# Linear Time Approximation Schemes for the Gale-Berlekamp Game and Related Minimization Problems 

Marek Karpinski ${ }^{*}$<br>Dept. of Computer Science<br>University of Bonn<br>marek@cs.uni-bonn.de

Warren Schudy ${ }^{\dagger}$<br>Dept. of Computer Science<br>Brown University<br>ws@cs.brown.edu


#### Abstract

We design a linear time approximation scheme for the GaleBerlekamp Switching Game and generalize it to a wider class of dense fragile minimization problems including the Nearest Codeword Problem (NCP) and Unique Games Problem. Further applications include, among other things, finding a constrained form of matrix rigidity and maximum likelihood decoding of an error correcting code. As another application of our method we give the first linear time approximation schemes for correlation clustering with a fixed number of clusters and its hierarchical generalization. Our results depend on a new technique for dealing with small objective function values of optimization problems and could be of independent interest.


## Categories and Subject Descriptors

F.2.2 [Analysis of Algorithms and Problem Complexity]: Nonnumerical Algorithms and Problems-Computations on discrete structures

## General Terms

Algorithms, Theory

## Keywords

Approximation Algorithms, Correlation and Hierarchical Clustering, Dense Instances, Gale-Berlekamp Game, Linear Time Approximation Schemes, Minimum Constraint Satisfaction, Nearest Codeword Problem

## 1. INTRODUCTION

The Gale-Berlekamp Switching Game (GB Game) was introduced independently by Elwyn Berlekamp [10, 23] and David Gale [23] in the context of coding theory. This game

[^0][^1]is played using of a $m$ by $m$ grid of lightbulbs. The adversary chooses an arbitrary subset of the lightbulbs to be initially "on." Next to every row (resp. column) of lightbulbs is a switch, which can be used to invert the state of every lightbulb in that row (resp. column). The protagonist's task is to minimize the number of lit lightbulbs (by flipping switches). This problem was proven very recently to be NP-hard [21]. Let $\Phi=\{-1,1\} \subset \mathbb{R}$. For matrices $M, N$ let $d(M, N)$ denote the number of entries where $M$ and $N$ differ. It is fairly easy to see that the GB Game is equivalent to the following natural problems: [21]

- Given matrix $M \in \Phi^{m \times m}$ find column vectors $x, y \in$ $\Phi^{m}$ minimizing $d\left(M, x y^{T}\right)$.
- Given matrix $M \in \mathbb{F}_{2}^{m \times m}$ find $x, y \in \mathbb{F}_{2}^{m}$ minimizing $\sum_{i j} \mathbb{1}\left(M_{i j} \neq x_{i} \oplus y_{j}\right)$ where $\mathbb{F}_{2}$ is the finite field over two elements with addition operator $\oplus$.
- Given matrix $M \in \Phi^{m \times m}$ find column vectors $x, y \in$ $\Phi^{m}$ maximizing $x^{T} M y$.

We focus on the equivalent minimization versions and prove existence of linear-time approximation schemes for them.

Theorem 1. For every $\epsilon>0$ there is a randomized $1+$ $\epsilon$-approximation algorithm for the Gale-Berlekamp Switching Game (its minimization version) with runtime $O\left(m^{2}\right)+$ $2^{O\left(1 / \epsilon^{2}\right)}$.

A constraint satisfaction problem (CSP) consists of $n$ variables over a domain of constant-size $d$ and a collection of arity- $k$ constraints ( $k$ constant). The objective of MIN$k$ CSP (MAX- $k \mathrm{CSP}$ ) is to minimize the number of unsatisfied (maximize the number of satisfied) constraints. An (everywhere) dense instance is one where every variable is involved in at least a constant times the maximum possible number of constraints, i.e. $\Omega\left(n^{k-1}\right)$. For example, the GB Game is a dense MIN-2CSP since each of the $n=2 m$ variables is involved in precisely $m=n / 2$ constraints. It is natural to consider generalizing Theorem 1 to all dense MIN-CSPs, but unfortunately many such problems have no PTASs unless $\mathrm{P}=\mathrm{NP}$ [7] so we must look at a restricted class of MIN-CSPs. A constraint is fragile if modifying any variable in a satisfied constraint makes the constraint unsatisfied. A CSP is fragile if all of its constraints are. Clearly the GB Game can be modeled as a fragile dense MIN-2CSP. Our results generalize to all dense fragile MIN- $k$ CSPs.

We now formulate our general theorem.

Theorem 2. For every $\epsilon>0$ there is a randomized $1+\epsilon$ approximation algorithm for dense fragile MIN-kCSPs with runtime $O\left(n^{k}\right)+2^{O\left(1 / \epsilon^{2}\right)}$.

Any approximation algorithm for MIN- $k$ CSP must read the entire input to distinguish between instances with optimal value of 1 and 0 and hence the $O\left(n^{k}\right)$ term of the runtime cannot be improved $\sqrt{1}$ It is fairly easy to see that improving the second term (to $2^{o\left(1 / \epsilon^{2}\right)}$ ) would imply a $O\left(n^{2}\right)+2^{o\left(1 / \epsilon^{2}\right)}$ time PTAS for average-dense max cut. Over a decade worth of algorithms [5, 6, 16, 2, 20] for MAX-kCSP all have dependence on $\epsilon$ of at best $2^{O\left(1 / \epsilon^{2}\right)}$, so any improvement to the runtime of Theorem 2 would be surprising.

We begin exploring applications of Theorem 2 by generalizing the Gale-Berlekamp game to higher dimensions $k$ and to arbitrary $k$-ary equations. Given $n$ variables $x_{i} \in \mathbb{F}_{2}$ and $m$ linear equations of the form $x_{i_{1}} \oplus x_{i_{2}} \oplus \ldots \oplus x_{i_{k}}=0$ (or $=1$ ), the $k$-ary Nearest Codeword Problem (NCP) consists of finding an assignment minimizing the number of unsatisfied equations. As the name suggests, the Nearest Codeword Problem can be interpreted as maximum likelihood decoding for linear error correcting codes. The Nearest Codeword Problem has fragile constraints so Theorem 2 implies a linear-time PTAS for the dense $k$-ary Nearest Codeword Problem.

The Unique Games Problem (UGP) [12, 18] consists of solving MIN-2CSPs where the constraints are permutations over a finite domain $D$ of colors; i.e. a constraint involving variables $x_{u}$ and $x_{v}$ is satisfied iff $x_{u}=\pi_{u v}\left(x_{v}\right)$ for permutation $\pi_{u v}$. These constraints are clearly fragile, so Theorem 2 implies also a linear-time PTAS for the dense Unique Game Problem (with a constant number of colors).

The multiway cut problem, also known as MIN-dCUT, consists of coloring an undirected graph with $d$ colors, such that each of $d$ terminal nodes $t_{i}$ is colored with color $i$, minimizing the number of bichromatic edges. The requirement that the terminal nodes must be colored particular colors does not fit in our dense fragile MIN-CSP framework, so we use a work-around: let the constraint corresponding to an edge be satisfied only if it is monochromatic and the endpoint(s) that are terminals (if any) are colored correctly.

As another application, consider MIN-kSAT, the problem of minimizing the number of satisfied clauses of a boolean expression in conjunctive normal form where each clause has $k$ variables (some negated). We consider the equivalent problem of minimizing the number of unsatisfied conjunctions of a boolean expression in disjunctive normal form. A conjunction can be represented as a fragile constraint indicating that all of the negated variables within that constraint are false and the remainder are true, so Theorem 2 applies to MIN- $k$ SAT as well.

