

Decentralized Constraint Satisfaction

K. R. Duffy¹, C. Bordenave² and D. J. Leith¹.

1: Hamilton Institute, National University of Ireland Maynooth

2: Department of Mathematics, University of Toulouse.

Abstract—We show that several important resource allocation problems in wireless networks fit within the common framework of Constraint Satisfaction Problems (CSPs). Inspired by the requirements of these applications, where variables are located at distinct network devices that may not be able to communicate but may interfere, we define natural criteria that a CSP solver must possess in order to be practical. We term these algorithms decentralized CSP solvers. The best known CSP solvers were designed for centralized problems and do not meet these criteria. We introduce a stochastic decentralized CSP solver, proving that it will find a solution in almost surely finite time, should one exist, and also showing it has many practically desirable properties. We benchmark the algorithm’s performance on a well-studied class of CSPs, random k-SAT, illustrating that the time the algorithm takes to find a satisfying assignment is competitive with stochastic centralized solvers on problems with order a thousand variables despite its decentralized nature. We demonstrate the solver’s practical utility for the problems that motivated its introduction by using it to find a non-interfering channel allocation for a network formed from data from downtown Manhattan.

I. INTRODUCTION

A Constraint Satisfaction Problem (CSP) consists of N variables, $\vec{x} := (x_1, \dots, x_N)$, and M clauses, i.e. $\{0, 1\}$ -valued functions, $(\Phi_1(\vec{x}), \dots, \Phi_M(\vec{x}))$. An assignment \vec{x} is a solution if all clauses simultaneously evaluate to 1. In the context of wireless networks, we show that CSPs provide a unifying framework that encompasses many important resource allocation tasks. Examples include: allocation of radio channels so that transmissions by neighbouring WLANs (Wireless Local Area Networks) or mobile phone cells do not interfere; the selection of packets to be XORed on each link in network coding; and finding a non-colliding schedule of time-slots in a WLAN. Unlike in traditional CSPs, however, in network applications each constrained variable x_i is typically co-located with a physically distinct device such as an access point, base-station or a link. The resulting communication constraints impose severe restrictions on the nature of the algorithm that can be used for solving the CSP. These restrictions are violated by existing CSP solvers, leading us to define a new class of algorithms that we term *decentralized CSP solvers*. We constructively establish the existence of decentralized solvers by introducing a family of randomized algorithms belonging to this class.

Roughly speaking, decentralized CSP solvers must be capable of finding a satisfying assignment, \vec{x} , while updating each variable x_i based solely on knowledge of whether all of the clauses in which x_i participates are satisfied or at least one

clause is unsatisfied. The following example gives a concrete illustration of the origins of these constraints.

A. Motivating example: channel allocation

802.11 WLANs are ubiquitous, but their deployments are unstructured and they operate in an unlicensed frequency band. Within this band, an 802.11 WLAN can select one of several channels, typically 11, to operate on. In the vast majority of WiFi routers channel selection is, at present, based on operator selection without any quantitative instruction from the router. Self-selection of this channel is the task we consider here. Co-ordinated selection is hampered by the fact that the interference range of a typical 802.11 device is substantially larger than its communication range. Consequently, WLANs can interfere but may be unable to decode each others messages (illustrated schematically in Fig. 1). Discovery by a WLAN of the existence of interfering WLANs via its wired back-haul may be prevented by firewalls and, in any case, these WLANs may not know their physical location. These practical restrictions mandate a decentralized channel-selection algorithm.

We can identify the task of each WLAN self-selecting a non-interfering channel with a CSP with communication constraints imposed by the network topology. Let x_i be the channel selected by WLAN $i \in \{1, \dots, N\}$ and define $M = N(N - 1)/2$ clauses, one for each pair of WLANs, that evaluates to one if the WLANs are on non-interfering channels or are out of interference range and is zero otherwise. Due to the unstructured nature of the deployment, the lack of shared administrative control and communication constraints, a WLAN cannot rely on knowing the number or the identity of interfering WLANs or the channels that they currently have selected. A practical CSP solver for this task can only rely on each WLAN being able to measure whether: (i) all of the neighbouring WLANs have selected a different channel from it; or (ii) if one or more neighbours have selected the same channel.

In certain settings, a limited amount of communication of variable values between the stations hosting them may be feasible. For example, it may be possible for a stations hosting a link or for a WLAN to overhear beacons or traffic from a subset of its interferers, or for an access point to communicate with other, nearby access points. Algorithms for channel selection that are proposed in the literature assume the existence of end-to-end communication for centralized solutions, for message passing in decentralized or gossiping solutions, or to co-ordinate global restart in simulated annealing proposals e.g. [1], [2], [3], [4] and references therein. This

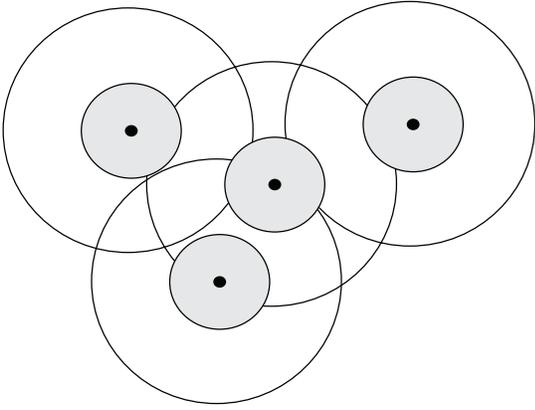


Fig. 1. An illustration of overlapping interference regions in neighbouring WLANs. Black dot indicates location of a wireless access point, shaded area indicates region within which communication with access point can take place and outer circles indicate regions within which transmissions interfere.

information, however, tends to be opportunistic in nature and cannot be relied on *a priori*. We leave how best to exploit such partial information to future work, focusing here on the most challenging, fundamental cases where fully decentralized algorithms are required.

B. Contribution

The primary contributions of this paper are fourfold. Firstly, we show that CSPs provide a unifying framework that encompasses several important problems in wireless networks, including channel selection, inter-session network coding and decentralized scheduling in a WLAN. Each of these problems is generally thought of as being distinct and requiring different solution approaches. We show that they are fundamentally related as CSPs so that the decentralized CSP solver introduced in this paper can be used for all of them. Secondly, we define a new class of CSP algorithms which we term decentralized CSP solvers and, extending our earlier approach to graph colouring [5][6], introduce a novel stochastic decentralized CSP solver proving that it will find a solution in almost surely finite time, should one exist, and showing it has many practically desirable properties. Thirdly, we benchmark the algorithm's performance on a well-studied class of application-agnostic CSPs, random k-SAT. For problems with up to a thousand variables, which are large for our motivating examples, we find that the time the algorithm takes to find a satisfying assignment is close to that of WalkSAT [7], a well regarded, efficient centralized CSP solver, while also possessing desirable features of Survey Propagation [8]. That is, despite its decentralized nature, the algorithm is fast. Fourthly, we demonstrate the solver's practical utility in a complex, wireless resource allocation case study.

The rest of the paper is organized as follows. In Section II we define decentralized CSP solvers and show that well-known algorithms for solving CSPs fail to meet these criteria. In Section III we introduce a decentralized CSP solver. We prove that it finds a solution in finite time with probability one whenever a feasible solution exists, obtaining an upper bound on the algorithm's convergence rate in the process. In Section

IV we consider its speed at identifying solutions of instances of random k-SAT. In Section V we demonstrate the algorithm's utility in solving the problems that motivated its introduction. Section VI contains discussion and closing remarks.

II. DECENTRALIZED CSP SOLVERS

We begin by formalizing the criteria that CSP solvers must possess to be of use for problems with communication constraints such as those outlined above. We call algorithms that meet these criteria decentralized CSP solvers.

