of link i cancels the transmission on link j, and $a^r_{ij} = 0$ otherwise. As above, we let $a^t_{ii} = 1 \; \forall i$, and $a^r_{ij} = 0$ if $G_c[j, i] = 0$. Equation (5) can then be combined across all receivers as:

$$A_r S \leq K^r, \tag{6}$$

where A_r is the Boolean matrix of a^r_{ij} values.

Without loss of generality, we assume interference cancellation is *coordinated* such that, for any link l_i interfering with another link l_j, either the transmitter of l_i nulls its signal at the receiver of l_j, or the receiver of l_j cancels the signal from the transmitter of l_i, but not both. Having both transmitter and receiver cancel the same interference uses unnecessary resources and any solution to the defined problem having such a property can be directly transformed into a solution where only one side cancels the interference by just setting one of the two variables, a^t_{ij} or a^r_{ji}, to zero. With this assumption, matrices A_t and A_r are related such that any choice of A_t completely determines A_r, and vice-versa, according to the following equation:

$$A_t = G_c + I - A_r^T, \tag{7}$$

where I is the identity matrix.

The matrix formulation of the unilateral feasibility problem is formally defined below.

Input: A set $\mathcal{L} = \{(t_1, r_1), \ldots, (t_L, r_L)\}$ of primary-interference-free links, a stream allocation vector S for $\mathcal{L}$, antenna element vectors K^t and K^r, and a conflict graph $G_c = (\mathcal{L}, E_c)$.

Output: **True** if S and $\mathcal{L}$ are feasible, and **False** otherwise. S and $\mathcal{L}$ are defined to be feasible if $\mathcal{L}$ is free of primary interference and there exist matrices A_t and A_r such that:

1) $A_t S \leq K^t$,
2) $A_r S \leq K^r$, and
3) $A_t = G_c + I - A_r^T$,

where I is the identity matrix.

The matrix pair (A_t, A_r) is called an interference cancellation (IC) assignment for the link set $\mathcal{L}$. The stream vector S is said to be feasible over $\mathcal{L}$ if there exists at least one valid IC assignment (A_t, A_r) that satisfies Conditions 1–3 above. We then say that (A_t, A_r) *supports* the stream vector S over $\mathcal{L}$, i.e., all interference between the links in $\mathcal{L}$ can be removed by using the available MIMO resources. Finally, the *feasible space* of the network can be obtained by identifying the set of all feasible stream allocation vectors.

B. Accuracy of Feasibility Definition

In this section, we prove that the conditions specified in our unilateral feasibility problem definition are necessary and locally sufficient for the calculation of beamforming and combining weights that support the given stream allocation vector. By local sufficiency, we mean that, for each node, if the beamforming or combining weights of every other node are fixed, then there is a solution to the weights of the local node that cancels all interference to/from nodes it is assigned and can support the allocated number of streams on its link.

Theorem 1: Existence of interference cancellation assignment matrices A_t and A_r satisfying Conditions 1–3 in the feasibility problem definition is necessary and locally sufficient for a stream vector to be feasible with unilateral interference cancellation in a rich scattering environment.

Proof:
Necessity:
Consider what happens if there are no matrices A_t and A_r that satisfy all of Conditions 1–3. Thus, for every possible way of performing unilateral cancellation, at least one condition is violated. If Condition 1 or Condition 2 is violated, then some node is assigned more interference than it can cancel, according to its antenna element constraint (Inequality 2 and (A_t, A_r) does not support the stream vector. If Condition 3 is violated, then there is some interference that is not accounted for and, again, (A_t, A_r) does not support the stream vector.

Local sufficiency: Assume that there are interference cancellation assignment matrices A_t and A_r that satisfy Conditions 1–3 in the feasibility definition and that we have a rich scattering environment.

We show local sufficiency for an arbitrary receiver under the given assumptions. The analysis for an arbitrary transmitter is essentially identical. Consider an arbitrary receiver r_i with s_i streams allocated on r_i's link. Without loss of generality, let the transmitters that r_i is assigned to cancel interference from in A_r be $t_1, t_2, \ldots, t_m$, where $s_i + \sum_{j=1}^{m} s_j \leq k_i^r$, because Condition 2 is satisfied. Let the beamforming weights of t_j be V_j. Let the interference channel for r_i be defined as:

$$H_{\text{int}} = [H_{1,i}V_1 | H_{2,i}V_2 | \cdots H_{m,i}V_m]$$

The interference channel contains the combined interference from all transmitters that r_j is assigned to cancel. The local problem is now to calculate combining weights U_i such that:

$$U_i^T H_{\text{int}} = 0 \tag{8}$$

and

$$\text{rank}(U_i^T H_{i,i} V_i) = s_i$$

Now, U_i^T is of size $s_i \times k_i^r$ and H_{int} is of size $k_i^r \times \sum_{j=1}^{m} s_j$. Thus, the right hand side of Equation 8 is a matrix of all 0's of size $s_i \times \sum_{j=1}^{m} s_j$. From the preceding dimensionalities, Equation 8 is a system of $s_i \times \sum_{j=1}^{m} s_j$ equations with $s_i \times k_i^r$ unknowns. This system has more unknowns than equations, because $k_i^r \geq s_i + \sum_{j=1}^{m} s_j > \sum_{j=1}^{m} s_j$. Thus, in a rich scattering environment, where H_{int} is full rank, there are multiple solutions to Equation 8. The difference between the number of equations and unknowns is, in fact, s_i, which means that Equation IV-B is also satisfied.

■

Some discussion of Theorem 1 is warranted. This is not an exact characterization of unilateral feasibility, because there might be some cases where local sufficiency at every node does not yield an overall solution. In fact, an exact characterization is not possible since there are always choices of channel matrices that prevent solution even where MIMO resources are sufficient everywhere for interference cancellation. This is true even when working with the direct algebraic system of Equation 1 as pointed out in [38]. Nevertheless, it is still important to understand the complexity of applying this characterization to real systems, since it has been shown to have good agreement with solutions generated by numerical techniques, as pointed out in Section III-B, and it has been widely adopted in the MIMO networking community [4], [12], [23], [30], [31], [35]. Furthermore, if feasibility algorithms are used to prune the search space of possible stream allocation vectors for maximizing throughput, as was done in [8], if one or two infeasible vectors are evaluated due to an unlucky combination of channels, they will be rejected for low throughput and result in only a slight increase in execution time. In this situation, it is preferable to consider too many possibilities than to have a conservative model, which rejects some feasible vectors that might result in very high throughput.

C. Unilateral Feasibility as a Boolean Satisfiability Problem

The matrix formulation of unilateral feasibility suggests that it is related to Boolean satisfiability. The A_t and A_r matrices contain sets of Boolean variables that must satisfy certain constraints (Conditions 1–3 in the matrix formulation). We call Condition 3 the interference constraint, because it says that interference between every pair of interfering links must be cancelled. We call Conditions 1 and 2 the antenna element constraints, because they limit the number of streams that can be multiplexed and cancelled by a given node based on the number of antenna elements of the node.

In a Boolean satisfiability problem, values must be found for a set of Boolean variables, which simultaneously satisfy a set of disjunctive clauses. Thus, to formulate feasibility as a satisfiability problem, Conditions 1–3 must be rewritten as sets of disjunctive clauses. Condition 3, the interference constraint, is relatively straightforward. It says that for every edge (i, j) in the conflict graph, meaning for every link l_i that interferes with any other link l_j,

$$a_{ij}^t \vee a_{ji}^r \tag{9}$$

This simply states that either the transmitter of l_i or the receiver of l_j must cancel the interference from l_i to l_j. This constraint leaves open the possibility that both a_{ij}^t *and* a_{ji}^r are set to true. However, we reiterate that if there is a feasible solution with both these variables set to true, then there is also a feasible solution with only one of them set to true. Taking

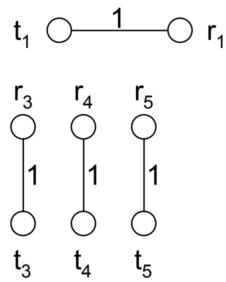

Fig. 1. Example used to Illustrate Antenna Constraint in Disjunctive Form

Expression 9 over all edges of the conflict graph yields a set of disjunctive clauses that together are equivalent to Condition 3 of the matrix formulation.

The antenna element constraints are not as straightforward. The exact set of clauses corresponding to one of these constraints, given by Equation 3 or Equation 5, depends on the conflict graph and the stream allocation vector. Consider a simple example for one transmitter, depicted in Figure 1. In this figure, t_1 causes interference only on r_2, r_3, r_4, and r_5, but no other receivers. Assume that all links in the figure are allocated one stream in the stream allocation vector and that t_1 has 3 antenna elements. Then, Equation 3 simplifies to:

$$\sum_{j=2}^{5} a_{1j}^t \leq 2$$

This equation states that at least two of the a_{1j}^t variables must be false (zero). Leaving off the t superscript for simplicity, we can rewrite this equation as a Boolean expression in the following way:

$$\overline{a_{12}}\,\overline{a_{13}} \vee \overline{a_{12}}\,\overline{a_{14}} \vee \overline{a_{12}}\,\overline{a_{15}} \vee \overline{a_{13}}\,\overline{a_{14}} \vee \overline{a_{13}}\,\overline{a_{15}} \vee \overline{a_{14}}\,\overline{a_{15}}$$

Applying DeMorgan's Theorem to this and simplifying yields the following set of disjunctive clauses that are equivalent to this expression:

$$\begin{gathered} \overline{a_{12}} \vee \overline{a_{13}} \vee \overline{a_{14}} \vee \overline{a_{15}} \;, \\ \overline{a_{12}} \vee \overline{a_{13}} \vee \overline{a_{14}} \;, \\ \overline{a_{12}} \vee \overline{a_{13}} \vee \overline{a_{15}} \;, \\ \overline{a_{12}} \vee \overline{a_{14}} \vee \overline{a_{15}} \;, \\ \overline{a_{13}} \vee \overline{a_{14}} \vee \overline{a_{15}} \end{gathered}$$

The first clause dictates that at least one of the a_{1j} variables must be false. Setting one of the variables to false makes three of the remaining clauses true. This leaves the fourth clause to be satisfied, which requires one of the remaining a_{1j} variables to be false. Thus, these clauses together ensure that at least two of the a_{1j} variables are set to false (zero).