Finally we consider correlation clustering with a fixed number of clusters $[17,1]$. The input to the correlation clustering problem consists of a complete graph where each edge is labeled either "+" or "-". The objective is to color the vertices with $d$ colors minimizing the number of bichromatic " + " edges plus the number of monochromatic "-" edges. Correlation clustering with two clusters is equivalent to the following symmetric variant of the Gale-Berlekamp game: given a symmetric matrix $M \in \Phi^{m \times m}$ find a column vector $x \in \Phi^{m}$

[^2]minimizing $d\left(M, x x^{T}\right)$. Like the GB game, correlation clustering with 2 clusters is fragile and Theorem 2 gives a lineartime approximation scheme. For $d>2$ correlation clustering is not fragile but has properties allowing for a PTAS anyway. We also solve a generalization of correlation clustering called hierarchical clustering [1]. We prove the following theorem.

Theorem 3. For every $\epsilon>0$ there is a randomized $1+\epsilon$ approximation algorithm for correlation clustering and hierarchical clustering with fixed number of clusters $d$ with running time $n^{2} 2^{O\left((\log d) d^{6} / \epsilon^{2}\right)}$.

The above results improves on the running time $n^{O\left(9^{d} / \epsilon^{2}\right)}$. $\log n=n^{O\left(9^{d} / \epsilon^{2}\right)}$ of the previous PTAS for correlation clustering by Giotis and Guruswami [17] in two ways: first the polynomial is linear in the size of the input and second the exponent is polynomial in $d$ rather than exponential. Our result for hierarchical clustering with a fixed number of clusters is the first PTAS for this problem.

We prove Theorem 2 in Sections 2 and 3 and Theorem 3 in Sections 4 and 5.

## Related Work

Elwyn Berlekamp built a physical model of the GB game with either $m=8$ or $m=10[10,23]$ at Bell Labs in the 1960s motivated by the connection with coding theory and the Nearest Codeword Problem. Several works [15, 10 investigated the cost of worst-case instances of the GB Game; for example the worst-case instance for $m=10$ has cost 35 [10]. Roth and Viswanathan [21] showed very recently that the GB game is in fact NP-hard. They also give a linear-time algorithm if the input is generated by adding random noise to a cost zero instance. Replacing $\Phi$ with $\mathbb{R}$ in the third formulation of the GB Game yields the problem of computing the 1-rigidity of a matrix. Lower bounds on matrix rigidity have applications to circuit and communication complexity [19].

The Nearest Codeword Problem is hard to approximate in general [4, 11] better than $n^{\Omega(1 / \log \log n)}$. It is hard even if each equation has exactly 3 variables and each variable appears in exactly 3 equations [9]. There is a $O(n / \log n)$ approximation algorithm $[8,3]$.

Over a decade ago two groups [6, 13] independently discovered polynomial-time approximation algorithms for MAXCUT achieving additive error of $\epsilon n^{2}$, implying a PTAS for average-dense MAX-CUT instances. The fastest algorithms [2, 20] have constant runtime $2^{O\left(1 / \epsilon^{2}\right)}$ for approximating the value of any MAX- $k$ CSP over a binary domain $D$. This can be generalized to an arbitrary domain $D$. To see this, note that we can code $D$ in binary and correspondingly enlarge the arity of the constraints to $k\lceil\log |D|\rceil$. A random sample of $\tilde{O}\left(1 / \epsilon^{4}\right)$ variables suffices to achieve an additive approximation [2, 20, 22]. These results extend to MAXBISECTION [14].

Arora, Karger and Karpinski [6] introduced the first PTASs for dense minimum constraint satisfaction problems. They give PTASs with runtime $n^{O\left(1 / \epsilon^{2}\right)}$ [6] for min bisection and multiway cut (MIN-d-CUT). Bazgan, Fernandez de la Vega and Karpinski [7] designed PTASs for MIN-SAT and the nearest codeword problem (which generalizes the GB Game) with runtime $n^{O\left(1 / \epsilon^{2}\right)}$. Giotis and Guruswami [17] give a

PTAS for correlation clustering with $d$ clusters with runtime $n^{O\left(9^{d} / \epsilon^{2}\right)}$. We give linear-time approximation schemes for all of the problems mentioned in this paragraph except for the MIN-BISECTION problem.

## 2. FRAGILE-DENSE ALGORITHM

### 2.1 Intuition

Consider the following scenario. Suppose that our nemesis, who knows the optimal solution to the Gale-Berlekamp problem shown in Figure 1, gives us a constant size random sample of it to tease us. How can we use this information to construct a good solution? One reasonable strategy is to set each variable greedily based on the random sample. Throughout this section we will focus on the row variables; the column variables are analogous. For simplicity our example has the optimal solution consisting of all of the switches in one position, which we denote by $\alpha$. For row $v$, the greedy strategy, resulting in assignment $x^{(1)}$, is to set switch $v$ to $\alpha$ iff $\hat{b}(v, \alpha)<\hat{b}(v, \beta)$, where $\hat{b}(v, \alpha)$ (resp. $\hat{b}(v, \beta))$ denotes the number of light bulbs in the intersection of row $v$ and the sampled columns that would be lit if we set the switch to position $\alpha$ (resp. $\beta$ ).

With a constant size sample we can expect to set most of the switches correctly but a constant fraction of them will elude us. Can we do better? Yes, we simply do greedy again. The greedy prices analogous to $\hat{b}$ are shown in the columns labeled with $b$ in the middle of Figure 1. For the example at hand, this strategy works wonderfully, resulting in us reconstructing the optimal solution exactly, as evidenced by the $b\left(x^{(1)}, v, \alpha\right)<b\left(x^{(1)}, v, \beta\right)$ for all $v$. In general this does not reconstruct the optimal solution but provably gives something close.

Some of the rows, e.g. the last one, have $b\left(x^{(1)}, v, \alpha\right)$ much less than $b\left(x^{(1)}, v, \beta\right)$ while other rows, such as the first, have $b\left(x^{(1)}, v, \alpha\right)$ and $b\left(x^{(1)}, v, \beta\right)$ closer together. We call variables with $\left|b\left(x^{(1)}, v, \alpha\right)-b\left(x^{(1)}, v, \beta\right)\right|>\Theta(n)$ clearcut. Intuitively, one would expect the clearcut rows to be more likely correct than the nearly tied ones. In fact, we can show that we get all of the clearcut ones correct, so the remaining problem is to choose values for the rows that are close to tied. However, those rows have a lot of lightbulbs lit, suggesting that the optimal value is large, so it is reasonable to run an additive approximation algorithm and use that to set the remaining variables.

Finally observe that we can simulate the random sample given by the nemesis by simply taking a random sample of the variables and then doing exhaustive search of all possibly assignments of those variables. We have just sketched our algorithm.

Our techniques differ from previous work [7, 6, 17] in two key ways:

1. Previous work used a sample size of $O\left((\log n) / \epsilon^{2}\right)$, which allowed the clearcut variables to be set correctly after a single greedy step. We instead use a constantsized sample and run a second greedy step before identifying the clearcut variables.
2. Our algorithm is the first one that runs the additive error algorithm after identifying clearcut variables. Previous work ran the additive error algorithm at the beginning.