Definition 1 (CSP): A CSP with N variables, $\{x_1, \dots, x_N\}$, and M clauses is defined as follows. The variables each take values in a finite set $\mathcal{D} = \{1, \dots, D\}$ and $\vec{x} := (x_1, \dots, x_N) \in \mathcal{D}^N$. Each clause $m \in \mathcal{M} = \{1, \dots, M\}$ is defined by a function $\Phi_m : \mathcal{D}^N \mapsto \{0, 1\}$ where for an assignment of variables, \vec{x} , $\Phi_m(\vec{x}) = 1$ if clause m is satisfied and $\Phi_m(\vec{x}) = 0$ if clause m is not satisfied. An assignment \vec{x} is a solution to the CSP if all clauses are simultaneously satisfied. That is,

$$\vec{x} \text{ is a satisfying assignment iff } \min_{m \in \mathcal{M}} \Phi_m(\vec{x}) = 1. \quad (1)$$

This encompasses all of our examples of interest as well as k-SAT, the most well-studied general class of CSPs.

A CSP solver is an algorithm that can find a satisfying variable assignment for any solvable CSP.

Definition 2 (CSP Solver): A CSP solver takes a CSP as input and determines a sequence $\{\vec{x}(t)\}$ such that for any CSP that has satisfying assignments:

- (D1) for all t sufficiently large $\vec{x}(t) = \vec{x}$ for some satisfying assignment \vec{x} ;
- (D2) if t' is the first entry in the sequence $\{\vec{x}(t)\}$ such that $\vec{x}(t')$ is a satisfying assignment, then $\vec{x}(t) = \vec{x}(t')$ for all $t > t'$.

To provide criteria that classify CSP solvers as being decentralized, we begin with the following definition.

Definition 3 (Clause participation): We say that variable x_i participates in clause $\Phi_m(\vec{x})$ if the value of x_i can influence the clause's satisfaction for at least one assignment of the rest of the variables in \vec{x} .

For each variable x_i , let \mathcal{M}_i denote the set of clauses in which it participates:

$$\mathcal{M}_i = \bigcup_{j \neq i} \bigcup_{x_j \in \mathcal{D}} \left\{ m : \min_{x_i \in \mathcal{D}} \Phi_m(x_i, \{x_j\}) = 0 \right. \\ \left. \text{and } \max_{x_i \in \mathcal{D}} \Phi_m(x_i, \{x_j\}) = 1 \right\}.$$

Re-writing the left hand side of eq. (1) in a way that focuses on the satisfaction of each variable we have:

$$\vec{x} \text{ is a satisfying assignment iff } \min_i \min_{m \in \mathcal{M}_i} \Phi_m(\vec{x}) = 1. \quad (2)$$

A decentralized CSP algorithm can be thought of as having intelligence co-located with each variable. The intelligence at variable x_i can only determine whether all of the clauses that x_i participates in are satisfied or that at least one clause is unsatisfied and must make update decisions locally based solely on this knowledge.

Definition 4 (Decentralized CSP Solver): A decentralized CSP solver is a CSP solver that for each variable x_i , $i \in \{1, \dots, N\}$, must select its next value based only on an evaluation of

$$\min_{m \in \mathcal{M}_i} \Phi_m(\vec{x}). \quad (3)$$

That is, for each variable x_i the solver must make a decision on the next value of x_i based solely on knowing whether that variable is currently satisfied or not without explicitly knowing:

- (D3) the assignments of other variables, x_j for any $j \neq i$;
- (D4) the set of clauses that any variable, including itself, participates in, \mathcal{M}_j for $j \in \{1, \dots, N\}$;
- (D5) the functions that define clauses, Φ_m for $m \in \{1, \dots, M\}$.

Note that the properties (D1) (D2) of any CSP solver mean that a decentralized CSP solver must, without explicit communication, settle upon a satisfying assignment the first time one is found. Communication to co-ordinate global stopping or restarting of the solver would be contrary to the nature of the natural constraints of these problems and so is forbidden.

The stochastic algorithm we define in Section III shall provably satisfy the properties of a decentralized CSP solver with probability one. Moreover, it will also have the desirable property that it will automatically restart its search, without communication, upon any change to the clauses of a CSP that makes the current variable assignment no longer satisfying. This is significant in practical applications as, for example, the arrival of new transmitters in a wireless network would induce a change the associated CSP.

A. Formulating Wireless Network Tasks As CSPs

Before proceeding, we demonstrate that several important resource allocation tasks in wireless networks fall within this CSP framework, which, therefore, provides a unifying framework for analyzing these tasks.

1) *Graph colouring.* As briefly described in Section I-A, the channel assignment problem corresponds to decentralized graph colouring. In the simplest model, the network is described by an undirected interference graph $G = (V, E)$ with vertices V and edges E . Each vertex represents a WLAN and an edge exists between two WLANs if they interfere with each other when on the same channel. Define $N = |V|$ and $M = |E|$ and let \mathcal{D} denote the set of available colours (channels) and let x_i be a variable with value equal to the colour selected by an intelligence co-located with each vertex $i \in V$. Each clause $m \in \{1, \dots, |E|\}$, which is an enumeration of the elements of E , can be identified with an edge $(i, j) \in E$. We define

$$\Phi_m(\vec{x}) = \Phi_m(x_i, x_j) = \begin{cases} 1 & \text{if } x_i \neq x_j \\ 0 & \text{otherwise.} \end{cases}$$

The participation set, \mathcal{M}_i , of a variable x_i consists of all clauses where the vertex i is one end of the associated edge,

$$\mathcal{M}_i = \{m \equiv (i, j) : (i, j) \in E\}.$$

A variable assignment satisfies this CSP if and only if it is also a proper colouring for the graph G . Proper colourings

correspond to channel assignments in which no two neighbouring WLANs in the interference graph have selected the same channel. To find a proper colouring, when one exists, a decentralized CSP solver requires only that an intelligence co-located with each vertex i can measure whether: (i) all of its neighbours are using a different colour from vertex i ; or (ii) at least one neighbour has selected the same colour as vertex i . This information is sufficient to evaluate eq. (3) and, in particular, the intelligence at each vertex does not need to know: (D3) the colour selected by any other vertex; (D4) its set of neighbours; or (D5) the exact nature of the clauses Φ_m , $m \in \{1, \dots, M\}$.

2) *Channel assignment with channel-dependent interference.* Whether or not transmitters interfere in the channel assignment task may depend on the radio channel selected. This can arise due to frequency dependent radio propagation or to channel dependent spectral masks. That is, for regulatory reasons, different spectral masks are typically used when transmitting on channels at the edge and on channels in the middle of a radio band. This problem then has a collection of conflict graphs, one for each possible radio channel. This version of the channel assignment task can also be formulated as a CSP even though it is no longer a graph colouring problem. Let $G(c) = (V, E(c))$, $c \in \mathcal{D}$ be a set of undirected graphs with the same vertices V but possibly differing edge sets $E(c)$. Again each vertex represents a WLAN, but the interference graph is channel-dependent. The Graph $G(c)$ is associated with radio channel c and an edge exists in $E(c)$ if two WLANs interfere while on channel c . Let x_i be a variable with value equal to the channel selected by each vertex $i \in V$. Each clause $m \in \{1, \dots, |E(1)| + \dots + |E(\mathcal{D})|\}$, which is an enumeration of all edges in all graphs, can be identified with a $c \in \mathcal{D}$ and an edge $(i, j) \in E(c)$ and is defined by

$$\Phi_m(\vec{x}) = \Phi_m(x_i, x_j) = \begin{cases} 0 & \text{if } x_i = x_j = c \\ 1 & \text{otherwise.} \end{cases}$$

The collection of clauses that variable x_i , associated with vertex i , participates in can be identified as

$$\mathcal{M}_i = \bigcup_{c \in \mathcal{D}} \{m \equiv (i, j, c) : (i, j) \in E(c)\}.$$

With x_i being the current channel selection of vertex i , for all $c \neq x_i$ the clauses associated with $m \equiv (i, j, c)$ are automatically satisfied. Thus to evaluate eq. (3) it is sufficient for the station to measure if no neighbour coincides with its current channel selection or if at least one does. To find an interference-free channel assignment, a decentralized CSP solver requires that each vertex i can measure whether: (i) all of its neighbours on the currently selected channel are using a different channel from itself; or (ii) at least one neighbour has selected the same channel.