Generalizing this example can generate a set of disjunctive clauses for every transmitter and every receiver in the network. For an arbitrary transmitter t_i, the procedure is given by Procedure FindAntennaConstraintClauses, as follows:

Procedure FindAntennaConstraintClauses

1) enumerate the minimal combinations of interfered-with receivers, such that setting those a_{ij}^t variables to zero will allow t_i to cancel its interference on the remaining receivers while staying within the antenna element constraints (this depends on k_i^t, s_i, and the s_j's of the receivers),
2) use DeMorgan's Theorem to convert the sum of products expression from Step 1 into a product of sums expression,
3) each "sum" term from Step 2 represents one disjunctive clause in the Boolean satisfiability problem to be solved.

The procedure for a receiver is completely symmetric to this.

The Boolean satisfiability formulation of the unilateral feasibility problem is formally defined below.

Input: A set $\mathcal{L} = \{(t_1, r_1), \ldots, (t_L, r_L)\}$ of primary-interference-free links, a stream allocation vector S for $\mathcal{L}$, antenna element vectors K^t and K^r, and a conflict graph $G = (\mathcal{L}, E)$.

Output: **True** if S and $\mathcal{L}$ are feasible, and **False** otherwise. S and $\mathcal{L}$ are defined to be feasible if there is an assignment of values to Boolean variables a_{ij}^t and a_{ij}^r, $\forall i, j \in [1, \ldots, L]$, that simultaneously satisfies the following Boolean clauses:

$$a_{ij}^t \vee a_{ji}^r, \forall (l_i, l_j) \in E$$

and all clauses generated according to Procedure FindAntennaConstraintClauses for each t_i and r_j that occur in $\mathcal{L}$.

One question to be answered is how many disjunctive clauses can be produced by Procedure FindAntennaConstraintClauses for an arbitrary node in the worst case? The number of a_{ij} variables to be considered is determined by the degree of the node in the conflict graph. For conflict graphs with high degree, say on the order of L, the number of clauses generated could be exponential in L. However, for conflict graphs with lower degree, the number is smaller. In particular, if the conflict graph maximum degree is bounded by a constant, the number of clauses per node is a constant and, if the conflict graph maximum degree is $O(\log L)$, the number of clauses per node is linear in L, so that the total number of clauses generated is polynomial.

D. Restricted Unilateral Feasibility as a Graph Problem

If we add additional restrictions to the unilateral feasibility problem, it becomes possible to formulate it as a simple graph problem. For this formulation only, we assume that no spatial multiplexing is performed, i.e. that every link is either inactive or carries exactly one stream. We also assume that the conflict graph is symmetric, i.e. if link l_i's transmission causes interference on the receiver of link l_j, then the transmission of l_j also causes interference on the receiver of l_i. Next, we assume that every node has the same number of antenna elements, which we denote by k. Finally, we assume that interference between two links that have an edge between them in the conflict graph is completely handled by one link or the other. By this, we mean that if l_i handles the interference, then the transmitter of l_i nulls its signal on the receiver of l_j and the receiver of l_i cancels interference coming from the transmitter of l_j. Instead, if l_j handles the interference, then

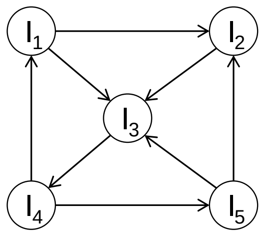

Fig. 2. Example of a Valid Conflict Graph Orientation for the Restricted Feasibility Problem with $k = 3$

the transmitter of l_j and the receiver of l_j are the two nodes assigned to do the interference cancellation.

With this restricted version of the problem, we can consider an IC assignment as an orientation of the conflict graph. In the example where l_i and l_j have an edge in the conflict graph and l_i is assigned to cancel the interference between the two links, then the edge in the conflict graph is directed from l_i to l_j. If l_j is assigned to cancel the interference, then the opposite orientation is given to the conflict graph edge. In this manner, the problem of whether a stream vector of all 1's is feasible is equivalent to asking whether there is an orientation of the conflict graph such that every edge is given a direction and no vertex in the conflict graph has more than $k-1$ outgoing edges.

The graph formulation of the restricted unilateral feasibility problem is formally defined below.

Input: A set $\mathcal{L} = \{(t_1, r_1), \ldots, (t_L, r_L)\}$ of primary-interference-free active links, a stream allocation vector $S = [1, 1, \ldots, 1]$ for $\mathcal{L}$, an antenna array size k, and an undirected conflict graph $G_c = (\mathcal{L}, E_c)$.

Output: **True** if S and $\mathcal{L}$ are feasible, and **False** otherwise. S and $\mathcal{L}$ are defined to be feasible if there is an orientation of G_c, call it G_d, in which each edge of E is assigned a direction and $\forall l_i \in \mathcal{L}, d_{\text{out}}(l_i) < k$, where $d_{\text{out}}(l_i)$ is the out-degree of l_i in G_d.

An example of an orientation showing feasibility of the all 1's stream allocation vector for a network with 5 active links where every node has 3 antenna elements is given in Figure 2. In this example, l_1 and l_5 do not interfere and l_2 and l_4 do not interfere, but all other pairs of links interfere. Note that no vertex in the graph has more than 2 outgoing edges in the given orientation, meaning that every node satisfies its antenna element constraint.

The general unilateral feasibility problem can also be considered as a type of graph problem, although it is not as simply and naturally specified as when the restrictions imposed in this subsection are added. As long as the conflict graph is symmetric and all work in cancelling interference between a pair of links is done completely by one link, we can still view an interference cancellation assignment as an orientation of the conflict graph. We can weight the directed edges by the number of streams that must be cancelled and we can weight each vertex by the number of streams multiplexed on the corresponding link. The antenna element constraints then dictate that the weight of a vertex plus the sum of the weights of all of its outgoing edges does not exceed the number of antenna elements on the transmitter and receiver of the link. If the conflict graph is asymmetric, then it is represented as a directed graph. We can still use a graph model in this situation but instead of orienting edges, we must think of marking them in some other way to indicate which link is responsible for cancelling interference. If we do not assume that all of the work is done by only one of the links in an interfering pair, then the conflict graph model is not sufficient.

In the remainder of the paper, we will only use the graph model when considering the restricted problem as outlined in this subsection.

V. Complexity of Unilateral Feasibility

In this section, we evaluate the complexity of checking the unilateral feasibility of a stream allocation vector in a MIMO network. Section V-A presents the NP-completeness proof, while the subsections that follow it present special cases solvable in polynomial time, along with their analyses.

A. General Case

Theorem 2: Evaluating feasibility of a stream allocation vector S and a link set $\mathcal{L}$ over an arbitrary MIMO network is NP-complete.

Proof: The problem is clearly in NP, because given an interference cancellation assignment (A_t, A_r), it can easily be verified in polynomial time whether (A_t, A_r) supports the stream vector by verifying Conditions 1–3 of the matrix formulation. Conditions 1 and 2 each require one matrix multiplication of $L \times L$ and $L \times 1$ matrices and comparison of L scalar quantities. Condition 3 requires simply checking that $a^t_{ij} = 1$ or $a^r_{ji} = 1$ for each of the at most $L(L-1)$ edges of the conflict graph.

To complete the proof, we present a reduction from the 3SAT problem, which was loosely inspired by [20]. We recall that in the 3SAT problem we are given a set of disjunctive boolean clauses each having 3 literals, and the problem to solve is determining whether there exists an assignment of truth values to literals such that all clauses evaluate to true (*satisfying truth assignment*). Given an instance I_{3SAT} of 3SAT, we build an instance I_{Feas} of *Feasibility* as follows. We recall that an instance of I_{Feas} is obtained by defining a set of links, an assignment of radio resources (antenna elements) to each node (link endpoint), a conflict graph describing interference relationships between links, and a stream allocation vector S describing the number of streams to be transmitted on each link.

Let $c_1, \ldots, c_m$ be the m clauses and $x_1, \ldots, x_n$ the n variables in I_{3SAT}, respectively. We recall that in 3SAT each clause is formed by exactly three literals, where a literal is a variable or its negation. The high level intuition of the construction is the following. First, we define disjoint sets of links assigned to clauses and literals in I_{3SAT}, respectively. We also define another disjoint set of links whose purpose is to mimic truth values assignment to literals. Finally, we assign radio resources to link endpoints, define the conflict graph, and choose a stream allocation vector S in such a way that S is feasible in I_{feas} if and only if the corresponding truth values assignment is satisfying for I_{3SAT}.

The three sets of node disjoint[4] links are defined as follows:

- **clause links**: a set C of links corresponding to the clauses in I_{3SAT} (one link per clause, m links in total); with a slight abuse of notation, we denote by c_i the link corresponding to clause c_i in I_{3SAT}.
- **literal links**: a set Lit of links corresponding to all possible literals (2 links per variable, $2n$ links in total); again slightly abusing notation, we denote by x_j the link corresponding to literal x_j, and by $\bar{x}_j$ the link corresponding to literal $\bar{x}_j$;
- **truth assignment links**: a set $A = A_1 \cup A_2 \cup A_3$ of links corresponding to truth assignment values (3 links per variable, $3n$ links in total); we denote by a_j, a_j^T and a_j^F the links corresponding to variable x_j.

Thus, the link set of I_{Feas} is $\mathcal{L} = C \cup Lit \cup A$, which has cardinality $m + 5n$.

The number of antenna elements available at the nodes is as follows:

- for each link c_i, the transmitter has 3 antenna elements, while the receiver has 1 antenna element;
- for each link corresponding to a literal x_j (or $\bar{x}_j$), the transmitter has 1 antenna element, while the receiver has $\ell + 1$ antenna elements;
- for each link a_j, both the transmitter and receiver have $\ell + 1$ antenna elements;
- for each link a_j^T and a_j^F, the transmitter has $\ell+1$ antenna elements, while the receiver has 1 antenna element.

In the above definitions, ℓ is a constant defined as follows. Let $\hat{h}_j = \max\{h_j, \bar{h}_j\}$, where h_j is the number of clauses to which literal x_j belongs, and $\bar{h}_j$ is the number of clauses to which literal $\bar{x}_j$ belongs; we let ℓ be an arbitrary integer larger than $\max_j\{\hat{h}_j\}$.

The stream allocation vector S allocates one stream on each link in $C \cup Lit$. For links in A, we allocate ℓ streams on the a_j links, and one stream on the a_j^T and a_j^F links.

The link set, resource allocation, and stream allocation vector corresponding to the following instance of 3SAT: $c_1 = (x_1 \vee x_2 \vee x_4)$, $c_2 = (\bar{x}_1 \vee x_3 \vee \bar{x}_4)$, is reported in Figure 3. Notice that, since $\max_j\{\hat{h}_j\} = 2$ in the instance of 3SAT at hand, we arbitrarily set $\ell = 3$ in the example.