Figure 1: An illustration of our algorithmic ideas on the Gale-Berlekamp Game.

The same ideas apply to all dense fragile CSPs. In the remainder of the paper we do not explicitly discuss the GB Game but present our ideas in the abstract framework of fragile-dense CSPs.

### 2.2 Model

We now give a formulation of MIN- $k$ CSP that is suitable for our purposes. For non-negative integers $n$, $k$, let $\binom{n}{k}=$ $\frac{n!}{k!(n-k)!}$, and for a given set $V$ let $\binom{V}{k}$ denote the set of subsets of $V$ of size $k$ (analogous to $2^{S}$ for all subsets of $S$ ). There is a set $V$ of $n$ variables, each of which can take any value in constant-sized domain $D$. Let $x_{v} \in D$ denote the value of variable $v$ in assignment $x$.

Consider some set of variables $I \in\binom{V}{k}$. There may be many constraints over these variables; number them arbitrarily. Define $p(I, \ell, x)$ to be 1 if the $\ell$ th constraint over $I$ is unsatisfied in assignment $x$ and zero otherwise. For $I \in\binom{V}{k}$, we define $p_{I}(x)=\frac{1}{\eta} \sum_{\ell} p(I, \ell, x)$, where $\eta$ is a scaling factor to ensure $0 \leq p_{I}(x) \leq 1$ (e.g. $\eta=2^{k}$ for MIN- $k$ SAT). For notational simplicity we write $p_{I}$ as a function of a complete assignment, but $p_{I}(x)$ only depends on $x_{u}$ for variables $u \in I$. For $I \notin\binom{V}{k}$ define $p_{I}(x)=0$.

Definition 4. On input $V, p$ a minimum constraint satisfaction problem (MIN-kCSP) is a problem of finding an assignment $x$ minimizing $\operatorname{Obj}(x)=\sum_{I \in\binom{V}{k}} p_{I}(x)$.

Let $R_{v i}(x)$ be an assignment over the variables $V$ that agrees with $x$ for all $u \in V$ except for $v$ where it is $i$; i.e. $R_{v i}(x)_{u}=\left\{\begin{array}{ll}i & \text { if } u=v \\ x_{u} & \text { otherwise }\end{array}\right.$. We will frequently use the identity $R_{v x_{v}}(x)=x$. Let $b(x, v, i)=\sum_{I \in\binom{V}{k}: v \in I} p_{I}\left(R_{v i}(x)\right)$ be the number of unsatisfied constraints $v$ would be in if $x_{v}$ were set to $i$ (divided by $\eta$ ).

We say the $\ell$ th constraint over $I$ is fragile if $p\left(I, \ell, R_{v i}(x)\right)+$ $p\left(I, \ell,\left(R_{v j}(x)\right) \geq 1\right.$ for all assignments $x$, variables $v$ and distinct values $i$ and $j$.

Definition 5. A Min-kCSP is $\delta$-fragile-dense (or simply fragile-dense) for some $\delta>0$ if $b(x, v, i)+b(x, v, j) \geq \delta\binom{n}{k-1}$ for all assignments $x$, variables $v$ and distinct values $i$ and $j$.

Lemma 6. For any $\delta>0$ an instance where every variable $v \in V$ participates in at least $\delta \eta\binom{n}{k-1}$ fragile constraints is $\delta$-fragile-dense.

Proof. By definitions:

$$
\begin{aligned}
& b(x, v, i)+b(x, v, j) \\
& \quad=\sum_{I \in\binom{V}{k}: v \in I}\left(p_{I}\left(R_{v i}(x)\right)+p_{I}\left(R_{v j}(x)\right)\right) \\
& =\sum_{I \in\binom{V}{k}: v \in I} \frac{1}{\eta} \sum_{\ell}\left(p\left(I, \ell, R_{v i}(x)\right)+p\left(I, \ell, R_{v j}(x)\right)\right) \\
& \geq \sum_{I \in\binom{V}{k}: v \in I} \frac{1}{\eta} \cdot(\text { The number of fragile constraints over } I) \\
& \geq \frac{\delta \eta}{\eta}\binom{n}{k-1}=\delta\binom{n}{k-1}
\end{aligned}
$$

We will make no further mention of individual constraints, $\eta$ or fragility; our algorithms and analysis use $p_{I}$ and the fragile-dense property exclusively.

### 2.3 Algorithm

Recall that $b\left(x^{*}, v, i\right)$ can be written as a sum, over $S \in$ $\binom{V}{k-1}$, of $p_{S \cup\{v\}}\left(R_{v i}\left(x^{*}\right)\right)$, which is a function of $x_{u}^{*}$ for $u \in S$. We estimate $b\left(x^{*}, v, i\right)$ with a random sample of $s=\frac{18 \log (480|D| k / \delta)}{\delta^{2}}$ of its terms. Let $S_{1}, S_{2}, \ldots, S_{s}$ be independent random samples of $k-1$ variables from $V$. Later (see Lemma 13) we show that

$$
\hat{b}(v, i)=\frac{\binom{n}{k-1}}{s} \sum_{j=1}^{s} p_{S_{j} \cup\{v\}}\left(R_{v i}\left(\hat{x^{*}}\right)\right)
$$

is an unbiased estimator of $b\left(x^{*}, v, i\right)$.
We now describe our linear-time Algorithm 1 for fragiledense MIN- $k$-CSPs. We first (lines $1-3$ ) dispose of the easy case $\mathrm{OPT}=\Omega\left(\binom{n}{k}\right)$ by running an additive error algorithm and returning the output if it is sufficiently costly. Second (lines 5-7) we take a random sample of variables. The optimal assignment $x^{*}$ is of course unknown to us so we simply try every possible assignment of the sampled variables (line 7). We then compute $\hat{b}$ (line 8) and then make a preliminary assignment $x^{(1)}$ by setting each variable greedily relative to $\hat{b}$ (line 9). To reduce noise we do a second greedy step, yielding assignment $x^{(2)}$ (line 10). While constructing $x^{(2)}$ we make a note when the best value for a variable is far better than the second best; we fix such clear-cut variables to their values in assignment $x^{(2)}$. We finally run an additive error algorithm (line 12) to set the remaining variables.

## 3. ANALYSIS OF ALGORITHM 1

We use one of the known additive error approximation algorithms for MAX- $k$ CSP problems.

Theorem 7. [20] For any MAX-kCSP (or MIN-kCSP) and any $\epsilon^{\prime}>0$ there is a randomized algorithm which returns an assignment of cost $Y$ such that $|Y-O P T| \leq \epsilon^{\prime} n^{k}$ in runtime $O\left(n^{k}\right)+2^{O\left(1 / \epsilon^{\prime 2}\right)}$.

Throughout the rest of the paper let $x^{*}$ denote an optimal assignment.