3) *Inter-session network coding.* Network coding has been the subject of considerable interest in recent years as it offers the potential for significant increases in network capacity [9], [10], [11]. In network coding, network elements combine packets together before transmission rather than forwarding them unmodified. The combined packets can be from indi-

vidual flows, known as intra-flow coding, or across multiple flows, inter-flow coding.

While intra-flow coding within multicast flows has been well studied, inter-flow coding between unicast flows has received less attention, yet is, perhaps, of more immediate relevance to current Internet traffic. Inter-flow coding is known to be challenging [12]. The task of a network finding a feasible linear network code in a distributed fashion, but with some global sharing of calculations, is investigated in [13] through the use of a genetic algorithm. This task can be formulated as a CSP.

Let $G = (V, E)$ denote a directed acyclic multi-graph representing the network with vertices V and edges E . Each vertex represents either a source of an information flow, a destination or router in the network. Edges represent physical connections between these elements. Time is slotted and each edge can transmit a single packet per slot. We allow multiple edges between vertex pairs in order to accommodate higher rate links. It is the goal of each link beyond the source to determine a linear combination of its incoming packets to forward so that ultimately all flows get to their destinations. As the combination of packets coded by each link will take care of routing, each flow $p \in P$ is defined by its source, $\sigma(p) \in V$, and its destination $\delta(p) \in V$. Each flow is assumed to be unit rate with higher rate flows accommodated by flows with the same source and destination vertices. We will treat the collection of source, $\bigcup_{p \in P} \{\sigma(p)\}$, and destination vertices, $\bigcup_{p \in P} \{\delta(p)\}$, as special, solely having one incoming and one outgoing edge per flow respectively.

For each edge $i \in E$, define $s(i)$, $s : E \mapsto V$, to be its source vertex and $t(i)$, $t : E \mapsto V$, to be its target vertex. For each vertex $v \in V$ we define the set of incoming edges $I_v := \{i \in E : t(i) = v\}$ and the set of outgoing edges $O_v := \{i \in E : s(i) = v\}$. For each edge not corresponding to source, $i \in E \setminus \bigcup_{p \in P} O_{\sigma(p)}$, we associate a variable $x_i \in \mathcal{D} := \{1, \dots, 2^{|P|}\}$, which is an enumeration of the power set of the set of flows P . We define a bijective map $\psi : \mathcal{D} \mapsto \{0, 1\}^{|P|}$ that determines a vector whose positive entries correspond to flows to be coded.

In total, we have $|E \setminus \bigcup_p O_{\sigma(p)}|$ clauses. One for every link that isn't an outgoing link from a source vertex. At each time slot, each link that doesn't correspond to a source or final destination link, $i \in E \setminus \bigcup_p (O_{\sigma(p)} \cup I_{\delta(p)})$, wishes to forward a packet consisting of XORed packets (corresponding to addition in the Galois field of two elements) from the flows indicated by $\psi(x_i)$. Each final link $i \in \bigcup_p I_{\delta(p)}$ wishes to forward packets from p , which we indicate by the vector γ_p with a 1 at the location p and zeros elsewhere. A link and its immediate upstream neighbouring links will be dissatisfied if it cannot do so.

More formally, for each edge not corresponding to the outgoing link from a source or incoming link to a destination, $i \in E \setminus \bigcup_{p \in P} (O_{\sigma(p)} \cup I_{\delta(p)})$, we have a clause $m \equiv i$, $\Phi_m(\vec{x})$, defined by

$$\Phi_m(x_i, x_{j:j \in I_{s(i)}}) = \begin{cases} 1 & \text{if } \exists \theta \text{ s.t. } \Psi(I_{s(i)})\theta^T = \psi(x_i)^T \\ 0 & \text{otherwise.} \end{cases}$$

where θ is a binary vector and $\Psi(I_{s(i)})$ is the rectangular matrix consisting of $\psi(x_j)^T$ for $j \in I_{s(i)} \setminus \bigcup_{p \in P} O_{\sigma(p)}$ and γ_p^T if $i \in O_{\sigma(p)}$ for some p .

To complete the CSP, we need to ensure that packets from flow p can be decoded at their destination $\delta(p)$. For $i \in \bigcup_{p \in P} I_{\delta(p)}$, then with $\delta(p') = t(i)$ we introduce a clause

$$\Phi_m(x_{j:j \in I_{s(i)}}) = \begin{cases} 1 & \text{if } \exists \theta \text{ s.t. } \Psi(I_{s(i)})\theta^T = \gamma_{p'}^T \\ 0 & \text{otherwise.} \end{cases}$$

where again θ is a binary vector and Ψ is the rectangular matrix defined above. The set of clauses a link variable x_i , $i \in E \setminus (\bigcup_{p \in P} O_{\sigma(p)} \cup I_{\delta(p)})$, participates in is

$$\mathcal{M}_i = \{m \equiv j : j = i \text{ or } i \in I_{s(j)}\}$$

Any variable assignment then satisfies this CSP if and only if it is a realizable network code satisfying all flow demands.

Hence, to find a proper assignment a decentralized CSP solver requires only that each link i can determine whether: (i) its own coding is realizable and the coding of its immediately down-stream links are realizable; or (ii) if at least one of these is not realizable. This is sufficient to evaluate eq. (3). Each link does not need to explicitly know: (D3) the code selected by any other link; (D4) the network topology; (D5) any details of how code realizability is determined at any link.

4) *Decentralized transmission scheduling.* In an Ethernet or WLAN it is necessary to schedule transmissions by stations. This might be achieved in a centralized manner using TDMA, but it can also be formulated as a decentralized problem. The classical CSMA/CA approach to decentralized scheduling never settles to a single schedule and some comes at a cost of the possibility of continual collisions. Recently, there has been interest in decentralized approaches for finding collision-free schedules, see [14], [15]. This task can be formulated as a CSP as follows. Let V denote the set of transmitters in the WLAN, T denote the set of available time slots and define $N = |V|$ and $M = N(N - 1)/2$. Let x_i be a variable with value equal to the transmission slot selected by transmitter $i \in V$. Define a clause $\Phi_m(\vec{x})$ associated with each pair of transmitters $m \equiv (i, j)$ such that

$$\Phi_m(x_i, x_j) = \begin{cases} 1 & \text{if } x_i \neq x_j \\ 0 & \text{otherwise.} \end{cases}$$

The participation sets are

$$\mathcal{M}_i = \bigcup_{j \neq i} \{m \equiv (i, j)\}.$$

Any variable assignment satisfies this CSP if and only if it is also a collision-free time-slot schedule. To find a collision-free schedule, when one exists, a decentralized CSP solver requires only that each transmitter i can measure whether: (i) all of its neighbours are using a different time-slot from transmitter i ; or (ii) at least one transmitter in the WLAN has selected the same time-slot as transmitter i . Again, this is all that is needed to evaluate eq. (3). Each transmitter does not need to know: (D3) the time-slot selected by any other transmitter; (D4) the set of transmitters; (D5) the clauses.

B. Related work - existing algorithms are not decentralized

The literature on general purpose CSP solvers is vast, typically focusing on solving k-SAT problems in conjunctive normal form, but they can be broadly classified into those based on: (i) the Davis-Putnam-Logemann-Loveland (DPLL) algorithm [16], [17]; (ii) Survey Propagation [8]; and (iii) on Stochastic Local Search (SLS) [7]. Each of these approaches has experienced substantial development and has its own merits, but none were motivated by problems where variables have a geographical sense of locality.

Algorithms developed from the DPLL approach have proved to be the quickest at SAT-Race and SAT Competition in recent years, e.g. ManySAT [18]. The DPLL approach ultimately guarantees a complete search of the solution space and so meets the (D1) and (D2) criteria. They are, however, based on a branching rule methodology, e.g. [19], that assumes the existence of a centralized intelligence that employs a backtracking search. The implicit assumptions of the information available to this intelligence breaks the conditions (D3) (D4) (D5) and so they are not decentralized CSP solvers.

Survey propagation, a development of belief propagation [20] from trees to general graphs, has proved effective in graphs that do not contain small loops [21]. For a given CSP, the fundamental structure of study is a called a factor graph. In order to generate this, it is necessary to know what clauses each variable participates in and the nature of each of the clauses, breaking the (D4) and (D5) criteria and so these are not decentralized CSP solvers.