We are now left with the definition of the conflict graph G, which has a node for each link in $\mathcal{L}$. The edge set is built as follows. For each clause link c_i, we add a directed edge between c_i and the links corresponding to the three literals in clause c_i. Furthermore, we add a directed edge between each truth assignment link a_j and the links x_j and $\bar{x}_j$. Finally, we add a directed edge between each link a_j^T and the links x_j and a_j, and between each link a_j^F and the links $\bar{x}_j$ and a_j. The conflict graph corresponding to the instance of I_{3SAT} mentioned above is reported in Figure 4.

Notice that in the instance of I_{Feas} at hand only links in $Lit \cup A_1$ are subject to interference, while links in $C \cup A_2 \cup A_3$ are not interfered by other links in the network (but they generate interference to the links in $Lit \cup A_1$). Thus, receivers of links in $C \cup A_2 \cup A_3$ can always correctly decode the incoming stream with the available antenna elements. Notice also that, given the antenna elements available at the nodes, the interference cancellation task can be undertaken only by transmitter nodes of links in C, by the receiver nodes of links in Lit, by the transmitter nodes of links in $A_2 \cup A_3$, and by both transmitter and receiver nodes of links in A_1.

We now show that I_{Feas} is feasible if and only if there exists a satisfiable truth assignment for I_{3SAT}.

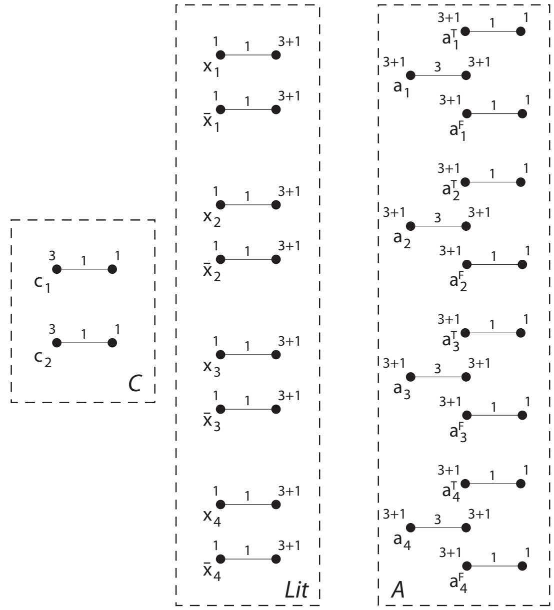

Fig. 3. Link set, resource allocation, and stream allocation corresponding to the following instance of I_{3SAT}: $c_1 = (x_1 \vee x_2 \vee x_4)$, $c_2 = (\bar{x}_1 \vee x_3 \vee \bar{x}_4)$. Each node is labeled with the number of available antenna elements. Links are labeled with the number of allocated streams in the stream allocation vector S.

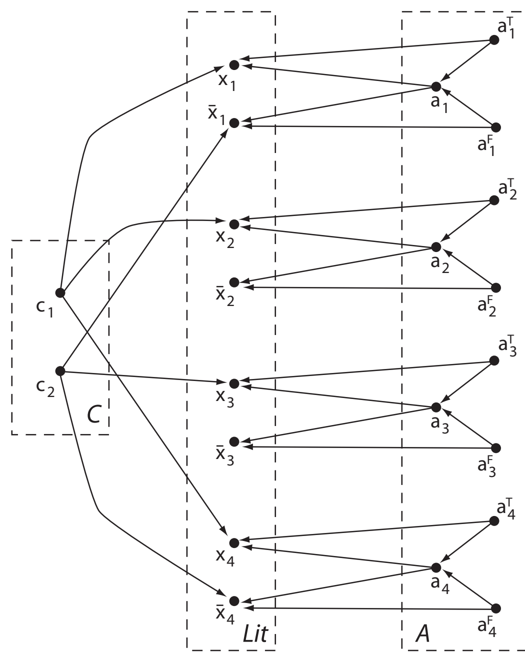

Fig. 4. Conflict graph corresponding to the following instance of I_{3SAT}: $c_1 = (x_1 \vee x_2 \vee x_4)$, $c_2 = (\bar{x}_1 \vee x_3 \vee \bar{x}_4)$.

[4]This is to ensure primary interference freeness in I_{Feas}.

If. Let us consider a satisfying truth assignment for I_{3SAT}. We assign antenna elements to cancel interference in I_{Feas} as follows. For each link a_j in A, the extra antenna element available at the transmitter is used to cancel interference at the receiver of the link corresponding to the **True** literal in the satisfying truth assignment (either x_j or $\bar{x}_j$), and the extra antenna element available at the receiver node is used to cancel interference incoming from either a_j^F (if x_j is **True** in the truth assignment) or a_j^T (otherwise). Assume w.l.o.g. that the **True** literal is x_j. The ℓ extra antenna elements at the receiver of x_j are used as follows: $h_j < \ell$ are used to cancel the interference incoming from the transmitter nodes of the links corresponding to the clauses c_i to which x_j belongs, and one DOF is used to cancel interference incoming from a_j^T. The ℓ extra antenna elements at the receiver of $\bar{x}_j$ (more in general, at the **False** literal) are all used to cancel the interference incoming from a_j. The ℓ extra antenna elements at the transmitter of link a_j^T are used to cancel interference at the receiver of link a_j. One of the ℓ extra antenna elements at the transmitter of link a_j^F is used to cancel interference at the receiver of link $\bar{x}_j$. Finally, the 2 extra antenna elements available at the transmitter nodes of links in C are used as follows: for link c_i, the extra antenna elements are used to cancel interference to the (at most two) links corresponding to **False** literals in c_i. If less than 2 literals are **False** in c_i, these extra antenna elements are left unused.

We now show that the above described IC assignment cancels all the interference, and hence the stream allocation vector S as defined above is feasible in I_{Feas}. We recall that only receivers of links in $Lit \cup A_1$ are subject to interference. Let us first consider links in Lit. Let r_j be the receiver of the link corresponding to an arbitrary literal x_j. Receiver r_j is interfered by the transmitter of link a_j, by the transmitter of link a_j^T (or a_j^F if the literal is negative), and by the transmitters of all clause links to which x_j belongs. In total, there are $h_j + \ell + 1$ interfering streams in the network. If the literal is **True** in the satisfying truth assignment, then the transmitter of link a_j cancels its interference at r_j, and r_j can use $h_j < \ell$ of its extra antenna elements to cancel interference incoming from the h_j transmitters of the clause links to which x_j belongs. Furthermore, since $h_j < \ell$, r_j is guaranteed to have at least one antenna element left to cancel the interference incoming from a_j^T. Thus, all the interference at r_j can be canceled in this case. Assume now the literal is **False** in the satisfying truth assignment. In this case, r_j has to use all its ℓ extra antenna elements to cancel interference generated by the transmitter of link a_j. Notice that r_j is still subject to interference generated by the transmitters of the clause links to which x_j belongs, and by transmitter of link a_j^T. We first notice that if x_j is **False**, then the transmitter of link a_j^T uses one of its extra antenna elements to cancel interference at r_j. Furthermore, for each clause c_i to which literal x_j belongs, it is the transmitter of link c_i that cancels its interference at r_j. Notice that the transmitter at c_i is guaranteed to have at least one antenna element available for that. In fact, each transmitter of a link in C has two extra antenna elements available for interference cancellation, and, since the truth assignment is satisfying for I_{3SAT}, we are guaranteed that at most two literals are false in any clause. We can then conclude that also in this case all the interference can be canceled at r_j. Let us now consider a link in A_1, say link a_j. The receiver ra_j of link a_j is subject to interference generated by the transmitters of links a_j^T and a_j^F. Assume w.l.o.g. that x_j is **True** in the satisfying truth assignment. In this case, the extra antenna element available at ra_j is used to cancel interference incoming from a_j^F, while the interference generated by the transmitter of link a_j^T is canceled at the transmitter side (notice that the transmitter node of link a_j^T has ℓ extra antenna elements available for that). We can then conclude that, also in this case, all the interference can be successfully canceled. This implies that a satisfying truth assignment for I_{3SAT} results in an IC assignment that supports stream allocation vector S for I_{Feas}.

Only If. Assume now that there exists a feasible IC assignment that support S in I_{Feas}. We first show that, for any feasible IC assignment, the transmitter of each link a_j must cancel its interference at the receiver of either link x_j or link $\bar{x}_j$, but not both. We first notice that the transmitter of link a_j cannot simultaneously cancel its interference at the receivers of both link x_j and $\bar{x}_j$, since that would require $\ell+2$ antenna elements at this node, which however has only $\ell+1$ antenna elements. We now prove that the transmitter of each link a_j must cancel its interference at the receiver of either link x_j or link $\bar{x}_j$. Suppose otherwise, i.e., that there exists a link a_j whose transmitter does not cancel its interference at any of the above mentioned receivers (call them r_j and $\bar{r}_j$). In this case, both receivers r_j and $\bar{r}_j$ must use all their ℓ extra antenna elements to cancel interference generated by the transmitter of link a_j. Hence, neither r_j nor $\bar{r}_j$ can use their antenna elements to cancel interference generated by link a_j^T (at r_j) and a_j^F (at $\bar{r}_j$). This implies that both transmitters of links a_j^T and a_j^F must use one antenna element to cancel their interference at r_j and $\bar{r}_j$, respectively. These transmitters are then left with $\ell-1$ antenna elements available for canceling interference, and they hence cannot cancel the interference they generate at the receiver ra_j of link a_j (which would require using ℓ antenna elements). It then follows that ra_j should cancel interference incoming from both links a_j^T and a_j^F, which would require using 2 antenna elements. However, ra_j has only 1 antenna element available, since the other ℓ antenna elements are needed to receive the ℓ streams transmitted on link a_j. We can then conclude that, for any feasible IC assignment for I_{Feas}, the transmitter of each link a_j must cancel its interference at the receiver of either link x_j or link $\bar{x}_j$.

Now, we define the truth assignment for I_{3SAT} as follows: we set literal x_j to **True** if a_j cancels its interference at the receiver of link x_j, and to **False** otherwise. It is then easy to prove that the resulting truth assignment is satisfying for I_{3SAT}. The proof proceeds by contradiction. Assume the truth assignment is not satisfying for I_{3SAT}, and let c_i be a **False** clause. Then, all the literals in c_i are false. Let r_1, r_2 and r_3 be the receivers of the respective literal links. Since all the three literals are **False**, the transmitters of the corresponding truth assignment links (call these links a_1, a_2, and a_3) cancel their interference at the receivers of the corresponding **True** literals. It follows that r_1, r_2, and r_3 must cancel the interference generated by links a_1, a_2, and a_3, respectively. For this, each of the r_j must use all the ℓ extra antenna elements available. It follows that, for all receivers r_1, r_2 and r_3, the interference generated by the transmitter of link c_i must be canceled at the *transmitter* side. The transmitter of link c_i has only 2 extra antenna elements available for interference cancellation, while 3 antenna elements in total are needed to cancel interference at receivers r_1, r_2 and r_3, resulting in an unfeasible IC allocation for S, which is a contradiction.