First consider Algorithm 1 when the "then" branch of the "if" is taken. Choose constants appropriately so that the additive error algorithm fails with probability at most $1 / 10$ and

```
Algorithm 1 Our \(1+\epsilon\) approximation for \(\delta\)-fragile-dense
MIN- \(k\) CSP with variables taking values from domain \(D\).
    Run a \(\frac{\epsilon}{1+\epsilon} \cdot \frac{\delta^{2}}{72 k}\binom{n}{k}\) additive approximation algorithm.
    if \(\operatorname{Obj}(\) answer \() \geq\binom{ n}{k} \delta^{2} /(72 k)\) then
        Return answer.
    else
        Let \(s=\frac{18 \log (480|D| k / \delta)}{\delta^{2}}\)
        Draw \(S_{1}, S_{2}, \ldots, S_{s}\) randomly from \(\binom{V}{k-1}\) with re-
        placement.
        for Each assignment \(\hat{x^{*}}\) of the variables in \(\bigcup_{j=1}^{s} S_{j}\)
        do
            For all \(v \in V\) and \(i \in D\) let
            \(\hat{b}(v, i)=\frac{\left(\begin{array}{c}n-1\end{array}\right)}{s} \sum_{j=1}^{s} p_{S_{j} \cup\{v\}}\left(R_{v i}\left(\hat{x^{*}}\right)\right)\)
            For all \(v \in V\) let \(x_{v}^{(1)}=\arg \min _{i} \hat{b}(v, i)\)
            For all \(v \in V\) let \(x_{v}^{(2)}=\arg \min _{i} b\left(x^{(1)}, v, i\right)\)
            Let \(C=\left\{v \in V: \quad b\left(x^{(1)}, v, x_{v}^{(2)}\right)<b\left(x^{(1)}, v, j\right)-\right.\)
            \(\delta\binom{n}{k-1} / 6\) for all \(\left.j \neq x_{v}^{(2)}\right\}\).
            Find \(x^{(3)}\) of cost at most \(\frac{\epsilon|V \backslash C| \delta}{3 n}\binom{n}{k}+\min [\operatorname{Obj}(x)]\)
            using an additive approximation algorithm, where
            the minimum ranges over \(x\) such that \(x_{v}=\)
            \(x_{v}^{(2)} \forall v \in C\). See Section 3.1 for details.
        end for
        Return the best assignment \(x^{(3)}\) found.
    end if
```

assume it succeeds. Let $x^{a}$ denote the additive-error solution. We know $\operatorname{Obj}\left(x^{a}\right) \leq \operatorname{Obj}\left(x^{*}\right)+\frac{\epsilon}{1+\epsilon} P$ and $\operatorname{Obj}\left(x^{a}\right) \geq P$ where $P=\binom{n}{k} \delta^{2} /(72 k)$. Therefore $\operatorname{Obj}\left(x^{*}\right) \geq P\left(1-\frac{\epsilon}{1+\epsilon}\right)=$ $\frac{P}{1+\epsilon}$ and hence $\operatorname{Obj}\left(x^{a}\right) \leq \operatorname{Obj}\left(x^{*}\right)+\frac{\epsilon}{1+\epsilon}(1+\epsilon) \operatorname{Obj}\left(x^{*}\right)=$ $(1+\epsilon) \operatorname{Obj}\left(x^{*}\right)$. Therefore if the additive approximation is returned it is a $1+\epsilon$-approximation.

The remainder of this section considers the case when Algorithm 1 takes the "else" branch. Define $\gamma$ so that $\operatorname{Obj}\left(x^{*}\right)=$ $\gamma\binom{n}{k}$. We have $\operatorname{Obj}\left(x^{*}\right) \leq \operatorname{Obj}\left(x^{a}\right)<\binom{n}{k} \delta^{2} /(72 k)$ so $\gamma \leq$ $\delta^{2} /(72 k)$. We analyze the $\hat{x}^{*}$ where we guess $x^{*}$, that is when $\hat{x_{v}^{*}}=x_{v}^{*}$ for all $v \in \bigcup_{i=1}^{s} S_{i}$. Clearly the overall cost is at most the cost of $x^{(3)}$ during the iteration when we guess correctly.

Lemma 8. $b\left(x^{*}, v, x_{v}^{*}\right) \leq b\left(x^{*}, v, j\right)$ for all $j \in D$.
Proof. Immediate from definition of $b$ and optimality of $x^{*}$.

Lemma 9. For any assignment $x$,

$$
\operatorname{Obj}(x)=\frac{1}{k} \sum_{v \in V} b\left(x, v, x_{v}\right)
$$

Proof. By definitions,

$$
b\left(x, v, x_{v}\right)=\sum_{I \in\binom{V}{k}: v \in I} p_{I}\left(R_{v x_{v}}(x)\right)=\sum_{I \in\binom{V}{k}: v \in I} p_{I}(x) .
$$

Write $\operatorname{Obj}(x)=\sum_{I \in\binom{V}{k}} p_{I}(x)=\sum_{I \in\binom{V}{k}} p_{I}(x)\left[\sum_{v \in I} \frac{1}{k}\right]$ and reorder summations.

Definition 10. We say variable $v$ in assignment $x$ is corrupted if $x_{v} \neq x_{v}^{*}$.

Definition 11. Variable $v$ is clear if $\left(x^{*}, v, x_{v}^{*}\right)<b\left(x^{*}, v, j\right)-$ $\frac{\delta}{3}\binom{n}{k-1}$ for all $j \neq x_{v}^{*}$. A variable is unclear if it is not clear.

Clearness is the analysis analog of the algorithmic notion of clear-cut vertices sketched in Section 2.1. Comparing the definition of clearness to Lemma 8 further motivates the terminology "clear."

Lemma 12. The number of unclear variables $t$ satisfies

$$
t \leq 3(n-k+1) \gamma / \delta \leq \frac{\delta n}{24 k}
$$

Proof. Let $v$ be unclear and choose $j \neq x_{v}^{*}$ minimizing $b\left(x^{*}, v, j\right)$. By unclearness, $b\left(x^{*}, v, x_{v}^{*}\right) \geq b\left(x^{*}, v, j\right)-$ $(1 / 3) \delta\binom{n}{k-1}$. By fragile-dense, $b\left(x^{*}, v, x_{v}^{*}\right)+b\left(x^{*}, v, j\right) \geq$ $\delta\binom{n}{k-1}$. Adding these inequalities we see

$$
\begin{equation*}
b\left(x^{*}, v, x_{v}^{*}\right) \geq \frac{1-1 / 3}{2} \delta\binom{n}{k-1}=\frac{1}{3} \delta\binom{n}{k-1} \tag{1}
\end{equation*}
$$

By Lemma 9 and (1),

$$
\begin{aligned}
O P T & =\gamma\binom{n}{k}=1 / k \sum_{u} b\left(x^{*}, u, x_{u}^{*}\right) \\
& \geq 1 / k \sum_{u: \text { unclear }} \frac{\delta}{3}\binom{n}{k-1}=\frac{\delta}{3 k}\binom{n}{k-1} t
\end{aligned}
$$

Therefore $t \leq \gamma\binom{n}{k} \frac{3 k}{\delta\binom{n}{k-1}}=\frac{3 \gamma}{\delta}(n-k+1)$.
For the second bound observe $3 n \gamma / \delta \leq \frac{3 n}{\delta} \frac{\delta^{2}}{72 k}=\frac{\delta n}{24 k}$.
Lemma 13. The probability of a fixed clear variable $v$ being corrupted in $x^{(1)}$ is bounded above by $\frac{\delta}{240 k}$.