SLS algorithms also depend fundamentally upon the exchange of information, mostly in an explicit manner breaking the (D3) condition by basing update decisions on relative rankings of the constraint variables but also in a more subtle fashion. To see this implicit requirement, consider the following algorithm for binary valued variables originally proposed by Papadimitriou [22] and developed further by Schöning [23]. Pick a random assignment of values for the constraint variables. Repeat the following: from all of the unsatisfied clauses, pick one uniformly at random, select one of the variables participating in that clause and negate its value, breaking the (D4) and (D5) conditions. The algorithm halts when all clauses are satisfied or a specified time limit expires. Although simple, this forms the basic building block for all SLS algorithms, including the well-studied WalkSAT algorithm [7]. It is important that a single unsatisfied clause is selected at each step and that a single variable within the clause is adjusted as it is this that leads to the algorithm behaving as a random walk [23]. Thus, again, solvers in this class are not decentralized.

III. A DECENTRALIZED CSP SOLVER

We now introduce an algorithm that satisfies the decentralized CSP solver criteria.

A. Communication-Free Learning Algorithm

An instance of the following Communication-Free Learning (CFL) algorithm is run in parallel for every variable. For each variable, $i \in \{1, 2, \dots, N\}$ it keeps a probability distribution,

$p_{i,j}$ over $j \in D$ as well as the current variable value x_i . In pseudocode, the CFL algorithm is:

ALGORITHM 1: Communication-Free Learning

- 1: Initialize $p_{i,j} = 1/d, j \in \{1, \dots, D\}$.
- 2: **loop**
- 3: Realize a random variable, selecting $x_i = j$ with probability $p_{i,j}$.
- 4: Evaluate $\min_{m \in \mathcal{M}_i} \Phi_m(\vec{x})$, returning *satisfied* if its value is 1, as this indicates all of variable i 's clauses are satisfied given the present assignment, and *unsatisfied* otherwise.
- 5: Update: If *satisfied*,

$$p_{i,j} = \begin{cases} 1 & \text{if } j = x_i \\ 0 & \text{otherwise.} \end{cases}$$

If *unsatisfied*,

$$p_{i,j} = \begin{cases} (1-b)p_{i,j} + a/(D-1+a/b) & \text{if } j = x_i \\ (1-b)p_{i,j} + b/(D-1+a/b) & \text{otherwise,} \end{cases}$$

where $a, b \in (0, 1]$ are design parameters.

- 6: **end loop**
-

To understand the logic behind CFL, note that for each variable a probability distribution over all possible variable-values is kept. The variable's value is then selected from this distribution. Should all the clauses that the variable participates in be satisfied with its current value, the associated probability distribution is updated to ensure that the variable value remains unchanged. If at least one clause is unsatisfied, then the probability distribution evolves by interpolating between it and a distribution that is uniform on all values apart from on the one that is presently generating dissatisfaction. Thus if all of a variable's clauses were once satisfied, the algorithm retains memory of this into the future. This has the effect that if a collection of variables are content, they can be resistant, but not impervious, to propagation of dissatisfaction from other variables.

CFL possesses two parameters. The parameter b determines how quickly the past is forgotten, while a determines the algorithm's aversion to a variable value should it be found to cause dissatisfaction. Even though an instance of CFL is run for each variable, we shall show that this completely decentralized solution is a CSP solver.

From here on we will assume that the update Step 5 is performed in a synchronized fashion across variables. Without synchronization, the fundamental character of the algorithm doesn't change, but the analysis becomes more involved. Synchronization solely requires that algorithm instances each have access to a shared sense of time and that this can be achieved without information-sharing or other communication between variables, i.e. between algorithm instances. A suitable clock is, for example, available to any Internet connected device via the Network Time Protocol (NTP).

B. CFL is a decentralized CSP solver

By construction, the only information used by the algorithm is $\min_{m \in \mathcal{M}_i} \Phi_i(\vec{x})$ in Step 4 and thus it satisfies the criteria (D3), (D4) and (D5). That is, it only needs to know if all

clauses in which variable i participates are satisfied or if one or more are not. The CFL algorithm also satisfies the (D2) criterion that it sticks with a solution from the first time one is found. To see this, note that the affect of Step 5 is that if a variable experiences success in all clauses that it participates in it continues to select the same value with probability 1. Thus if all variables are simultaneously satisfied in all clauses, i.e. if $\min_i \min_{m \in \mathcal{M}_i} \Phi_i(\vec{x}) = 1$, then the same assignment will be reselected indefinitely with probability 1.

To establish that the CFL algorithm is a decentralized CSP solver all that remains is to show that it meets the (D1) criterion, that if the problem has a solution it will be found, which is dealt with by the the following theorem. It provides an upper bound on the distribution of the number of iterations the algorithm requires to find a solution to any solvable CSP. It's proof exploits the iterated function system structure of the algorithm and can be found in the Appendix.

Theorem 1: For any satisfiable CSP, with probability greater than $1 - \epsilon \in (0, 1)$ the number of iterations for the CFL algorithm to find a satisfying assignment is less than

$$N \exp\left(\frac{N(N-1)}{2} \log(\gamma^{-1})\right) \log(\epsilon^{-1}),$$

$$\text{where } \gamma = \frac{\min(a, b)}{D - 1 + a/b}.$$

For a CSP corresponding to graph coloring, a tighter bound holds and a satisfying assignment will be found with probability greater than $1 - \epsilon$ in a number of iterations less than

$$N \exp(2N \log(\gamma^{-1})) \log(\epsilon^{-1}).$$

Theorem 1 proves that for fixed N the tail of the distribution of the number of iterations until the first identification of a satisfying assignment is bounded above by a geometric distribution and so all of its moments are finite. As Theorem 1 covers any arbitrary CSP that admits a solution, for any given instance these bounds are likely to be loose. They do, however, allow us to conclude the following corollary proving that if a solution exists, the CFL algorithm will almost surely find it.

Corollary 2: For any CSP that admits a satisfying assignment, the CFL algorithm will find a satisfying assignment in almost surely finite time.

Hence the CFL algorithm satisfies all of the criteria (D1) (D2) (D3) (D4) (D5), almost surely, and so is a decentralized CSP solver.

C. Parameterization

Theorem 1 establishes that the CFL algorithm provably identifies satisfying assignments for all values of its two design parameters, a and b . The value of a determines the algorithm's aversion to variable values for which clause failure has been experienced. The value of b impacts on the speed of convergence of the algorithm. Optimal values of a and b depend upon each problem and a performance metric. For fast convergence across a broad range of CSPs with distinct structure, we have found that small values of a and b , corresponding to strong aversion to a dissatisfying variable value and reasonably long memory, work well. Thus for simplicity, we set $a = b$ in all experimental examples. We use $b = 0.2$ for random 3-SAT,

$b = 0.1$ for random 4-SAT and $b = 0.05$ for random 5-SAT. We use $b = 0.1$ for our wireless networks example, a value which we have found to yield good performance across a range of k -SAT problems.

IV. APPLICATION-AGNOSTIC BENCHMARKING

A CSP associated with a specific networking task will possess structure induced by its topology. Before considering practically-motivated applications in Section V, as the CFL algorithm can solve any CSP but its design was subject to restrictions not normally considered it is prudent to investigate its performance as a CSP solver in an application-agnostic setting.

We present simulation data evaluating the performance of the CFL algorithm for random k -SAT, a well-studied class of CSPs. A k -SAT problem is a CSP in which all variables are binary-valued and clauses consist solely of logical disjunctions of no more than k variables. In random k -SAT, an instance of k -SAT is generated by drawing M such clauses uniformly at random [22]. Our investigation surprisingly reveals that the CFL algorithm is competitive with centralized and distributed solvers on problems of reasonable size (order 1000 variables).