■

B. Receiver-Side Cancellation

When CSI is available only at the receivers and not at the transmitters, then only receiver side interference cancellation can be done. Theorem 3 states that, in this special case, the feasibility problem has polynomial time complexity.

Theorem 3: Evaluating feasibility of a stream allocation vector S and a link set $\mathcal{L}$ over a MIMO network with receiver-side-cancellation only can be done in polynomial time.

The proof of Theorem 3 can be found in [31] and is not repeated here. The essence of it is that only having interference cancellation on one side removes the choice of how to handle a given interference relationship. Thus, one can simply assign all interference cancellation to receivers and then check whether the antenna element constraints are satisfied at every node. Since this check can easily be done in polynomial time, the result follows. In a similar fashion, it can be shown that checking feasibility for the transmitter-side-cancellation only case also has polynomial time complexity.

C. Conflict Graph Maximum In-Degree and Out-Degree of Two

In this section, we consider the special case where no link interferes with, or is interfered by, more than two other links. This results in a conflict graph with maximum in-degree and maximum out-degree of two. We use the Boolean satisfiability problem formulation in analyzing this special case.

Theorem 4: If the conflict graph has maximum in-degree and maximum out-degree of two, then all clauses in the Boolean satisfiability problem formulation of unilateral feasibility have at most two literals.

Proof: It is direct to see from the Boolean satisfiability formulation, that the clauses resulting from the interference constraint all contain two literals. It therefore remains to show that the clauses resulting from the antenna element constraints contain at most two literals.

Consider an arbitrary transmitter t_i. Since the out-degree of l_i is at most two, t_i interferes with receivers on at most two other links. In the worst case, assume that t_i interferes with two receivers r_j and r_k. The choices that exist for the variables a_{ij}^t and a_{ik}^t depend on k_i^t and the numbers of streams carried by l_i, l_j, and l_k, which are fixed by the given stream allocation vector S. There are three cases to consider.

Case 1: t_i has enough antenna elements to cancel its interference at both r_j and r_k simultaneously

In this case, any combination of values for a_{ij}^t and a_{ik}^t is possible and no clauses are generated for the antenna element constraints.

Case 2: t_i cannot cancel interference at one of the two receivers, regardless of what it does with the other receiver

Without loss of generality, assume that t_i cannot cancel interference at r_j, regardless of the value of a_{ik}^t. Thus, s_j is greater than $k_i^t - s_i$, so t_i simply does not have enough antenna elements to multiplex s_i streams and simultaneously cancel interference at r_j. This adds the clause $\overline{a_{ij}^t}$ with one literal, which simply says that a_{ij}^t must be zero (false), independent of a_{ik}^t. (Similarly, if t_i cannot cancel interference at r_k, regardless of a_{ij}^t, the clause $\overline{a_{ik}^t}$ is added.)

Case 3: t_i cannot cancel interference at either receiver

This is actually Case 2, happening simultaneously at both r_j and r_k. Here, we simply have two clauses with one literal each, $\overline{a_{ij}^t}$ and $\overline{a_{ik}^t}$.

Case 4: t_i can cancel interference at either of the two receivers individually but not at both simultaneously

In this case, the values of a_{ij}^t and a_{ik}^t that satisfy the antenna element constraints are given by the following sum of products expression:

$$\overline{a_{ij}^t}\ \overline{a_{ik}^t} \vee a_{ij}^t\ \overline{a_{ik}^t} \vee \overline{a_{ij}^t}\ a_{ik}^t$$

In other words, the only invalid combination is when both a_{ij}^t and a_{ik}^t are true. Using the procedure outlined in Section IV-C, this is converted to the following product of sums expression:

$$\overline{a_{ij}^t} \vee \overline{a_{ik}^t}$$

Thus, one clause with two literals is added in this case.

In each of the above cases, the number of literals appearing in the clauses added by the antenna element constraint at each transmitter is at most two. The exact same argument can be used to show that the maximum number of literals in a clause added by the antenna element constraint at each receiver is at most two. Therefore, every clause in the satisfiability problem formulation has at most two literals.

■

Corollary 1: Evaluating feasibility of a stream allocation vector S and a link set $\mathcal{L}$ over a MIMO network, where the conflict graph has maximum out-degree and maximum in-degree of two, can be done in polynomial time.

From Theorem 4, feasibility with a conflict graph maximum degree of two is an instance of the 2SAT problem. Furthermore, as discussed in Section IV-C, the total number of clauses in the satisfiability problem is polynomial when the graph degree is bounded by a constant. Since 2SAT can be solved in polynomial time when the number of clauses is polynomial [7], the corollary follows.

D. Antenna Array Size $k = 2$

Another interesting special case is when every node in the network has $k = 2$ antenna elements. For this case, we also assume that the conflict graph is symmetric and that all interference cancellation between two interfering links is handled entirely by one of the two links. Checking the unilateral feasibility of such a MIMO network can be done in polynomial time, even when transmitters and receivers are both capable of performing IC and independent of the maximum degree of the conflict graph. This result is stated in Theorem 5.

The proof of this theorem uses the graph formulation of the unilateral feasibility problem.

Theorem 5: Evaluating feasibility of a stream allocation vector S and a link set $\mathcal{L}$ over a MIMO network, where every node has $k = 2$ antenna elements and the conflict graph is symmetric, can be done in polynomial time.

Proof: The proof is constructive, i.e., we describe a polynomial time algorithm that, given inputs S and $\mathcal{L}$, returns **True** if and only if stream allocation vector S is feasible for link set $\mathcal{L}$ and returns **False** otherwise. The algorithm first checks whether $\mathcal{L}$ is primary-interference-free (as in the proof of Theorem 3). If $\mathcal{L}$ is not primary-interference-free, the algorithm returns **False**, otherwise it continues.

Let the conflict graph of the network be $G_c = (\mathcal{L}, E_c)$. We first eliminate links with zero or two streams. Inactive links (with zero streams) do not impact the problem and can be omitted from $\mathcal{L}$ and G_c. Let $\mathcal{L}_2 = \{l_i \in \mathcal{L} : s_i = 2\}$, i.e. $\mathcal{L}_2$ is the set of links that carry two streams. The algorithm checks whether any link in $\mathcal{L}_2$ interferes with any other active link. If such a link is found in $\mathcal{L}_2$, the algorithm returns **False**. If, instead, all links in $\mathcal{L}_2$ are isolated vertices in G_c, the algorithm drops all links in $\mathcal{L}_2$ from G_c.

Note that, if the algorithm has not returned **False** at this point, we are left with an instance of the graph problem formulation given in Section IV-D. This is because all links that remain at this point carry exactly one stream and all other assumptions match those presented in Section IV-D.

Denote the remaining links, all carrying one stream, by $\mathcal{L}_1$. Denote by G_1 the subgraph of G_c induced by node set $\mathcal{L}_1$, i.e. the conflict graph made up of only the links carrying one stream. Let $G^1, \ldots, G^h$ be the connected components of graph G_1. The algorithm checks whether for each $G^i = (\mathcal{L}^i, E^i)$, inequality $|E^i| \leq |\mathcal{L}^i|$ is satisfied; if the inequality is not satisfied for any of the G^i, the algorithm returns **False**, otherwise it returns **True** and terminates.

Clearly the above algorithm has polynomial time complexity. We now prove that, when the algorithm returns **False** on input S and $\mathcal{L}$, the stream allocation vector S is infeasible for $\mathcal{L}$. To prove this, we observe that the algorithm returns **False** if only if one of the following conditions hold:

$i)$ Set $\mathcal{L}$ is not primary-interference-free.

$ii)$ $\mathcal{L}_2$ contains at least one link, which is not an isolated vertex in G_c; denote such a link by l_i, and suppose it is adjacent to link l_j in the conflict graph. Since l_i carries two streams and every node has only two antenna elements, t_i and r_i cannot perform any interference cancellation. Link l_j carries at least one stream, and t_j and r_j therefore cannot cancel interference with l_i without violating their antenna element constraints (since $k = 2$ for every node). Hence, Condition (3) for feasibility cannot be satisfied for links l_i, l_j unless Conditions (1) and (2) are violated. This implies that stream assignment S is not feasible for link set $\mathcal{L}$.

$iii)$ There exists a connected component G^j of graph G_1 such that $|E^j| > |\mathcal{L}^j|$. A simple counting argument can be used to prove that S is not feasible for $\mathcal{L}$: for each link $l \in \mathcal{L}^j$, one antenna DOF is available both at the transmitter and at the receiver side. Thus, $2|\mathcal{L}^j|$ antenna DOFs in total are available to cancel interference within G^j. On the other hand, cancelling interference between any two adjacent links l_i, l_j in the conflict graph requires using 2 antenna DOFs: one for cancelling interference generated by t_i on r_j, and one for cancelling interference generated by t_j on r_i. Thus, $2|E^j|$ antenna DOFs in total are needed to cancel all interference between the $|\mathcal{L}^j|$ links in G^j. Hence, $|E^j| > |\mathcal{L}^j|$ implies that not enough MIMO resources are available within G^j to completely cancel interference, which proves that stream allocation vector S is infeasible for $\mathcal{L}$.

The next step is to prove that whenever none of conditions $i), ii), iii)$ hold on given input $S, \mathcal{L}$, then stream allocation vector S is feasible for $\mathcal{L}$, which implies correctness of our feasibility evaluation algorithm, which returns **True** in this situation. We prove this by showing a construction (IC assignment) that makes S feasible for $\mathcal{L}$ when none of the conditions $i), ii), iii)$ are satisfied. This construction uses the graph model of the feasibility problem and orients edges of the conflict graph, as discussed in Section IV-D.

If condition $iii)$ is not satisfied, then $|E^j| \leq |\mathcal{L}^j|$ for each connected component G^j of G_1. Observe that IC assignments for the G^js can be built independently, since links in different G_1 connected components do not interfere with each other. Hence, it suffices to show the construction for a single G^j, making the overall construction the result of the composition of IC assignments for the individual connected components. Given that G^j is connected and $|E^j| \leq |\mathcal{L}^j|$, the topology of G^j can take only one of the four following forms: $a)$ a single vertex, $b)$ a tree, $c)$ a simple cycle, $d)$ a connected graph containing a single simple cycle.