Proof. First we show that $\hat{b}(v, i)$ is in fact an unbiased estimator of $b\left(x^{*}, v, i\right)$ for all $i$. By definitions and in particular by the assumption that $p_{I}=0$ when $|I|<k$, we have for any $1 \leq j \leq s$ :

$$
\begin{aligned}
\mathbf{E}\left[p_{S_{j} \cup\{v\}}\left(R_{v i}\left(x^{*}\right)\right)\right] & =\frac{1}{\binom{n}{k-1}} \sum_{J \in\binom{V}{k-1}} p_{J \cup\{v\}}\left(R_{v i}\left(x^{*}\right)\right) \\
& =\frac{1}{\binom{n}{k-1}} \sum_{I \in\binom{V}{k}: v \in I} p_{I}\left(R_{v i}\left(x^{*}\right)\right) \\
& =\frac{1}{\binom{n}{k-1}} b_{v i}\left(x^{*}\right)
\end{aligned}
$$

Therefore

$$
\mathbf{E}[\hat{b}(v, i)]=s \frac{\binom{n}{k-1}}{s} \mathbf{E}\left[p_{S_{1} \cup\{v\}}\left(R_{v i}\left(x^{*}\right)\right)\right]=b\left(x^{*}, v, i\right) .
$$

Recall that $0 \leq p_{I}(x) \leq 1$ by definition of $p$, so by AzumaHoeffding,

$$
\begin{gathered}
\operatorname{Pr}\left[\left|\sum_{j=1}^{s} p_{S_{j} \cup\{v\}}\left(R_{v i}\left(x^{*}\right)\right)-\frac{s}{\left(k_{k-1}^{n}\right)} b\left(x^{*}, v, i\right)\right| \geq \lambda s\right] \\
\leq 2 e^{-2 \lambda^{2} s}
\end{gathered}
$$

for any $\lambda>0$ hence

$$
\operatorname{Pr}\left[\left|\hat{b}(v, i)-b\left(x^{*}, v, i\right)\right| \geq \lambda\binom{n}{k-1}\right] \leq 2 e^{-2 \lambda^{2} s}
$$

Choose $\lambda=\delta / 6$ and recall $s=\frac{18 \log (480|D| k / \delta)}{\delta^{2}}$, yielding.

$$
\operatorname{Pr}\left[\left|\hat{b}(v, i)-b\left(x^{*}, v, i\right)\right| \geq \frac{\delta}{6}\binom{n}{k-1}\right] \leq \frac{\delta}{240|D| k}
$$

By clearness we have $b\left(x^{*}, v, j\right)>b\left(x^{*}, v, x_{v}^{*}\right)+\delta\binom{n}{k-1} / 3$ for all $j \neq x_{v}^{*}$. Therefore, the probability that $\hat{b}\left(v, x_{v}^{*}\right)$ is not the smallest $\hat{b}(v, j)$ is bounded by $|D|$ times the probability that a particular $\hat{b}(v, j)$ differs from its mean by at least $\delta\binom{n}{k-1} / 6$. Therefore $\operatorname{Pr}\left[x_{v}^{(1)} \neq x_{v}^{*}\right] \leq|D| \frac{\delta}{240|D| k}=$ $\frac{\delta}{240 k}$.

Let $E_{1}$ denote the event that the assignment $x^{(1)}$ has at most $\delta n / 12 k$ corrupted variables.

Lemma 14. Event $E_{1}$ occurs with probability at least 1 1/10.

Proof. We consider the corrupted clear and unclear variables separately. By Lemma 12, the number of unclear variables, and hence the number of corrupted unclear variables, is bounded by $\frac{\delta n}{24 k}$.

The expected number of clear corrupted variables can be bounded by $\frac{\delta n}{240 k}$ using Lemma 13, so by Markov bound the number of clear corrupted variables is less than $\frac{\delta n}{24 k}$ with probability at least $1-1 / 10$.

Therefore the total number of corrupted variables is bounded by $\frac{\delta n}{24 k}+\frac{\delta n}{24 k}=\frac{\delta n}{12 k}$ with probability at least $9 / 10$.

We henceforth assume $E_{1}$ occurs. The remainder of the analysis is deterministic.

Lemma 15. For assignments $y$ and $y^{\prime}$ that differ in the assignment of at most $t$ variables, for all variables $v$ and values $i,\left|b(y, v, i)-b\left(y^{\prime}, v, i\right)\right| \leq t\binom{n}{k-2}$.

Proof. Clearly $p_{I}\left(R_{v i}(y)\right)$ is a function only of the variables in $I$ excluding $v$, so if $I-\{v\}$ consists of variables $u$ where $y_{u}=y_{u}^{\prime}$, then $p_{I}\left(R_{v i}(y)\right)-p_{I}\left(R_{v i}\left(y^{\prime}\right)\right)=0$. Therefore $b(y, v, i)-b\left(y^{\prime}, v, i\right)$ equals the sum, over $I \in\binom{V}{k}$ containing $v$ and at least one variable $u$ other than $v$ where $y_{u} \neq$ $y_{u}^{\prime}$, of $\left[p_{I}\left(R_{v i}(y)\right)-p_{I}\left(R_{v i}\left(y^{\prime}\right)\right)\right]$. For any $I, \mid p_{I}\left(R_{v i}(y)\right)-$ $p_{I}\left(R_{v i}\left(y^{\prime}\right)\right) \mid \leq 1$, so by the triangle inequality a bound on the number of such sets suffices to bound $\left|b(y, v, i)-b\left(y^{\prime}, v, i\right)\right|$. The number of such sets can trivially be bounded above by $t\binom{n}{k-2}$.

Lemma 16. Let $C=\left\{v \in V: b\left(x^{(1)}, v, x_{v}^{(2)}\right)<b\left(x^{(1)}, v, j\right)-\right.$ $\delta\binom{n}{k-1} / 6$ for all $\left.j \neq x_{v}^{(2)}\right\}$ as defined in Algorithm 1. If $E_{1}$ then:

1. $x_{v}^{(2)}=x_{v}^{*}$ for all $v \in C$.
2. $|V \backslash C| \leq \frac{3 n \gamma}{\delta}$.

Proof. Assume $E_{1}$ occurred. Using the definitions of corrupted and event $E_{1}$ together with Lemma 15 we have for any $v, i$ (for sufficiently large $n$ so that $\frac{n-k+1}{k-1} \geq \frac{n}{k}$ ):

$$
\begin{align*}
\left|b\left(x^{(1)}, v, i\right)-b\left(x^{*}, v, i\right)\right| & \leq \frac{\delta n}{12 k}\binom{n}{k-2} \\
& \leq \frac{\delta}{12}\binom{n}{k-1} \tag{2}
\end{align*}
$$

To prove the first statement of the Lemma, let $v$ be a clear-cut variable, i.e. $v \in C$. Apply (2), yielding:

$$
\begin{aligned}
& b\left(x^{*}, v, x_{v}^{(2)}\right) \\
& \quad \leq b\left(x^{(1)}, v, x_{v}^{(2)}\right)+\frac{\delta}{12}\binom{n}{k-1} \\
& \quad<b\left(x^{(1)}, v, j\right)-\frac{\delta}{6}\binom{n}{k-1}+\frac{\delta}{12}\binom{n}{k-1} \\
& \quad \leq b\left(x^{*}, v, j\right)+\left(-\frac{\delta}{6}+2 \frac{\delta}{12}\right)\binom{n}{k-1}=b\left(x^{*}, v, j\right) .
\end{aligned}
$$

So by Lemma 8, $x_{v}^{*}=x_{v}^{(2)}$.
To prove the second statement, let $u$ be clear and use (2) again:

$$
\begin{aligned}
b\left(x^{(1)}, u, x_{u}^{*}\right) & \leq b\left(x^{*}, u, x_{u}^{*}\right)+\frac{\delta}{12}\binom{n}{k-1} \\
& <b\left(x^{*}, u, j\right)-\frac{\delta}{3}\binom{n}{k-1}+\frac{\delta}{12}\binom{n}{k-1} \\
& \leq b\left(x^{(1)}, u, j\right)+\left(-\frac{\delta}{3}+2 \frac{\delta}{12}\right)\binom{n}{k-1} \\
& =b\left(x^{(1)}, u, j\right)-\frac{\delta}{6}\binom{n}{k-1}
\end{aligned}
$$

so $u \in C$ by definition of $C$. Therefore the second statement follows from bound on the number of unclear variables from Lemma 12.