Briefly, we first review current knowledge regarding random k -SAT. The behavior of random k -SAT is known to depend strongly on the parameter $r = M/N$ with phase transitions and associated thresholds $r_{-exist,k}$ and $r_{exist,k}$. If $r > r_{-exist,k}$ the constraint problem is unsatisfiable with high probability, while if $r < r_{exist,k}$ a satisfying assignments exist with high probability. Evidently, $r_{exist,k} \leq r_{-exist,k}$ and it is conjectured that $r_{exist,k} = r_{-exist,k}$ [25]. A simple argument gives $r_{-exist,k} \leq 2^k \log 2$ and this upper bound can be refined to obtain the values in the second row of Table I. Also shown in row four of this table are estimated values for $r_{-exist,k}$ derived from statistical physics considerations. These latter estimates are supported by experimental data for $k = 3$, e.g. [30], although there are fewer experimental studies for $k > 3$. It can be seen that the theoretical bound for $r_{-exist,k}$ and the estimated values are in good agreement for $k \geq 5$, and that both approach $2^k \log 2$ for large k . Recent mathematical results have established that $r_{exist,k} \geq 2^k \log 2 - k$ [24] and this lower bound can be tightened to obtain the theoretical values shown in the third row of Table I.

There also exists a threshold $r_{poly,k}$ below which a satisfying assignment can be found in polynomial time with high probability. This threshold has been the subject of much interest and the current best analytic lower bound for $r_{poly,k}$ is also indicated in Table I. Statistical physics considerations have led to the conjecture that $r_{poly,k}$ is equal to the value $r_{1RSB,k}$ at which the one-step replica symmetry breaking (1RSB) instability occurs [31]. Current best estimates for this value are given in Table I. For $k \geq 8$, it has been proven analytically that the set of satisfying assignments is grouped into widely separated clusters, which lends support to this conjecture [32]. For values of $k < 8$ the situation is less clear, with experimental evidence indicating that $r_{poly,k}$ lies above $r_{1RSB,k}$ for $k = 3, 4$ and 5 [33].

We now consider the performance of the CFL algorithm. Fig. 2 gives median measurements of normalized stopping

k	3	4	5	7	10	20
Theoretically derived bounds						
$2^k \log 2$	5.54	11.09	22.18	88.72	709.78	726,817
$r_{\neg exist,k}$ [24]	4.51	10.23	21.33	87.88	708.94	726,817
$r_{exist,k}$ [25], [26]	3.52	7.91	18.79	84.82	704.94	726,809
$2^k \log 2 - k$	2.54	7.09	17.18	81.72	699.78	706,817
$r_{poly,k}$ [27], [26]	3.52	5.54	9.63	33.23	172.65	95,263
Estimates						
$r_{\neg exist,k}$ [28]	4.267	9.93	21.12	87.79	708.91	-
$r_{1RSB,k}$ [28], [29]	4.15	9.38	19.16	62.5	-	-

TABLE I

CURRENT BEST UPPER/LOWER BOUNDS ON THE PHASE TRANSITION THRESHOLDS IN RANDOM K-SAT.

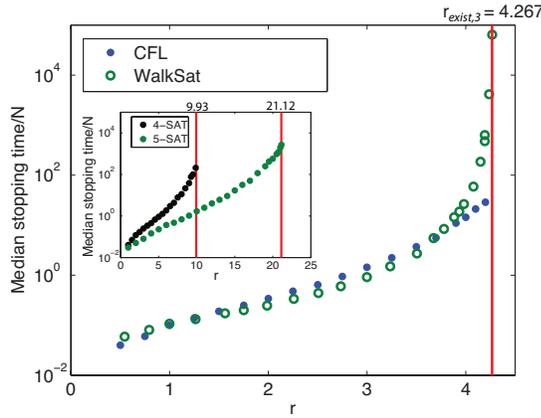


Fig. 2. Median normalized stopping time of CFL algorithm vs r for random 3-SAT and $N = 100$. Each point is derived from at least 1000 runs of the algorithm; algorithm terminated after 10^7 iterations if no satisfying assignment found. For comparison, stopping time measurements for WalkSAT taken from [30, Fig 2a] are also indicated. (Inset) CFL algorithm measurements for random 4-SAT and 5-SAT. Measurements shown are for values of r up to 4.20 (3-SAT), 9.90 (4-SAT), 21.10 (5-SAT); close to the current best estimates for $r_{exist,k}$ and well above $r_{1RSB,k}$.

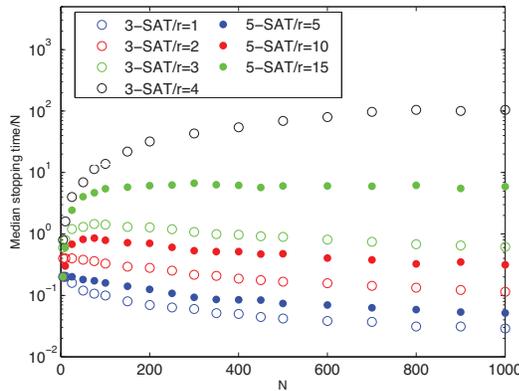


Fig. 3. Normalized stopping time of CFL algorithm vs N and r for 3-SAT and 5-SAT. Each point is derived from at least 1000 runs.

time (stopping time divided by N) of the CFL algorithm vs $r=M/N$ when the number of variables $N = 100$. Fig. 3 gives median measurements of normalized stopping time vs N with r held constant. Data is shown for random k -SAT with $k = 3, 4$ and 5. Observe two striking features from this data. Firstly, from Fig. 2 noting that a log scale is used on the y -axis, the median normalized stopping time increases exponentially with r when N is held constant. Secondly, from Fig. 3 the median normalized stopping time vs N , with r held constant, is upper bounded by a constant. That is, a satisfying assignment is found in a time with median value that increases no more than linearly with N . This linearity holds even when r is close to $r_{exist,k}$ (data is shown in Fig. 3 for 3-SAT with $r = 4$) and is of great practical importance as it implies that with high probability the CFL algorithm finds a satisfying assignment in polynomial time.

Comparing the performance of the CFL algorithm with the popular WalkSAT algorithm, we note that the WalkSAT algorithm also exhibits an exponential-like dependence of stopping time on r and linearity in N [30], [34], [33]. To allow comparison in more detail, median stopping time data taken from [30, Fig 2a] is marked on Fig. 2. For $r < 3.9$, the median stopping time is similar for CFL and WalkSAT. Above this value, however, the stopping time for WalkSAT diverges, increasing super-exponentially in r . This divergence is a feature not only of WalkSAT but also of other local search algorithms *e.g.* ChainSat [33]. It is not exhibited by the Survey Propagation algorithm [30], [8], which has been the subject of considerable interest as it creates the ability to operate close to the $r_{exist,k}$ threshold. Observe that the CFL algorithm also does *not* exhibit divergence as r approaches $r_{exist,k}$. These comparisons are encouraging as they indicate that the CFL algorithm is competitive with some of the most efficient general-purpose k -SAT algorithms currently available. It is also unexpected as information sharing is a key component of both WalkSAT and Survey Propagation, whereas CFL makes local decisions simultaneously with no use whatsoever of information-sharing, raising fundamental questions as to the role of message passing, the relationship between information exchange and algorithm performance and, in particular, what performance cost is imposed by constraining attention to decentralized operation.

V. CASE STUDY: CHANNEL ALLOCATION IN 802.11 WLANS

Having established general properties of CFL and determined its performance on random instances of k-SAT, we return to a problem of the sort that motivated its introduction. We consider the performance of the CFL algorithm in a realistic wireless network case study. From the online database WIGLE [35] we obtained the locations of WiFi wireless Access Points (APs) in an approximately $150m^2$ area at the junction of 5th Avenue and 59th Street in Manhattan¹. This space contains 81 APs utilizing the IEEE 802.11 wireless standard. It can be seen from Figure 4, which plots the mean number of APs lying within distance d of an AP, that within a 15m radius an AP has on average 3 neighbours and within a 30m radius it has on average 10 neighbours. Of the 11 channels available in the 802.11 protocol, only 3 are orthogonal. Thus managing interference in this dense deployment is a challenging task.