If G^j is of type $a)$, no interference cancellation is required. If G^j is a tree (type $b)$), perform the following procedure:

1. Designate some vertex in $\mathcal{L}^j$ to be the root.
2. For every edge $(l_i, l_k) \in E^j$, where l_k is deeper in the tree i.e. l_i is the parent and l_k is the child, direct the edge from l_k to l_i.

Since every vertex in a tree (except the root) has a single parent, each non-root node in G^j is assigned one outgoing edge. Since every edge is given a direction in this procedure (all interference is covered) and every node has at most $k-1 = 1$ outgoing edges, the stream allocation is feasible within G^j.

Now consider case $c)$. Here, it is sufficient to give either clock-wise or counterclock-wise orientation to all edges in E^j. Again, every edge is oriented and no node has more than one outgoing edge. Therefore, this construction makes S feasible (when restricted to G^j).

Finally, consider case $d)$. In this case, there is a single cycle with one or more tree components hanging off of the cycle. At each cycle node where a tree component hangs off, the node has degree greater than 2 in G_1 (two cycle edges and one or more edges into the tree component). We start by designating every vertex in $\mathcal{L}^j$ that is contained in the simple cycle and has degree higher than 3 as the root of the tree component it belongs to. Edges are then oriented by combining the construction for case $b)$ within the trees, with the construction for case $c)$ along the single simple cycle contained in G^j. It is clear that all nodes that are not root nodes still have at most one outgoing edge, as per the part b) and c) constructions. Note also that since the root nodes do not have parents in the part b) construction, they are not assigned any outgoing edges during that construction. Furthermore, each root node is assigned one outgoing edge in the cycle construction from part c). Therefore, these nodes also have one outgoing edge in the final oriented version of G^j. Since, again, all edges are covered and every node has at most one outgoing edge, the stream allocation is feasible within G^j.

Since these constructions are applied independently within each connected component and the components do not have any edges between them, every node in the overall oriented version of G_1 has at most one outgoing edge and the stream allocation is feasible overall. ■

To summarize, a stream allocation for a network with symmetric conflict graph G_c and with two antenna elements on every node is feasible if and only if all links carrying two streams are isolated vertices in G_c and every other connected component of G_c contains at most one simple cycle (or equivalently, has an average vertex degree of at most two).

VI. Feasibility Heuristics

A. Simple Greedy and Extended Greedy

Given that the general unilateral feasibility problem is NP-complete, heuristics for checking feasibility are necessary. Perhaps the most obvious heuristic is to check whether all interference can be eliminated by greedily allocating MIMO resources for interference cancellation. One possible implementation of the algorithm is as follows. Sort the links in order of non-increasing number of allocated streams. Begin with the first link and use its antenna DOFs to cancel interference on the links with which it interferes one by one until all its resources are used. Then, move on to the next link and continue until all interference is eliminated or all resources are exhausted, whichever comes first. If all interference can be removed with the available resources in the network, the allocation vector is declared to be feasible. We refer to this approach as Algorithm Simple Greedy. The time complexity of Algorithm SimpleGreedy is dominated by the time to initially sort the stream allocation vector and it is therefore $O(L \log L)$, where L is the number of active links.

In experimenting with Algorithm Simple Greedy, we found that it tends to concentrate resources among small groups of nodes, rather than more evenly distributing the resources across links in the network, and this causes it to frequently label feasible vectors as infeasible. To remedy this problem, we developed the algorithm in Figure 5, which we refer to as Algorithm Extended Greedy. This algorithm, when considering multiple candidate links, all carrying an equal number of streams, on which to cancel interference, chooses a target link in a way that tends to produce a better distribution of resources and outperforms Algorithm Simple Greedy. In Figure 5, note that the standard notation $< V, W >$ is used to represent the inner product of vectors V and W and that I is the identity matrix. The time complexity of Algorithm Extended Greedy is $O(L^2)$, because each operation inside the for loop requires $O(L)$ time and there are L iterations of the loop.

Input: Stream allocation vector S, link set $\mathcal{L}$, K^t, K^r, conflict graph $G_c = (V_c, E_c)$
Output: feasible $\in \{$ **true**, **false**$\}$, A^t, A^r

Order S in non-increasing fashion. Permute vertices in G_c accordingly.
$A^t = A^r = I_{L \times L}$
for $i = 1 \rightarrow L$
if $< A^t(i, 1:i), S^{1:i} > \leq K^t_i$, distribute $1's$ in $A^t(i, G_c(i, i+1:L))$ greedily, giving equal priority to columns of equal weight such that $< A^t_i, S > \leq K^t_i$
if $< A^r(i, 1:i), S^{1:i} > \leq K^r_i$, distribute $1's$ in $A^r(i, G_c(i, i+1:L))$ greedily, giving equal priority to columns of equal weight such that $< A^r_i, S > \leq K^r_i$
$A^r(m, i) = 1 - A^t(i, m)$ and $A^t(m, i) = 1 - A^r(i, m)$ $\forall m \geq i+1 : \ (i, m) \in E_c$
end for
if $(A^t S \leq K^t \wedge A^r S \leq K^r)$ **then** feasible **= true**, **else** feasible **= false**

Fig. 5. Algorithm Extended Greedy

Both Algorithm Simple Greedy and Algorithm Extended Greedy are safe, in that they always label infeasible vectors as infeasible. However, they are non-optimal in that they each label some feasible vectors as infeasible. The accuracy of the two heuristics is evaluated in Section VI-D.

B. Distributed Implementation

Algorithms SimpleGreedy and ExtendedGreedy can be implemented fairly simply in a distributed fashion using a token passing algorithm among the transmitters of the links on which a stream allocation is being considered. Any node that initially has the token can select the links on which it wants to cancel interference according to the SimpleGreedy or ExtendedGreedy method. When some node, other than the first, receives the token, it can also apply the greedy technique, being sure to first cancel interference on all links it has a conflict with that have already selected their cancellations but did not choose to cancel with this node's link. At the end of the token passing cycle, if all necessary cancellations have been assigned, the stream allocation vector is feasible. Otherwise, the last node in the cycle labels it as infeasible.

A problem with token passing is that it serializes feasibility calculation. Next, we sketch an alternate approach that is parallel but more complex due to looser coordination between cancellation assignments of different nodes. Nodes can start feasibility checking at any point, possibly in parallel with other nodes. A node that starts the checking greedily constructs an interference cancellation assignment for itself and sends it to all of its neighbors, e.g. with a single broadcast message. After initiation, the algorithm basically proceeds like the token passing algorithm. By this, we mean that nodes that do cancellation assignments factor in all assignments that they have received from other nodes by that time, which might dictate that they perform certain cancellations and then they greedily assign their remaining resources according to SimpleGreedy or ExtendedGreedy. After choosing an assignment, a node combines it with all other information it has received about assignments of other nodes and broadcasts the information to its neighbors. Nodes also rebroadcast new information they receive from other nodes so that all cancellation assignments are eventually disseminated to all nodes.

With this approach, because the cancellation assignments are only loosely coordinated, it could happen that two links i and j assign themselves to cancel interference with each other. Without loss of generality, assume $i < j$. Once a node that is part of link j detects that it and link i have chosen to cancel interference on each other, link j (the higher-numbered link) will replace its cancellation of i with cancellation of a higher-numbered link. Since these redundancies are always resolved by moving resources to higher-numbered links, eventually the resolution process will end. If all necessary cancellations have been assigned after a sufficient number of steps has elapsed for all nodes to do cancellation assignments and resolve redundancies, then the stream allocation vector is feasible. Otherwise, any node that detects missing cancellations at that point can label the vector as infeasible.

C. *isFeasibile3 Heuristic for $k \leq 3$*

Consider a MIMO network where every node has $k = 3$ antenna elements. In this section, we extend the approach of Section V-D, which solved the $k = 2$ case in polynomial time. For $k = 3$, this approach does not yield a polynomial-time exact solution but it does lead to an efficient heuristic. Since the Section V-D approach uses the graph formulation of the feasibility problem, we adopt the assumptions of that model in this section, as well. In particular, we assume the conflict graph is symmetric. As in the proof of Theorem 5, we include in the undirected conflict graph G only active links that are not at full capacity. Thus, we eliminate links with zero or three streams. For feasibility, all links with three streams must be isolated vertices in G, which can be easily checked in polynomial time.

We begin by analyzing the case where the active links not at full capacity all carry one stream. This matches the problem setting of Section IV-D and, following the formulation in that section, the problem can be stated as:

Can the edges of the conflict graph be directed such that every vertex has at most two outgoing edges?

Theorem 6 states that checking feasibility when $k = 3$ and where every link carries exactly one stream, is equivalent to checking whether the conflict graph G contains a subgraph of average degree greater than 4.

Theorem 6: Let D_1 be the property of a graph $G = (V, E)$, whereby every vertex induced subgraph of G has an average degree at most equal to four. D_1 is necessary and sufficient for the edges of G to be directed such that every vertex has at most two outgoing edges.

Proof:

Necessary condition: Assume the edges of G can be directed such that every vertex has at most two outgoing edges. We will prove that Property D_1 holds, i.e. that all subgraphs of G have average degree no greater than 4. Consider an arbitrary subgraph G_1 with n_1 vertices. Since in some complete labeling of G each vertex has at most two outgoing edges, the total number of edges in the subgraph can be at most $2n_1$. Since each edge is incident on two vertices, the average degree is at most $\frac{2 \cdot 2n_1}{n_1} = 4$.

Sufficient condition: Suppose the given graph G satisfes D_1. We prove the sufficient condition through construction, by determining a direction for all of G's edges such that every vertex has at most two outgoing edges. The construction is described in Procedure Proc I, which is given after the following definitions. The quantities defined by these definitions are dynamic, i.e., they are recalculated dynamically as the construction proceeds.