### 3.1 Computation of $x^{(3)}$

This subsection has been overhauled since the proceedings version. The previous reduction had excessive runtime.
Now we give the details of the computation of $x^{(3)}$. Let $T=V \backslash C$. We call $C$ the clear-cut vertices and $T$ the tricky vertices. If $|T| \geq n / 2$ the desired approximation factor can be achieved using a direct application of any additive error algorithm so suppose $|T|<n / 2 \sqrt{2}$ It is natural to substitute the values of the clear-cut variables into the constraints, but doing these leaves constraints with arity less than $k$. To remedy this we create placeholder variables as follows. We divide $C$ into $|T|$ sets $C_{1}, C_{2}, \ldots, C_{|T|}$ all of size at most $2|C| /|T|$. We create new placeholder variables $P=\left\{\nu_{1}, \ldots, \nu_{|T|}\right\}$ to represent the variables in $C_{1}, C_{2}, \ldots, C_{|T|}$ respectively.

We create a new CSP over variables $T \cup P$ by replacing each clear-cut variable in $C$ with its representative in $P$. Constraints with insufficient arity are padded with arbitrary placeholder variables from $P$. Constraints that formerly depended only on variables in $C$ can sefely be discarded.

Formally, for any $H \subseteq C$ let $\mu(H)$ denote the variables representing $H$, i.e. $\mu(H)$ is an arbitrary subset of $P$ of size $|H|$ satisfying $H \subseteq \bigcup_{\nu_{i} \in \mu(H)} C_{i}$.

For any assignment $y$ of $T \cup P, M \subseteq T$ and $L \subseteq P$ with $|M|+|L|=k$, define

$$
q_{M, L}(y)=\mathbb{1}(|L| \leq k-1) \sum_{H \in\binom{C}{|L|}: \mu(H)=L} p_{M \cup H}\left(R_{T y}\left(x^{(2)}\right)\right)
$$

The following Lemma follows easily from the definitions:

[^3]Lemma 17. For any $y \in D^{2|T|}$ we have

$$
\operatorname{Obj}\left(R_{T y}\left(x^{(2)}\right)\right)=\sum_{K \in\binom{T}{k}} q_{K}(y)+\sum_{I \in\binom{C}{k}} p_{I}\left(x^{(2)}\right)
$$

Lemma 18.

$$
0 \leq q_{K}(y) \leq O\left(\left(\frac{|C|}{|T|}\right)^{k-1}\right)
$$

Proof. Let $M=K \cap T$ and $L=K \cap P$. If $|L|=k$ the Lemma is trivial so the interesting case is $|L| \leq k-1$. Recall that $\rho(H)=L$ implies $|H|=|L| \leq k-1$, so any $H$ contributing to the sum is a subset of $\bigcup_{\nu_{i} \in L} C_{i}$, which has size at most $2|L||C| /|T|$. Therefore the number of such $H$ is at most $\binom{2|L||C| /|T|}{k-1}=O\left(\left(\frac{|C|}{|T|}\right)^{k-1}\right)$.

After using Lemma 18 to normalize the weights, Theorem 7 with an error parameter of $\epsilon^{\prime}=\Theta(\epsilon)$ yields an additive error of $O\left(\epsilon|T|^{k}(|C| /|T|)^{k-1}\right)=O\left(\epsilon(|T| / n) n^{k}\right)$ for the problem of minimizing $\sum_{K \in\binom{\text { TUP }}{k}} q_{K}(y)$. By Lemma 17 this is also an additive error $O\left(\epsilon(|T| / n) n^{k}\right)$ for $\operatorname{Obj}\left(R_{T y}\left(x^{(2)}\right)\right)$ as stated in the algorithm. Using Lemma 16 we further bound the additive error $O\left(\epsilon(|T| / n) n^{k}\right)$ by $O\left(\epsilon \gamma n^{k}\right)$. Lemma 16 implies that $x^{*}=R_{T y}\left(x^{(2)}\right)$ for some $y$, so this yields an additive error $O\left(\epsilon \gamma n^{k}\right)=\epsilon O P T$ for our original problem of minimizing $\operatorname{Obj}(x)$ over all assignments $x$.

### 3.2 Runtime

In this subsection we heavily exploit our choice to hide $\delta$, $|D|$ and $k$ within $O(\cdot)$. Note $s=O(1)$, hence the number
 which we may safely ignore when analyzing runtime. The two calls to the additive error algorithm take, by Theorem 7, $O\left(n^{k}\right)+2^{O\left(1 / \epsilon^{2}\right)}$ time. Observe that $b(x, v, i)$ can be easily computed in time $O\left(n^{k-1}\right)$ directly from its definition, hence $b\left(x^{(1)}, v, i\right)$ can be computed for all all $v, i$ in $|D| n O\left(n^{k-1}\right)=O\left(n^{k}\right)$ time. Computing the cost of a particular assignment can also be done in $O\left(n^{k}\right)$ time. It is easy to see that the tasks not specifically mentioned take $O(n)$ time, so the overall runtime is $O\left(n^{k}\right)+2^{O\left(1 / \epsilon^{2}\right)}$ as claimed.


Figure 2: An illustration of correlation clustering and the rigidity property. Only " + " edges are shown.

## 4. CLUSTERING ALGORITHM

### 4.1 Intuition

As we noted previously in Section 1, correlation clustering constraints are not fragile for $d>2$. Indeed, the constraint corresponding to a pair of vertices that are connected by a "-" edge can be satisfied by any coloring of the endpoints as long as the endpoints are colored differently. Fortunately there is a key observation in [17] that allows for the construction of a PTAS. Consider the cost-zero clustering shown on the left of Figure 2. Note that moving a vertex from a small cluster to another small one increases the cost very little, but moving a vertex from a large cluster to anywhere else increases the cost a lot. Fortunately most vertices are in big clusters so, as in [17], we can postpone processing the vertices in small clusters. We use the above ideas, which are due to [17], the fragile-dense ideas sketched above, plus some additional ideas, to analyze our correlation clustering algorithm.

To handle hierarchical clustering (c.f. [1]) we need a few more ideas. Firstly we abstract the arguments of the previous paragraph to a CSP property rigidity. Secondly, we note that the number of trees with $d$ leaves is a constant and therefore we can safely try them all. We remark that all fragile-dense problems are also rigid.