Imagine a worst-case scenario where after a power-outage all these APs are switched back on. The aim of each AP is to select its radio channel in such a way as to ensure that it is sufficiently different from nearby WLANs. This can be written as a CSP where we have $N = 81$ APs and N variables x_i corresponding to the channel of AP i , $i = 1, 2, \dots, N$. As per the 802.11 standard [36] and FCC regulations, each AP can select from one of 11 radio channels in the 2.4GHz band and so the x_i , $i = 1, 2, \dots, N$ take values in $\mathcal{D} = \{1, 2, \dots, 11\}$. To avoid excessive interference each AP requires that: (a) no other AP within a 5m distance operates closer than 3 channels away; (b) no AP within a 10m distance operates closer than 2 channels away and (c) no AP with a 30m distance operates on the same channel. We can realize this as a CSP with 3 constraint clauses per AP, giving $3N$ in total, $\Phi_m(\vec{x})$, $m = 1, \dots, 3N$. With $e(i, j)$ denoting the Euclidean distance between APs i and j in metres, for $m = 1, \dots, N$

$$\Phi_m(\vec{x}) = \begin{cases} 1 & \text{if } \min_{j:e(m,j)<5} |x_m - x_j| \geq 3 \\ 0 & \text{otherwise,} \end{cases}$$

for $m = N + 1, \dots, 2N$

$$\Phi_m(\vec{x}) = \begin{cases} 1 & \text{if } \min_{j:e(m-N,j)<10} |x_{m-N} - x_j| \geq 2 \\ 0 & \text{otherwise,} \end{cases}$$

and for $m = 2N + 1, \dots, 3N$

$$\Phi_m(\vec{x}) = \begin{cases} 1 & \text{if } \min_{j:e(m-2N,j)<30} |x_{m-2N} - x_j| \geq 1 \\ 0 & \text{otherwise,} \end{cases}$$

The attenuation between adjacent channels is -28dB [36]. Taking the radio path loss with distance as d^α , where d is distance in meters and $\alpha = 4$ the path loss exponent, then these constraints ensure > 60 dB attenuation between APs.

Assuming all APs use the maximum transmit power of 18dBm allowed by the 802.11 standard, this means that the SINR is greater than 20dB within a 10m radius of each AP which is sufficient to sustain a data rate of 54Mbps when the

¹The extracted (x,y,z) coordinate data used is available online at www.hamilton.ie/net/xyz.txt

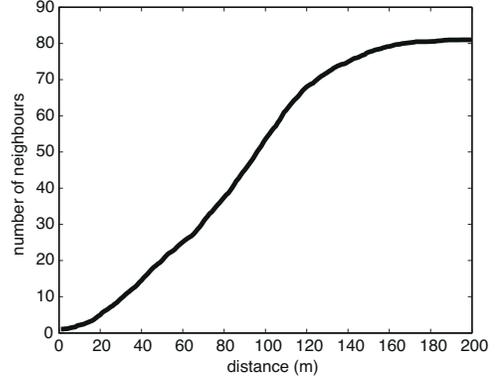


Fig. 4. Mean number of APs within distance d of an AP vs distance d .

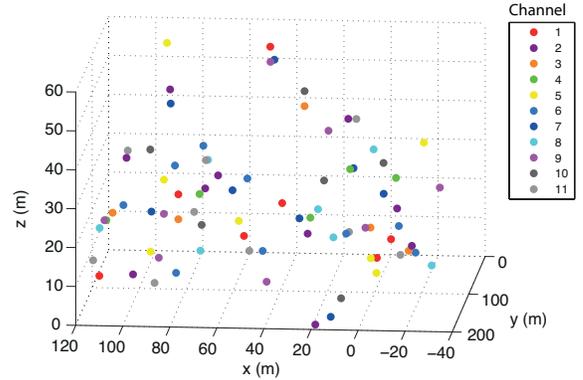


Fig. 5. A satisfying assignment of radio channels. Each dot marks the location of a WiFi wireless access point and is based on measurements taken at the junction of 5th Avenue/59th Street in Manhattan. The color of a dot indicates the radio channel. To avoid interference between transmissions, nearby access points need to operate on radio channels that are spaced sufficiently far apart. There are 11 radio channels available to choose from, but the frequencies of these channels overlap so that only 3 orthogonal/non-overlapping channels are available; there are 81 wireless access points in total.

connection is line of sight and channel noise is Gaussian [36], [37].

Observe that each variable x_i is located at a different AP. The APs do not belong to a single administrative domain and so security measures such as firewalls prevent communication over the wired network. An AP cannot rely on decoding wireless transmissions to communicate with all of its interferers or even just to identify them. This is because interferers may be too far away to allow their transmissions to be decoded and yet still their aggregate transmission power may be sufficiently powerful to create significant interference. That is, an AP cannot know the channel selections of other APs, condition (D3) and, moreover, cannot reliably identify or enumerate the clauses in which it participates, condition (D4). A decentralized algorithm is therefore mandated.

Fig. 5 shows an example satisfying assignment of radio channels obtained by running an instance of the CFL algorithm at all APs. The complexity of the topology generated by the physical location of the APs and the non-uniformity of the clauses it causes is apparent.

Fig. 6 shows the measured distribution of number of itera-

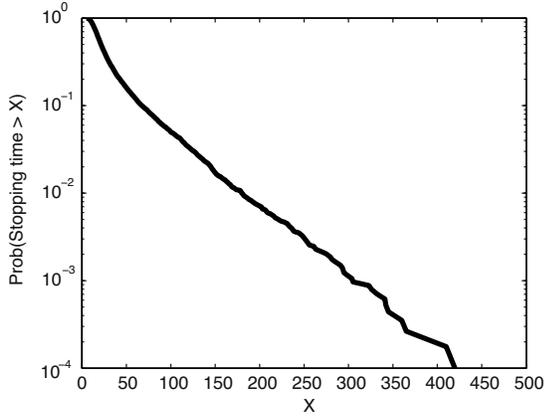


Fig. 6. Log of the empirical complementary cumulative distribution of convergence time in iterations, based on 12,000 runs of CFL algorithm for 5th Avenue data. The median is 21 and the 95% percentile is 98 iterations. In current hardware one iteration can readily be performed in under 10 seconds, leading to a median time to convergence of less than 4 minutes.

tions required to find a satisfying assignment, whereupon the algorithm natural halts in a decentralized fashion. The median value is 21 iterations and the 95% centile is 98 iterations. Note that during this convergence period, although the network is operating sub-optimally, it does not cease to function. In a prototype lab set-up we have shown that a CFL update interval of less than 10 seconds is feasible on current hardware. Thus the median time to convergence is under 4 minutes. This is a reasonable time-frame for practical purposes, particularly as subparts of the network are functioning during this convergence period, and thus the CFL algorithm offers a pragmatic solution to this difficult decentralized CSP for which existing solvers could not be employed.

VI. CONCLUSIONS

We have shown that apparently distinct problems in networking can be placed within the framework of CSPs, where - unlike with traditional motivating examples - the variables are co-located with devices that may not be able to communicate. We define the criteria that practical solvers must possess in order to be suitable for these problems, labeling them decentralized CSP solvers.

As existing solvers fail to meet one or more of these criteria, we introduce a decentralized algorithm for solving CSPs. We prove that it will almost surely find a satisfying solution if one exists. In doing so, we generate bounds for the speed of convergence of the algorithm. We suspect, however, that our bounds for a general CSP are not tight and conjecture that the real bound should be closer to the one we have for the specific case of CSPs corresponding to graph coloring.

Given how much information decentralization is sacrificing, surprisingly an experimental investigation of solving random k-SAT instances suggests that the algorithm is competitive with two of the most promising centralized k-SAT solvers, WalkSAT and Survey Propagation, on instances of random k-SAT with order one thousand variables. This raises the question: what is the performance cost of decentralized operation? This is particularly pertinent as the decentralized nature of the

CFL algorithm lends itself to parallelized computation as it has a small, fixed memory requirement per variable, making it suitable for use in circumstances where a centralized algorithm could also be used but would be difficult to implement in a distributed fashion.