1) Let the quantity n_v denote the number of remaining edges that can be marked as outgoing from vertex v in V. Initially, $n_v = 2$ for all vertices v. At any intermediate point during the construction, this value equals two minus the number of edges that have already been marked as outgoing from v. The construction does not allow more than 2 edges to be marked as outgoing for any vertex and, therefore, $0 \leq n_v \leq 2$ always holds. [5]
2) Define for any subgraph $G_{sub} = (V_{sub}, E_{sub})$ of G the quantity *ED* (*ED stands for 'extra DOF's'*). Let $E_{um} \subseteq E_{sub}$ be the set of edges of G_{sub} that are not yet marked with a direction.
$$ED(G_{sub}) = \sum_{v \in G_{sub}} n_v - |E_{um}|$$
Property D_1 implies that ED *is greater than equal or to zero for every subgraph*, at the beginning of Procedure Proc 1.
3) Define for any edge (v, v') in E, the boolean quantity *DO* (*DO stands for 'directable outwards'*) :
$$DO(v, v') = \bigwedge_{\forall G_{sub}} \big(ED(G_{sub}) > 0\big)$$
where G_{sub} is a vertex-induced sub-graph of G/v' such that it contains vertex v and $\bigwedge$ refers to the Boolean AND operation.

The DO definition is illustrated in Figure 6. If all subgraphs containing v but not v' have "extra DOFs", then it is safe to direct the edge (v, v') outwards.

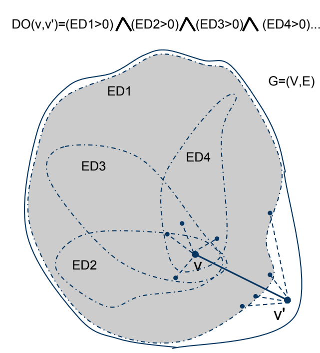

Fig. 6. Illustration of definition: $DO(v, v')$

[5]The "number of remaining edges n_v that can be marked as outgoing from v", refers to the number of DOF's that are available for interference suppresion at the transmitter and at the receiver of link v in the MIMO feasibility problem.

Keeping the above definitions in mind, apply Procedure Proc 1 to a graph $G = (V, E)$ which satisfies D_1. By definition, *every* subgraph of G has an ED value greater than equal to zero at the beginning of the procedure.

Begin Procedure Proc I

Input: $G = (V, E)$ satisfying Property D_1
Output: $f : E \to \{0, 1, \ldots, n\}$, where $f((u_i, u_j)) = i$ indicates that the edge is directed from u_i to u_j and $f((u_i, u_j)) = j$ indicates that the edge is directed from u_j to u_i

1. **Repeat:** If any vertex v_i in V has all but p edges marked as incoming, where $p \in \{1, 2\}$, mark these p edges as outgoing, i.e., $f((v_i, v_j)) = i$ for these edges, and set $n_{v_i} = n_{v_i} - p$
   Until: No such vertex exists
2. V_n = the set of all vertices with at least one unmarked edge
3. **while** there exists a vertex v_i in V_n with $n_{v_i} = 2$ (i.e. with *no* outgoing edges assigned)

3a. Let $v_j,\ j = 1, 2, 3...$ be the neighbors of v_i connected by unmarked edges
3b. $j = 1$
3c. **while** $(n_{v_i} = 2)$
  if $(DO(v_j, v_i) =$ TRUE$)$
    Mark (v_i, v_j) as outgoing, i.e. $f(v_i, v_j) = i$
    $n_{v_i} = n_{v_i} - 1$
  else
    Mark (v_i, v_j) as incoming, i.e. $f(v_i, v_j) = j$
    $n_{v_j} = n_{v_j} - 1$
  end if
  if (v_i has only two unmarked edges and $n_{v_i} = 2$)
    Mark the remaining two edges incident to v_i as outgoing, i.e. $f((v_i, v_j)) = i$ for these two edges
    $n_{v_i} = n_{v_i} - 2$
  end if
  $j = j + 1$
 end while
3d. **Repeat:** If any vertex v_i in V has all but $p \leq n_{v_i}$ edges marked as incoming, mark these p edges as outgoing, i.e. $f((v_i, v_j)) = i$ for these edges, and set $n_{v_i} = n_{v_i} - p$
 Until: No such vertex exists
3e. **end while**

4. **while** there exists a vertex v_i in V_n with $n_{v_i} = 0$
   mark the remaining unmarked edges incident to v as incoming, i.e., $f(v_i, v_j) = j$ for these edges
   end while
5. Let $G' = (V', E')$ be the graph that results from removing all marked edges from G and removing all vertices that have every edge marked
6. Direct the edges of G' according to the procedure used for the case $k = 2$.

End Procedure

The analysis of this procedure is presented in the Appendix, including the derivation of the following results.

Lemma L1: *$(DO(v, v') \vee DO(v', v))$ is always equal to TRUE at every iteration of Step 3c.*

Invariant 1 - $ED \geq 0$ for every subgraph of G at every iteration of Step 3, including after the final iteration.

Lemma L2: *At every stage of the procedure, a vertex $v \in V_n$ can have at most n_v neighbors that are contained in subgraphs that exclude v and with ED value equal to zero. Moreover, if v has exactly n_v such neighbors, then it is part of an average-degree-four-subgraph in the original graph G.*

Invariant 2 - $n_v \geq 0$ for every $v \in V$ at every stage of the procedure.

Lemma L3: *If the input graph G satisfies Property D_1 then every connected component of the resulting graph $G' = (V', E')$ at the end of Step 4 will be a tree or a simple cycle.* ■

We use Procedure Proc I, in the proof of Theorem 6, as a starting point to build an approximate test for feasibility when the array size is limited to three antennas. The algorithm is called *isFeasible3* and it works by approximately computing the value of $DO(v_j, v_i)$ in step 3*c*. It also extends the approach to deal with links that carry either one or two streams. The pseudo-code for isFeasible3 is omitted due to space constraints, but it is very similar in structure to Procedure Proc I. Algorithm isFeasible3 runs in $O(d^2 L) = O(L^3)$ time, where L is the number of active links and d is the maximum degree of the conflict graph. The running time is dominated by the while loop in Step 3 of Procedure Proc I and its inner while loop in Step 3c. The outer loop runs at most L times. For each considered link, the inner loop runs once for each neighbor of the link in the conflict graph (at most d times) and the procedure to approximate $DO(v_j, v_i)$ is $O(d)$.

Unfortunately, isFeasible3 appears to be quite centralized in nature and, at this time, we do not see an efficient way of implementing it in a distributed manner.

D. Accuracy of Feasibility Heuristics

We begin by comparing the accuracies of Algorithms Simple Greedy and Extended Greedy. The first set of results assumes a uniform antenna array size k on every node and a single collision domain. In a single collision domain, every link interferes (strongly) with every other link and the conflict graph is symmetric and complete. In this situation, the maximum number of active links has been shown to be $2k - 1$ [30]. We study $k = 8, 12, 16$. To evaluate the accuracy of the heuristics, we calculated the entire feasible space for network sizes up to 15 links, using a brute-force algorithm. The results are shown in Figure 7. Note that the Extended Greedy heuristic is significantly more accurate than Simple Greedy. Extended Greedy is inaccurate at most 5% of the time with $k = 8$ and $k = 12$ and at most 10% of the time with $k = 16$, for the network sizes studied here.

Greedy algorithms do not work as well when antenna array sizes are highly variable. Figure 8 shows the accuracies of the two greedy heuristics when antenna array sizes are randomly chosen between 2 and 8 for every node, where each array size is equally likely to be any of the 7 possible values. Here, the inaccuracy of Extended Greedy peaks at about 13%, which is more than twice the peak value when array sizes are uniform. However, Extended Greedy is still significantly better than Simple Greedy, which has a peak inaccuracy of 18%.

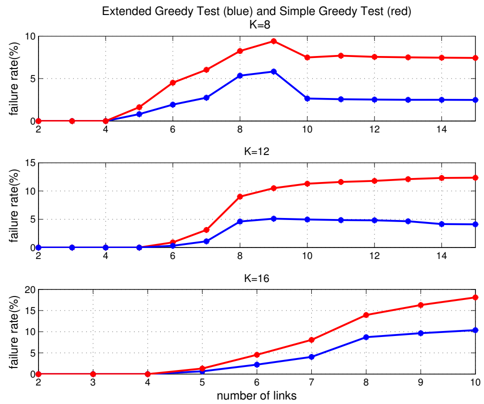

Fig. 7. Failure Rates of Simple Greedy and Extended Greedy Heuristics: Single Collision Domain, Uniform Array Size

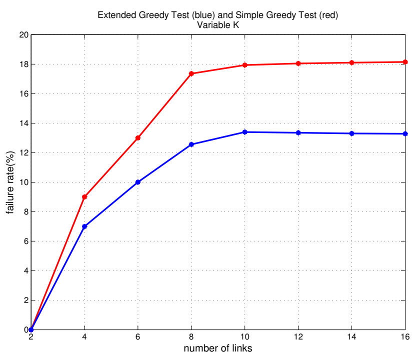

Fig. 8. Failure Rates of Simple Greedy and Extended Greedy Heuristics: Single Collision Domain, Array Size Random between 2 and 8

Both algorithms perform better when array sizes are non-uniform but multiples of a base value. In Figure 9, array sizes are randomly set to either 4 or 8. Here, Extended Greedy's peak inaccuracy is less than 2% up to 12 links.

Next, we consider an antenna array size of 3 so that we can evaluate Algorithm isFeasible3's performance. With a single collision domain, this would only allow us to consider network sizes up to 5 links, so in this set of results, we relax the single collision domain assumption. Here, we distributed links in order to produce conflict graphs with a certain average degree.

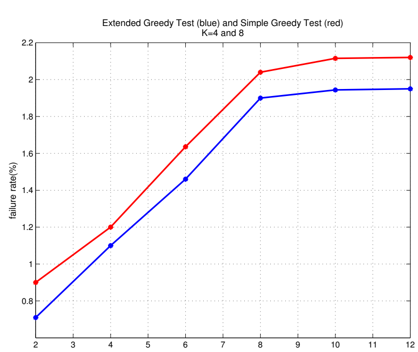

Fig. 9. Failure Rates of Simple Greedy and Extended Greedy Heuristics: Single Collision Domain, Array Size either 4 or 8

We varied the average degree to range from low interference (small conflict graph degree) to high interference (high conflict graph degree). Interference is assumed to be symmetric so that the conflict graph model used by isfeasible3 applies. In these results, we also considered an improved version of Extended Greedy, where after every iteration of the for loop, a round of constraint propagation was performed. In constraint propagation, any links that have only one remaining interfering link try to cancel the interference and, if they do not have sufficient antenna DOFs to do so, cancellation responsibility is assigned to the other link. In addition, for any links that have exhausted their antenna DOFs, cancellation responsibility is assigned to the links that interfere with the exhausted link. After each iteration, these assignments are propagated as far as possible before beginning the next iteration.