### 4.2 Reduction to Rigid MIN-2CSP

We now define hierarchical clustering formally (following [1]). For integer $M \geq 1$, an $M$-level hierarchical clustering of $n$ objects $V$ is a rooted tree with the elements of $V$ as the leaves and every leaf at depth (distance to root) exactly $M+1$. For $M=1$, a hierarchical clustering has one node at the root, some "cluster" nodes in the middle level and all of $X$ in the bottom level. The nodes in the middle level can be identified with clusters of $V$. We call the subtree induced by the internal nodes of a $M$-level hierarchical clustering the trunk. We call the leaves of the trunk clusters. A hierarchical clustering is completely specified by its trunk and the parent cluster of each leaf.

For a fixed hierarchical clustering and clusters $i$ and $j$, let $f(i, j)$ be the distance from $i$ (or $j$ ) to the least common ancestor of $i$ and $j$. For example when $M=1, f(i, j)=$ $\mathbb{1}(i=j)$.

We are given a function $F$ from pairs of vertices to $\{0,1, \ldots$, $M\}{ }^{3}$ The objective of hierarchical clustering is to output a $M$-level hierarchical clustering minimizing $\left.\sum_{u, v} \frac{1}{M} \right\rvert\, F(u, v)-$

[^4]$f($ parent $(u)$, parent $(v)) \mid$. Hierarchical clustering with $d$ clusters is the same except that we restrict the number of clusters (recall that equals the number of nodes whose children are leaves) to at most $d$. It is easy to see that the special case of hierarchical clustering with $M=1$ is equivalent to the correlation clustering problem described in the introduction.

Lemma 19. The number of possible trunks is at most $d^{(M-1) d}$.

Proof. The trunk can be specified by giving the parent of all non-root nodes. There are at most $d$ nodes on each of the $M-1$ non-root levels so the lemma follows.

We now show how to reduce hierarchical clustering with a constant number of clusters to the solution of a constant number of min-2CSPs with cardinality of $D$ equal to $d$. We use a variant of the notation used in Sections 2 and 3 that is specialized for MIN-2CSPs. For vertices $u, v$ and values $i, j$, let $p_{u, v}(i, j)$ be the cost of putting $u$ in cluster $i$ and $v$ in cluster $j$. This is the same concept as $p_{I}$ for the fragile case, but this notation is more convenient here. Define $b(x, v, i)=$ $\sum_{u \in V: u \neq v} p_{u, v}\left(x_{u}, i\right)$, which is identical to $b$ of the fragiledense analysis but expressed using different notation.

Definition 20. A MIN-2CSP is rigid if for some $\delta>0$, all $v \in V$ and all $j \neq x_{v}^{*}$

$$
b\left(x^{*}, v, x_{v}^{*}\right)+b\left(x^{*}, v, j\right) \geq \delta\left|\left\{u \in V \backslash\{v\}: x_{u}^{*}=x_{v}^{*}\right\}\right|
$$

Observe that $\left|\left\{u \in V \backslash\{v\}: x_{u}^{*}=x_{v}^{*}\right\}\right| \leq|V|=\binom{|V|}{2-1}$ hence any fragile-dense CSP is also rigid.

Lemma 21. If the trunk is fixed, hierarchical clustering can be expressed as a $1 / M$-rigid MIN-2CSP with $|D|=d$.

Proof. (C.f. Figure 2) Let $D$ be the leaves of the trunk (clusters). It is easy to see that choosing

$$
p_{u, v}(i, j)=\frac{1}{M}|f(i, j)-F(u, v)|
$$

yields the correct objective function. To show rigidity, fix vertex $v$, define $i=x_{v}^{*}$ and $\mathcal{C}_{i}=\left\{u \in V: x_{u}^{*}=i\right\}$. Fix $j \neq i$ and $u \in \mathcal{C}_{i} \backslash\{v\}$. Clearly $|f(i, i)-f(i, j)| \geq 1$, hence by triangle inequality $|F(u, v)-f(i, i)|+|F(u, v)-f(i, j)| \geq 1$, hence $p_{u, v}(i, i)+p_{u, v}(i, j) \geq 1 / M$. Summing over $u \in \mathcal{C}_{i} \backslash\{v\}$ we see

$$
\begin{aligned}
b\left(x^{*}, v, x_{v}^{*}\right)+b\left(x^{*}, v, j\right) & \geq \frac{1}{M}\left|\mathcal{C}_{i} \backslash\{v\}\right| \\
& =\delta\left|\left\{u \in V \backslash\{v\}: x_{u}^{*}=x_{v}^{*}\right\}\right|
\end{aligned}
$$

for $\delta=1 / M$.
Lemmas 21 and 19 suggest a technique for solving hierarchical clustering: guess the trunk and then solve the rigid MIN-2CSP. We now give our algorithm for solving rigid MIN-2CSPs.

### 4.3 Algorithm for Rigid MIN-2CSP

Algorithm 2 solves rigid MIN-2CSPs by identifying clearcut variables, fixing them to optimal values, and then recursing on the remaining "tricky" variables $T$. The recursion terminates when the remaining subproblem is sufficiently expensive for an additive approximation to suffice.

```
Algorithm 2 Approximation Algorithm for Rigid MIN-
2CSPs.
Return \(\mathrm{CC}(V\), blank assignment, 0\()\)
CC(tricky variables \(T\), assignment \(y\) of \(V \backslash T\), recursion depth
depth):
1: Find assignment of cost at most \(\frac{\epsilon}{1+\epsilon} \cdot \frac{\delta^{3}|T|^{2}}{6 \cdot 72^{2}|D|^{3}}+\)
    \(\min _{x: x_{v}=y_{v} \forall v \in V \backslash T}[\operatorname{Obj}(x)]\) using an additive approxi-
    mation algorithm.
    if \(\operatorname{Obj}(\) answer \() \geq \frac{\delta^{3}|T|^{2}}{6.72^{2}|D|^{3}}\) or depth \(\geq|D|+1\) then
        Return answer.
    else
        Let \(s=\frac{432^{2}|D|^{4} \log \left(1440|D|^{3} / \delta\right)}{2 \delta^{4}}\)
        Draw \(v_{1}, v_{2}, \ldots, v_{s}\) randomly from \(T\) with replace-
        ment.
        for Each assignment \(\hat{x^{*}}\) of the variables
        \(\left\{v_{1}, v_{2}, \ldots, v_{s}\right\}\) do
            For all \(v \in T\) and \(i\) let
            \(\hat{b}(v, i)=\frac{|T|}{s} \sum_{j=1}^{s} p_{v_{j}, v}\left(\hat{x}_{v_{j}}^{*}, i\right)+\sum_{u \in V \backslash T} p_{u, v}\left(y_{u}, i\right)\)
            For all \(v \in V\) let
                \(x_{v}^{(1)}= \begin{cases}y_{v} & \text { If } v \in V \backslash T \\ \arg \min _{i} \hat{b}(v, i) & \text { Otherwise }\end{cases}\)
                For all \(v \in T\) let \(x_{v}^{(2)}=\arg \min _{i} b\left(x^{(1)}, v, i\right)\)
                Let \(C=\left\{v \in T: b\left(x^{(1)}, v, x_{v}^{(2)}\right)<b\left(x^{(1)}, v, j\right)-\right.\)
                \(\frac{\delta|T|}{12|D|}\) for all \(\left.j \neq x_{v}^{(2)}\right\}\).
                Let \(T^{\prime}=T \backslash C\)
                Define assignment \(y^{\prime}\) by
                \(y_{v}^{\prime}=\left\{\begin{array}{ll}y_{v} & \text { If } v \in V \backslash T \\ x_{v}^{(2)} & \text { If } v \in C \\ \text { Undefined } & \text { If } v \in T \backslash C\end{array}\right.\).
                If \(\operatorname{CC}\left(T^{\prime}, y^{\prime}\right.\), depth +1\()\) is the best clustering so far,
                update best.
        end for
        Return the best clustering found.
    end if
```


## 5. ANALYSIS OF ALGORITHM 2

### 5.1 Runtime

Theorem 22. For any T, y, an assignment of cost at most $\epsilon^{\prime}|T|^{2}+\min _{x: x_{v}=y_{v} \forall v \in V \backslash T}[\operatorname{Obj}(x)]$ can be found in time $n^{2} 2^{O\left((\log |D|) / \epsilon^{\prime 2}\right)}$.