APPENDIX

For each $i \in \{1, \dots, N\}$ let $\vec{p}_i(t) \in [0, 1]^D$, $t \in \mathbb{N}$, be the CFL probability vector for variable i at time t and $x_i(t)$ be the variable's value selected stochastically from $\vec{p}_i(t)$. Let $P(t) = (\vec{p}_1(t), \dots, \vec{p}_N(t))$ and $X(t) = (x_1(t), \dots, x_N(t))$ record the overall state of the CSP-CFL system.

The state of the probability vectors $\{P(t)\}$ forms a Markov Chain, or Iterated Function System, with place dependent probabilities [38]. The convergence time of the algorithm is the first time the chain enters an absorbing state representing a valid solution to the CSP. The absorbing states are identifiable in terms of the variable values $\{X(t)\}$:

$$A = \bigcup \{\vec{x} : \Phi_m(\vec{x}) = 1 \text{ for all } m \in \mathcal{M}\},$$

as if $X(t) \in A$, then $P(t+1) = (\delta_{x_1}, \dots, \delta_{x_N})$ so that $X(t+1) = X(t)$ almost surely and hence a solution to the CSP has been found. The algorithm's stopping time is

$$\tau := \inf_{t \geq 0} \{X(t) \in A\}.$$

Define $\gamma = \min(a, b)/(D - 1 + a/b)$ and let S be the set of states such that for all $i \in \mathcal{N}$ either $p_i = \delta_k$ for some $k \in \{1, \dots, D\}$ or $(p_i)_k \geq \gamma$ for all $k \in \{1, \dots, M\}$. The next lemma is obvious.

Lemma 1: For any integer $t \geq 0$, $P(t) \in S$.

Proof: [Theorem 1] The method of proof for both statements is similar: we create a sequence of events over $N - 1$ iterations that, regardless of the initial configuration, lead to a satisfying assignment with a probability that we find a lower bound for. Due to the Markovian nature of the algorithm and the independence of the probability of this event on its initial conditions, if this event does not occur in $N - 1$ iterations, it has the same probability of occurring in the next $N - 1$ iterations. This is what leads to the geometric nature of the bounds. The difference between a general CSP and those corresponding to coloring is that for the latter the variables are ensured to experience fewer dissatisfaction events before finding a satisfying assignment.

First consider a general CSP. Select an arbitrary valid solution $\vec{a} \in A$ with components $(\vec{a})_i$. For each $m \in \mathcal{M}$ define \mathcal{N}_m to be the variables that participate in clause m . For each $t \geq 0$, define the set of unsatisfied variables at time t by

$$U_t := \bigcup_{m \in \mathcal{M}} \{\mathcal{N}_m : \Phi_m(X(t)) = 0\}$$

and the number of unsatisfied variables to be $n_t = |U_t|$. At time 0 assume $X(0) = \vec{x}(0)$, some $\vec{x}(0)$. If $U_0 = \emptyset$ then $X(0) \in A$, the algorithm has found a solution and $\tau = 0$. If

$U_0 \neq \emptyset$, the algorithm has not yet found a solution. Define the components, $\vec{x}(1)_i$, of the vector $\vec{x}(1)$ by

$$\vec{x}(1)_i = \begin{cases} \vec{a}_i & \text{if } i \in U_0 \\ \vec{x}(0)_i & \text{if } i \notin U_0 \end{cases}$$

and, by Lemma 1,

$$P(X(1) = \vec{x}(1)|X(0) = \vec{x}(0)) \geq \gamma^{n_0}.$$

If $U_1 = \emptyset$ then $X(1) \in A$, the algorithm has found a solution and $\tau = 1$. If $U_1 \neq \emptyset$, the algorithm has not yet found a solution and we proceed in an iterative fashion. For $t \geq 0$, define the components, $\vec{x}(t+1)_i$, of the vector $\vec{x}(t+1)$ by

$$\vec{x}(t+1)_i = \begin{cases} \vec{a}_i & \text{if } i \in U_t \\ \vec{x}(t)_i & \text{if } i \notin U_t \end{cases}$$

and, by Lemma 1,

$$P(X(t+1) = \vec{x}(t+1)|X(t) = \vec{x}(t)) \geq \gamma^{n_t}.$$

Let $t^* = \inf\{t : n_t = 0\}$ and note that

$$P(X(t^*) \in A) \geq$$

$$P(X(t^*) = \vec{x}(t^*), \dots, X(0) = \vec{x}(0)) \geq \gamma^{\sum_{t=0}^{t^*} n_t}.$$

We wish, therefore, to place an upper bound on the sum in the exponent. Observe that $t^* \leq N-1$, as one starts with at least one unsatisfied variable at each stage we must include at least one new variable and, therefore, this procedure must terminate in no more than $N-1$ steps with a valid assignment.

In general, the $\{n_t^*\}$ sequence that maximizes this sum, even though this sequence may not be feasible for a given CSP, is

$$n_0^* = 1, n_1^* = 2, \dots, n_{N-1}^* = N-1,$$

This occurs if we start with one unsatisfied variable and, in changing to the valid solution, at each iteration one additional variable has cause to be unsatisfied. On taking its new value, the unsatisfied variable included at time t triggers the failure of $t-1$ clauses each of which contains 3 variables: itself, the variable included immediately previously and one of the other variables that has already experienced a failure. This gives

$$\sum_{t=0}^{t^*} n_t \leq \sum_{t=0}^{N-1} n_t^* = \frac{(N-1)N}{2}.$$

and therefore

$$P(X(N-1) \in A) \geq \gamma^{(N-1)N/2}.$$

For a CSP corresponding to graph coloring with D colours, the advantage is that variables cannot be dissatisfied indefinitely by newly dissatisfied variables. Instead this procedure generates a flame-front of dissatisfied variables where those variables sufficiently far within the interior of this flame-front can select their final value without further disturbance. For graph colouring, clauses occur only for each pair of neighbors, i and j , such that $\Phi(\vec{x}) = 0$ if $(\vec{x})_i = (\vec{x})_j$ and $\Phi(\vec{x}) = 1$ otherwise. For any initial configuration $\vec{x}(0)$ with dissatisfaction, we must have $n_0 \geq 2$. In a sequence of events analogous to those in the proof of [6][Theorem 3], these two variables select their final values, \vec{a}_i and $(\vec{a})_j$, at the next

round, causing their initial clause to be satisfied as $\vec{a} \in A$. If either of the variables taking its final value causes no further clause to be dissatisfied, then in this sequence of events it will never appear in a dissatisfied clause again and stays with its value with probability 1. If it does cause one or more clause to be triggered, once these are satisfied it will experience no more dissatisfaction. Thus each variable experiences at most 2 rounds of dissatisfaction so that for graph colouring:

$$\sum_{t=0}^{t^*} n_t \leq \sum_{t=0}^{N-1} n_t^* = 2N.$$

■

As an example that illustrates the difference between a general CSP and graph colouring, consider an instance of 3-SAT starting with one unsatisfied variable, x_1 . The following sequence of clauses is possible in general, but not with graph colouring. Starting with an unsatisfied x_1 , changing its value to satisfy the first clause causes x_2 to be dissatisfied, which in turn causes x_1 and x_2 and x_3 to be dissatisfied:

$$x_1, \neg x_1 \vee x_2 \text{ and } \neg x_1 \vee \neg x_2 \vee x_3.$$

Satisfying these clauses for x_1 , x_2 and x_3 causes two more clauses with x_4 to trigger, keeping all variables unsatisfied:

$$\neg x_1 \vee \neg x_2 \vee x_4 \text{ and } \neg x_1 \vee \neg x_3 \vee x_4.$$

Satisfying these two clauses causes three further clauses involving x_5 to become unsatisfied:

$$\neg x_1 \vee \neg x_2 \vee x_5, \neg x_1 \vee \neg x_3 \vee x_5 \text{ and } \neg x_1 \vee \neg x_4 \vee x_5$$

and so forth triggering clauses with three variables at a time such that all variables are unsatisfied until a solution is found.