Results for average conflict graph degrees of 5.5 and 7.5 are shown in Figures 10a and 10b, which have logarithmic scales on the y axes. isFeasible3 is inaccurate at most 0.005% of the time whereas the improved Extended Greedy is inaccurate at most 7.0% of the time for an average conflict graph degree of 5.5 at 20 links. These numbers are respectively 0.2% and 13% for an average conflict graph degree of 7.5. Thus, the graph model for unilateral feasibility yields a very accurate feasibility test for $k = 3$. Extending this approach to handle larger antenna array sizes is a subject for future research.

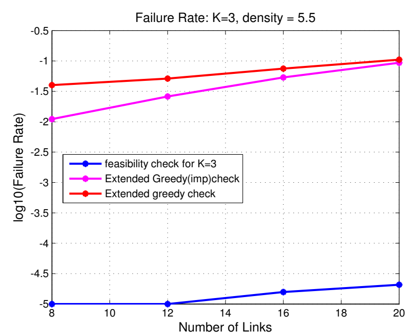

(a) Average Conflict Graph Degree = 5.5

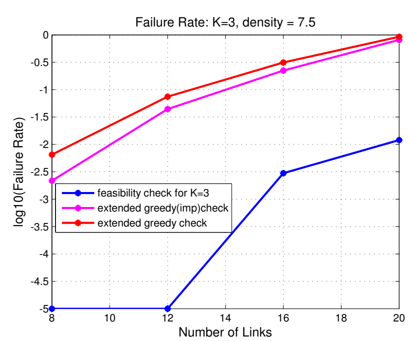

(b) Average Conflict Graph Degree = 7.5

Fig. 10. Failure Rates for Extended Greedy and isFeasible3

VII. Conclusion

We studied the feasibility problem in MIMO networks with unilateral interference cancellation. Despite proving that the unilateral feasibility problem is NP-complete in the general case, we showed that it has polynomial time complexity for several important special cases such as single-sided interference cancellation, small array sizes, and small conflict graph degrees. We have also presented two computationally efficient heuristic algorithms that exhibit good accuracy in testing for feasibility in more general MIMO networks.

ACKNOWLEDGMENT

D. Blough was supported by Grant CNS–1017248 from the National Science Foundation. P. Santi was supported by the MIMOnet Project of the Italian ccTLD Registration Authority.

REFERENCES

[1] M. Arslan, K. Sundaresan, S.V. Krishnamurthy, and S. Rangarajan, "Design and Implementation of an Integrated Beamformer and Uplink Scheduler for OFDMA Femtocells," *Proc. of MobiHoc,* pp. 215–224, 2012.

[2] J. Barry, E. Lee, and D. Messerschmitt, *Digital Communication,* Springer, 2003.

[3] R. Bhatia and L. Li, "Throughput Optimization of Wireless Mesh Networks with MIMO Links," *Proc. of Infocom,* pp. 2326–2330, 2007.

[4] D. Blough, G. Resta, P. Santi, R. Srinivasan, and L.M. Cortés–Peña, "Optimal One-Shot Scheduling for MIMO Networks," *Proc. of SECON,* pp. 377–385, 2011.

[5] V. Cadambe and S. Jafar, "Interference Alignment and Degrees of Freedom of the K-user Interference Channel," *IEEE Trans. on Information Theory,* Vol. 54, pp. 3425–3441, Aug. 2008.

[6] P. Casari, M. Levorato, and M. Zorzi, "DSMA: An Access Method for MIMO Ad Hoc Networks based on Distributed Scheduling," *Proc. of International Wireless Communications and Mobile Computing Conference,* 2006.

[7] S. Cook, "The Complexity of Theorem-Proving Procedures," *Proc. of STOC,* pp. 151–158, 1971.

[8] L.M. Cortés–Peña, D. Blough, and J. Barry, "The Performance Loss of Unilateral Interference Cancellation," *Proc. of ICC,* pp. 5718–5723, 2012.

[9] E. Gelal, G. Jakllari, and S.V. Krishnamurthy, "Exploiting Diversity Gain in MIMO Equipped Ad hoc Networks," *Proc. of Asilomar Conference on Signals and Systems,* 2006.

[10] A. Goldsmith, European Wireless 2010, Keynote presentation.

[11] K. S. Gomadam, V. R. Cadambe, and S. A. Jafar, "A Distributed Numerical Approach to Interference Alignment and Applications to Wireless Interference Networks," *IEEE Transactions on Information Theory,* Vol. 57, pp. 3309–3322, June 2011.

[12] B. Hamdaoui and K.G. Shin, "Characterization and Analysis of Multi-Hop Wireless MIMO Network Throughput", *Proc. of MobiHoc,* pp. 120–129, 2007.

[13] M. Hu and J. Zhang, "MIMO Medium Access Control and Routing in Ad Hoc Networks: A Holistic Perspective," *Proc. of Workshop on High Performance Switching and Routing,* pp. 236–240, 2005.

[14] R. Illtis, S. Kim, and D. Hoang, "Noncooperative Iterative MMSE Beamforming Algorithms for Ad Hoc Networks," *IEEE Trans. on Communications,* Vol. 54, pp. 748–759, April 2006.

[15] S. Jafar and S. Shamai, "Degrees of Freedom Region of the MIMO X Channel," *IEEE Transactions on Information Theory,* Vol. 54, pp. 151–170, Jan. 2008.

[16] D. Kim, M. Torlak, "Optimization of Interference Alignment Beamforming Vectors," *IEEE Journal on Selected Areas in Communications,* Vol. 28, pp. 1425–1434, December 2010.

[17] J. Liu, Y. Shi, and Y.T. Hou, "A Tractable and Accurate Cross-Layer Model for Multi-hop MIMO Networks," *Proc. of ICC,* 2010.

[18] M. Maddah-Ali, A. Motahari, and A. Khandani, "Communication over MIMOXchannels: Interference Alignment, Decomposition, and Performance Analysis," *IEEE Transactions on Information Theory,* Vol. 54, pp. 3457–3470, August 2008.

[19] M. Maddah-Ali, A. Motahari, and A. Khandani, "Signaling over MIMO Multibase Systems: Combination of Multiaccess and Broadcast Schemes," *Proc. of International Symposium on Information Theory,* 2006.

[20] B. Mumey, J. Tang, and T. Hahn, "The Stream Control Problem in Multihop Wireless Networks with MIMO Links," *Proc. of ICC,* 2008.

[21] F. Negro, S. Shenoy, I. Ghauri, and D. Slock, "On the MIMO Interference Channel," *Proceedings of the Information Theory and Applications Workshop,* pp. 1–9, 2010.

[22] H. Ning, C. Ling, and K.K. Leung, "Feasibility Condition for Interference Alignment With Diversity," *IEEE Transactions on Information Theory,* Vol. 57, pp. 2902–2912, May 2011.

[23] J. Park, A. Nandan, M. Gerla, and H. Lee, "Space-MAC: Enabling Spatial Reuse using MIMO Channel-Aware MAC," *Proc. of ICC,* pp. 3642–3646, 2005.

[24] S. Peters and R. Heath, "Cooperative Algorithms for MIMO Interference Channels," *IEEE Transactions on Vehicular Technology,* Vol. 60, pp. 206–218, Jan. 2011.

[25] S. Peters and R. Heath, "Interference Alignment via Alternating Minimization," *Proc. of International Conference on Acoustics, Speech and Signal Processing,* pp. 2445–2448, 2009.

[26] M. Razaviyayn, M. Sanjabi, and Z.-Q. Luo, "Linear Transceiver Design for Interference Alignment: Complexity and Computation," *IEEE Transactions on Information Theory,* to appear. (available as arXiv:1009.3481v1)

[27] I. Santamaria, O. Gonzalez, R. Heath, and S. Peters, "Maximum Sum-Rate Interference Alignment Algorithms for MIMO Channels," *Proc. of Globecom,* pp. 1–6, 2010.

[28] D. Schmidt, W. Itschick, and M. Honig, "Beamforming Techniques for Single-Beam MIMO Interference Networks," *Proc. of Allerton Conference,* pp. 1182–1187, 2010.

[29] Y. Shi, J. Liu, C. Jiang, C. Gao, and Y.T. Hou, "An Optimal Link Layer Model for Multi-hop MIMO Networks," *Proc. of Infocom,* pp. 1916–1925, 2011.

[30] R. Srinivasan, D. Blough, and P. Santi, "Optimal One-Shot Stream Scheduling for MIMO Links in a Single Collision Domain", *Proc. of SECON,* pp. 700–708, 2009.

[31] R. Srinivasan, D. Blough, L.M. Cortes-Peña, and P. Santi, "Maximizing Throughput in MIMO Networks with Variable Stream Rates", *Proc. of European Wireless,* pp. 551–559, 2010.

[32] K. Sundaresan and R. Sivakumar, "A Unified MAC Framework for Ad-hoc Networks with Smart Antennas," *IEEE/ACM Transactions on Networking,* Vol. 15, pp. 546–559, June 2007.

[33] K. Sundaresan and R. Sivakumar, "Routing in Ad-hoc Networks with MIMO Links: Optimization Considerations and Protocols," *Computer Networks,* Vol. 52, pp. 2623–2644, Oct. 2008.

[34] K. Sundaresan, R. Sivakumar, M. A. Ingram, and T-Y. Chang, "Medium Access Control in Ad-hoc Networks with MIMO Links: Optimization Considerations and Algorithms," *IEEE Trans. on Mobile Computing,* Vol. 3, pp. 350–365, Oct. 2004.

[35] K. Sundaresan, W. Wang, and S. Eidenbenz, "Algorithmic Aspects of Communication in Ad-hoc Networks with Smart Antennas," *Proc. of MobiHoc,* 2007.

[36] S. Ye and R. Blum, "Optimized signaling for MIMO interference systems with feedback," *IEEE Transactions on Signal Processing,* vol. 51, no. 11, pp. 2839–2848, Nov. 2003.

[37] J. Wang, et al., "Joint Bandwidth Allocation, Element Assignment and Scheduling for Wireless Mesh Networks with MIMO Links," *Computer Communic.,* Vol. 31, pp. 1372–1384, 2008.

[38] C. Yetis, T. Gou, S. Jafar, and A. Kayran, "On Feasibility of Interference Alignment in MIMO Interference Networks," *IEEE Trans. on Signal Processing,* Vol. 58, pp. 4771–4782, Sep. 2010.

[39] M. Zorzi, et al., "Cross-Layer Issues in MAC Protocol Design

for MIMO Ad Hoc Networks," *IEEE Wireless Communications,* pp. 62–76, Aug. 2006.

Appendix

Analysis of Procedure Proc 1:

Since the input graph G is modified during the procedure, let us assume there is a copy of the input graph, which we refer to as G_{copy}, that is kept unmodified as the procedure executes.