REFERENCES

- [1] H. Wu, Fan Y. Yang, K. Tan, J. Chen, Q. Zhang, and Z. Zhang, "Distributed channel assignment and routing in multiradio multichannel multihop wireless networks," *IEEE Trans. Wireless Commun.*, vol. 9, no. 10, pp. 1972–1983, 2006.
- [2] B. Kauffmann, F. Baccelli, A. Chaintreau, V. Mhatre, K. Papagiannaki, and C. Diot, "Self organization of interfering 802.11 wireless access networks," in *IEEE INFOCOM*, 2007.
- [3] J. Crichigno, M.-Y. Wu, and W. Shu, "Protocols and architectures for channel assignment in wireless mesh networks," *Ad Hoc Netw.*, vol. 6, no. 7, pp. 1051–1077, Sept. 2008.
- [4] A. P. Subramanian, H. Gupta, S. R. Das, and J. Cao, "Minimum interference channel assignment in multiradio wireless mesh networks," *IEEE Trans. Mobile Comput.*, vol. 7, no. 12, pp. 1459–1473, Dec. 2008.
- [5] D. J. Leith and P. Clifford, "A self-managed distributed channel selection algorithm for WLANs," in *Proceedings of RAWNET, Boston, MA*, 2006.
- [6] K. Duffy, N. O'Connell, and A. Sapozhnikov, "Complexity analysis of a decentralised graph colouring algorithm," *Inform. Process. Lett.*, vol. 107, no. 2, pp. 60–63, 2008.
- [7] B. Selman, H. Kautz, and B. Cohen, "Local search strategies for satisfiability testing," in *DIMACS SERIES IN DISCRETE MATHEMATICS AND THEORETICAL COMPUTER SCIENCE*, 1995, pp. 521–532.
- [8] A. Braunstein, M. Mezard, and R. Zecchina, "Survey propagation: An algorithm for satisfiability," *Rand. Struct. Algorithms*, vol. 27, no. 2, pp. 201–226, 2005.
- [9] R. Koetter and M. Médard, "An algebraic approach to network coding," *IEEE/ACM Trans. Network.*, vol. 11, pp. 782–795, 2003.
- [10] S. Deb, M. Effros, T. Ho, D. R. Karger, R. Koetter, D. S. Lun, M. Médard, and N. Ratnakar, "Network coding for wireless applications: A brief tutorial," in *IWWAN*, 2005.
- [11] S. Katti, H. Rahul, W. Hu, D. Katabi, M. Médard, and J. Crowcroft, "XORs in the air: practical wireless network coding," *SIGCOMM Comput. Commun. Rev.*, vol. 36, no. 4, pp. 243–254, Aug. 2006.

- [12] S. Sengupta, S. Rayanchu, and S. Banerjee, "Network coding-aware routing in wireless networks," *IEEE/ACM Trans. Netw.*, vol. 18, no. 4, pp. 1158–1170, Aug. 2010.
- [13] M. Kim, M. Médard, U.-M. O'Reilly, and D. Traskov, "An evolutionary approach to inter-session network coding," in *INFOCOM*, 2009, pp. 450–458.
- [14] J. Barcelo, B. Bellalta, C. Cano, A. Sfaïropoulou, M. Oliver, and K. Verma, "Towards a collision-free WLAN: Dynamic parameter adjustment in CSMA/E2CA," *EURASIP Journal on Wireless Communications and Networking*, 2011.
- [15] M. Fang, D. Malone, K.R. Duffy, and D.J. Leith, "Decentralised learning MACs for collision-free access in WLANs," *Wirel. Netw.*, vol. To appear, 2012.
- [16] M. Davis and H. Putnam, "A computing procedure for quantification theory," *J. ACM*, vol. 7, pp. 201–215, July 1960.
- [17] M. Davis, G. Logemann, and D. Loveland, "A machine program for theorem-proving," *Commun. ACM*, vol. 5, pp. 394–397, July 1962.
- [18] Y. Hamadi and S. Jabbour, "ManySAT: a parallel SAT solver," *JSAT*, vol. 6, no. 4, pp. 245–262, 2009.
- [19] Ming Ouyang, "How good are branching rules in dpll?," *Discrete Appl. Math.*, vol. 89, pp. 281–286, December 1998.
- [20] J. Pearl, "Reverend Bayes on inference engines: a distributed hierarchical approach," in *Proceedings of American Association for Artificial Intelligence National Conference on AI*, 1982, pp. 133–136.
- [21] M. Mézard and R. Zecchina, "Random k -satisfiability problem: From an analytic solution to an efficient algorithm," *Phys. Rev. E*, vol. 66, no. 5, pp. 056126, Nov 2002.
- [22] C.H. Papadimitriou, "On selecting a satisfying truth assignment," in *IEEE Symposium on Foundations of Computer Science (FOCS)*. 1991, pp. 162–169, IEEE Computer Society, New York.
- [23] U. Schoning, "A probabilistic algorithm for k -SAT and constraint satisfaction problems," in *IEEE Symposium on Foundations of Computer Science (FOCS)*. 1999, pp. 410–414, IEEE Computer Society, New York.
- [24] D. Achlioptas, A. Naor, and Y. Peres, "Rigorous location of phase transitions in hard optimization problems," *Nature*, vol. 435, no. 9, pp. 759–764, 2005.
- [25] D. Achlioptas and Y. Peres, "The threshold for random k -SAT is $2k \log 2 - O(k)$," *J. Amer. Math. Soc.*, vol. 17, no. 4, pp. 947–973, 2004.
- [26] A.C. Kaporis, L.M. Kirousis, and E.G. Lalas, "The probabilistic analysis of a greedy satisfiability algorithm," *Rand. Struct. Algorithms*, vol. 28, no. 4, pp. 444–480, 2006.
- [27] A. M. Frieze and S. Suen, "Analysis of two simple heuristics on a random instance of k -SAT," *J. Algorithms*, vol. 20, no. 2, pp. 312–355, 1996.
- [28] S. Mertens, M. Mezard, and R. Zecchina, "Threshold values of random k -SAT from the cavity method," *Rand. Struct. Algorithms*, vol. 28, no. 3, pp. 340–37, 2005.
- [29] F Krzakala, A. Montarani, Ricci-Tersenghi F., G. Semerijan, and L. Zdeborova, "Gibbs states and the set of solutions of random constraint satisfaction problems," *Proc. Natl. Acad. Sci. USA*, vol. 104, no. 25, pp. 10318–10323, 2007.
- [30] E. Aurell, U. Gordon, and S. Kirkpatrick, "Comparing beliefs, surveys and random walks," in *Advances in Neural Information Processing Systems (NIPS)*. 2004, vol. 17, pp. 49–56, MIT Press, Boston.
- [31] M. Mezard, G. Parisi, and R. Zecchina, "Analytic and algorithmic solution of random satisfiability problems," *Science*, vol. 297, pp. 812–815, 2002.
- [32] M. Mezard, T. Mora, and R. Zecchina, "Clustering of solutions in the random satisfiability problem," *Phys. Rev. Lett.*, vol. 94, pp. 197–205, 2005.
- [33] M. Alava, J. Ardeliu, E. Aurell, P. Kaski, S. Krishnamurthy, and P. Orponen, "Circumspect descent prevails in solving random constraint satisfaction problems," *Proc. Natl. Acad. Sci. USA*, vol. 105, no. 4, pp. 15253–15257, 2008.
- [34] M. Alekhovich and E. Ben-Sasson, "Linear upper bounds for random walk on small density 3-CNFs," *SIAM J. Comput.*, vol. 36, no. 5, pp. 1248–1263, 2006.
- [35] WIGLE.NET, " <http://www.wigle.net/>, 2010.
- [36] *Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications. Further Higher Data Rate Extension in the 2.4GHz Band*, IEEE, New York, 2003.
- [37] A. Goldsmith, *Wireless Communications*, Cambridge University Press, 2005.
- [38] M. F. Barnsley, S. G. Demko, J. H. Elton, and J. S. Geronimo, "Invariant measures for Markov processes arising from iterated function systems with place-dependent probabilities," *Ann. Inst. H. Poincaré Probab. Statist.*, vol. 24, no. 3, pp. 367–394, 1988.