Step 1: At the start of the procedure, every subgraph of G has $ED \geq 0$. During this step, all edges of degree-one and degree-two vertices in V are marked.

Step 2: V_n contains only vertices in V with degree > 2.

Step 3: If during the procedure, we determine $DO(v, v')$ to be TRUE, that means that *every* subgraph containing v but not v' has $ED > 0$. The procedure is allowed to direct vertex v towards v' *if and only if* $DO(v, v') =$ TRUE. Consequently, after the procedure directs v outwards to v', *every* subgraph containing v but not v' will have $ED \geq 0$. Moreover, all subgraphs containing v' as well as all subgraphs with neither v nor v' will have an unchanged ED value.

On the other hand, if $DO(v, v') = FALSE$ and $DO(v', v) = TRUE$ then the procedure directs v' outwards to v. The procedure is allowed to direct vertex v' towards v *if and only if* $DO(v', v) =$ TRUE. By the same argument as above, after doing this, *every* subgraph containing v' but not v will have $ED \geq 0$ and other subgraphs will not be affected.

Therefore, if the $DO(v, v_i) = TRUE$ condition in Step 3c is entered or the else condition is entered with $DO(v_i, v) = TRUE$, then all subgraphs will maintain the property that $ED \geq 0$.

Since the procedure can progress by marking an edge *only* when at least one of $DO(v, v')$ and $DO(v', v)$ is TRUE (the cases adressed above), *every* subgraph of G will have $ED \geq 0$, making this an invariant. This is true at every stage of the procedure.

Next, we show that at no point during the procedure, can both $DO(v, v')$ *and* $DO(v', v)$ simultaneously evaluate to FALSE. This is stated in Lemma L1.

Lemma L1: *$DO(v, v') \vee DO(v', v)$ is always equal to TRUE at every iteration of Step 3c.*

Proof of lemma:

If we were to have $DO(v, v') = FALSE$ and $DO(v', v) = FALSE$, that would imply that v is contained in some subgraph $G_1 = (V_1, E_1)$ such that it does not contain v' and has $ED = 0$. Also, v' is contained in some subgraph $G_1^{'} = (V_1^{'}, E_1^{'})$ such that it does not contain v and has $ED = 0$. And since an edge exists between v and v', the subgraph induced in G by $V_1 + V_1'$ has $ED = -1$. Since *every* subgraph of the input graph G_{copy} has $ED \geq 0$, this means that at some previous point in the procedure, the subgraph induced in G by $V_1 + V_1'$ had $ED \geq 0$. Consider the point in the procedure when the value of this ED was exactly equal to zero. Suppose without loss of generality that G_1 had an ED value equal to 1 and G_1' had an ED value equal to 0 at this point (this would make the ED of the subgraph induced in G by $V_1 + V_1'$ equal to zero). Now, in order for the ED value of the subgraph induced in G by $V_1 + V_1'$ to become -1 from 0, there must have been some vertex v_a from $G_1 = (V_1, E_1)$ that was directed outwards by the procedure to a vertex v_b where

v_b is *not* contained in V_1 [6] and also *not* contained in $V_1^{'}$ [7] If v_a was directed outwards to such a v_b, that would imply that the procedure was applied incorrectly since clearly there is a subgraph containing v_a, namely the subgraph induced in G by $V_1 + V_1'$ which has an ED value exactly equal to zero, meaning that we are not allowed to direct v_a to v_b.

This contradicts the fact that the procedure always maintains a positive value of ED for *every* subgraph of G. Therefore, this situation where $DO(v, v')$ is FALSE *and* $DO(v', v)$ is FALSE can *never* occur so long as the input graph G_{copy} satisfies D_1.

End of Lemma L1 proof.

Note that the final if statement in Step 3c does not cause any ED to become less than zero, because if v_i satisfies the condition and its last two edges are marked, this can only increase ED for subgraphs containing v_i.

Since $DO(v, v')$ is true or $DO(v', v)$ is true at all times and $ED \geq 0$ for every subgraph of G in either of these cases and the final if condition does not cause the ED condition to be violated, we have the following invariant:

Invariant 1 - $ED \geq 0$ for every subgraph of G at every iteration of Step 3, including after the final iteration.

Next, we prove another invariant of the procedure. If the original graph G_{copy} satisfies D_1, the procedure will never decrement the value of n_v below zero for all $v \in V$. Therefore, $n_v \geq 0$ is also an invariant. This follows from Lemma L1 and Lemma L2, stated next.

Since we started out by having a graph G_{copy} with *every* subgraph having $ED \geq 0$, this means that at some previous point in the procedure, the subgraph induced in G by $V_1 + V_2$ had $ED \geq 0$. Consider the point in the procedure when the value of this ED was exactly equal to zero. Suppose, without loss of generality that G_1 had an ED equal to one and G_2 had an ED equal to zero at this point (this would make the ED of the subgraph induced in G by $V_1 + V_2$ equal to zero). Now consider some vertex v_2 in $G_1 = (V_1, E_1)$. In order for the ED value of the subgraph induced in G by $V_1 + V_2$ to become -1 from 0, the vertex v_2 would have to be directed outwards by the procedure to a vertex v_a where v_a is *not* contained in G_1 [8] and also *not* contained in $G_1^{'}$ [9] If v_2 was directed outwards to such a v_a, that would imply that the procedure was applied incorrectly since clearly there is a subgraph containing V_2, namely the subgraph induced in G by $V_1 + V_2$ which has an ED value exactly equal to zero, meaning that we are not allowed to direct v_2 to v_a.

[6]if v_b was in V_1, then the ED value of G_1 would remain one and the ED value of the subgraph induced in G by $V_1 + V_1'$ would still remain zero.

[7]if v_b was in $V_1^{'}$, that would violate the assumption that at this point, the ED value of the subgraph induced in G by $V_1 + V_1'$ is zero (because, since we have said that ED of G_1 is one and ED of $G_1^{'}$ is zero, the fact that an edge between v and v' exists and the fact that an edge between v_a and v_b exists would make the ED value of this subgraph equal to -1).

[8]if v_a was in G_1, then the ED value of G_1 would remain one and the ED value of the subgraph induced in G by $V_1 + V_2$ would still remain zero.

[9]if v_a was in $G_1^{'}$, that would violate the assumption that at this point, the ED value of the subgraph induced in G by $V_1 + V_2$ is zero (because, since we have said that ED of G_1 is one, ED of $G_1^{'}$ is zero, the fact that an edge between v and v' exists nad the fact that an edge between v_2 and v_a exists would make the ED value of this subgraph equal to -1)

Lemma L2: *At every stage of the procedure, a vertex $v \in V_n$ can have at most n_v neighbors that are contained in subgraphs that exclude v and with ED value equal to zero. Moreover, if v has exactly n_v such neighbors, then it is part of an average-degree-four-subgraph in the input graph G_{copy}.*

Proof of lemma:

Suppose first that $n_v = 0$. Consider a subgraph $G_1 = (V_1, E_1)$ that contains a neighbor of v, but not v itself. Let the ED value of G_1 be ED_1. Therefore, the subgraph induced by $V_1 + v$ in G will have $ED = ED_1 - 1$. Since by Lemma L1, the ED value of *every* subgraph of G is always greater than equal to zero, we must have $ED_1 \geq 1$. Therefore, v has n_v (= 0) neighbors that are contained in subgraphs that do not contain v and with ED=0.

Similarly, for $n_v > 0$, suppose that v has p such neighbors, contained in subgraphs $G_1 = (V_1, E_1), ...G_p = (V_p, E_p)$. The subgraph induced by $V_1 + .. + V_p + v$ in G will have $ED = ED_1 + ... + ED_p - p + n_v$. Since by Lemma L1, ED of *every* subgraph of G is greater than equal to zero, we must have $p \leq n_v$.

End of Lemma L2 proof.

Note 1 - *Lemmas L1 and L2 preclude the following from occurring at any stage of the procedure: $DO(v, v_i) = FALSE$ and $n_{v_i} = 0$. This is because we must have: 1) $DO(v, v_i) = TRUE$ and/or $DO(v, v_i) = TRUE$ by Lemma L1, and 2) v_i has $n_{v_i} = 0$ neighbors in subgraphs with $ED = 0$ by Lemma L2. This means that if $n_{v_i} = 0$, then $DO(v, v_i) = TRUE$ will always hold and v can be marked outwards to v_i.*

This establishes the following invariant.

Invariant 2 - $n_v \geq 0$ for every $v \in V$ at every stage of the procedure.

Step 4: When Step 3 terminates, every vertex $v \in V_n$ has n_v equal to either one or zero. This follows because, at Step 3c, vertex v with $n_v = 2$ will have at least one of its incident edges marked as outgoing, and hence its n_v value will be decremented by at least one.

During Step 4, every vertex $v \in V_n$ with $n_v = 0$ is fully marked. This is done by marking every incident edge towards v. After this, all marked vertices and edges are removed. Therefore, the remaining vertices v all have $n_v = 1$.

Step 4 is justified by the observation that we made in Note 1 above based on Lemmas L1 and L2, that the following can never occur at any stage of the procedure: $DO(v, v_i) = FALSE$ and $n_{vi} = 0$. In other words, whenever $n_v = 0$, $DO(v', v)$ is always TRUE where v' is a neighbor of v. Therefore, we are allowed to mark all edges incident to v as incoming.

Step 5: ***Lemma L3:*** *If the input graph G satisfies Property D_1 then every connected component of the resulting graph $G' = (V', E')$ at the end of Step 4 will be a tree or a simple cycle.*

To see this, note that the invariant $ED \geq 0$ holds true for every subgraph of G'. Moreover, every vertex v in V' has $n_v = 1$. Let $G_i^{'} = (V_i^{'}, E_i^{'})$ be the i^{th} connected component of G. We have $ED(G_i^{'}) = \sum_{v \in V_i^{'}} n_v - |E_i^{'}| = |V_i^{'}| - |E_i^{'}|$. Since by Corollary C1, $ED(G_i^{'}) \geq 0$, we have $|V_i^{'}| \geq |E_i^{'}|$, making every connected component of G' a tree or a simple cycle.

Step 6: When we reach Step 6, every vertex $v \in V'$ has $n_v = 1$ and every connected component is a tree or a simple cycle. In Section V-D, we proved that such a graph can be directed such that every vertex has at most one outgoing edge and we provided a procedure to do so. By applying this procedure to the remaining graph G', the result is that every edge of the original graph G is directed such that every vertex has at most two outgoing edges.