Proof. The problem is essentially a CSP on $T$ variables but with an additional linear cost term for each variable. It is fairly easy to see that Algorithm 1 in [20] has error proportional to the misestimation of $b$ and hence is unaffected by arbitrarily large linear cost terms. On the other hand, the more efficient Algorithm 2 in [20] needs to estimate the objective value from a constant-sized sample as well and hence does not seem to work for this type of problem.

In this subsection $O(\cdot)$ hides only absolute constants and $\tilde{O}(\cdot)$ hides one factor of $\log |D|$. Algorithm 2 has recursion depth at most $|D|+1$ and branching factor $|D|^{s}$, so the number of recursive calls is at most $\left(|D|^{s}\right)^{|D|+1}=2^{s(|D|+1) \log |D|}=$ $2^{\tilde{O}\left(|D|^{5} / \delta^{4}\right)}$. Each call spends $O\left(|D| n^{2}\right)$ time on miscellaneous tasks such as computing the objective value plus time required to run the additive error algorithm, which is $n^{2}$.
$2^{\tilde{O}\left(|D|^{6} / \epsilon^{2} \delta^{6}\right)}$ by Theorem 22. Therefore the runtime of Algorithm 2 is $n^{2} 2^{\tilde{O}\left(\frac{|D|^{6}}{\epsilon^{2} \delta^{\sigma}}\right)}$, where the $2^{\tilde{O}\left(|D|^{5} / \delta^{4}\right)}$ from the size of the recursion tree got absorbed into the $2^{\tilde{O}\left(\frac{|D|^{6}}{\epsilon^{2} \delta^{6}}\right)}$ from Theorem 22. For hierarchical clustering, $\delta=1 / M$ yields a runtime of $n^{2} 2^{\tilde{O}\left(\frac{|D|^{6} M^{6}}{\epsilon^{2}}\right)} \cdot|D|^{(M-1)|D|}=n^{2} 2^{O\left(\frac{(\log |D|)|D|^{6} M^{6}}{\epsilon^{2}}\right)}$.

As noted in the introduction this improves on the runtime of $n^{O\left(\frac{g^{|D|}}{\epsilon^{2}}\right)}$ of [17] for correlation clustering in two ways: the degree of the polynomial is independent of $\epsilon$ and $|D|$, and the dependence on $|D|$ is singly rather than doubly exponential.

### 5.2 Approximation

We fix optimal assignment $x^{*}$. We analyze the path through the recursion tree where we always guess $\hat{x}^{*}$ correctly, i.e. $\hat{x}_{v}^{*}=x_{v}^{*}$ for all $v \in\left\{v_{1}, v_{2}, \ldots, v_{s}\right\}$. We call this the principal path.

We will need the following definitions.
Definition 23. Variable $v$ is $m$-clear if $b\left(x^{*}, v, x_{v}^{*}\right)<$ $b\left(x^{*}, v, j\right)-m$ for all $j \neq x_{v}^{*}$. We say a variable is clear if it is $m$-clear for $m$ obvious from context. A variable is unclear if it is not clear.

Definition 24. A variable is obvious if it in cluster $\mathcal{C}$ in $O P T$ and it is $\delta|\mathcal{C}| / 3$-clear.

Definition 25. A cluster $\mathcal{C}$ of OPT is finished w.r.t. $T$ if $T \cap \mathcal{C}$ contains no obvious variables.

Lemma 26. With probability at least 8/10, for any triple ( $T, y$, depth) encountered on the principle path,

1. $y_{v}=x_{v}^{*}$ for all $v \in V \backslash T$ and
2. The number of finished clusters w.r.t. $T$ is at least depth.

Before proving Lemma 26 let us see why it implies Algorithm 2 has the correct approximation factor.

Proof. Study the final call on the principal path, which returns the additive approximation clustering. The second part of Lemma 26 implies that depth $\leq|D|$, hence we must have terminated because $\operatorname{Obj}($ answer $) \geq \frac{\delta^{3}|T|^{2}}{6.72^{2}|D|^{3}}$. By the first part of Lemma 26 the additive approximation gives error at most

$$
\frac{\epsilon}{1+\epsilon} \cdot \frac{\delta^{3}|T|^{2}}{6 \cdot 72^{2}|D|^{3}}+O P T
$$

so the approximation factor follows from an easy calculation.

Now we prove Lemma 26 by induction. Our base case is the root, which vacuously satisfies the inductive hypothesis since $V \backslash T=\{ \}$ and depth $=0$. We show that if a node ( $T, y$, depth) (in the recursion tree) satisfies the invariant then its child $\left(T^{\prime}, y^{\prime}\right.$, depth +1$)$ does as well. We hereafter analyze a particular ( $T, y$, depth) and assume the inductive hypothesis holds for them. There is only something to prove if a child exists, so we hereafter assume the additive error answer is not returned from this node and hence

$$
\begin{equation*}
O P T \leq O b j(\text { answer })<\frac{\delta^{3}|T|^{2}}{6 \cdot 72^{2}|D|^{3}} \tag{3}
\end{equation*}
$$

We now prove a number of Lemmas in this context, from which the fact that $T^{\prime}, y^{\prime}$, depth +1 satisfies the inductive hypothesis will trivially follow.

Lemma 27. The number of $\delta^{2}|T| /\left(216|D|^{2}\right)$-clear variables that are corrupted in $x^{(1)}$ is at most $\delta|T| /(72|D|)$ with probability at least $1-1 /(10|D|)$.

Proof. Essentially the same proof as for fragile MIN$k$ CSP shows that $\hat{b}(v, i)$ is an unbiased estimator of $b\left(x^{*}, v, i\right)$. This time the Azuma-Hoeffding inequality yields

$$
\operatorname{Pr}\left[\left|\hat{b}(v, i)-b\left(x^{*}, v, i\right)\right| \geq \lambda|T|\right] \leq 2 e^{-2 \lambda^{2} s}
$$

Choose $\lambda=\frac{\delta^{2}}{432|D|^{2}}$ and recall $s=\frac{432^{2}|D|^{4} \log \left(1440|D|^{3} / \delta\right)}{2 \delta^{4}}$, yielding.

$$
\operatorname{Pr}\left[\left|\hat{b}(v, i)-b\left(x^{*}, v, i\right)\right| \geq \frac{\delta^{2}|T|}{432|D|^{2}}\right] \leq \frac{\delta}{720|D|^{3}}
$$

By clearness we have $b\left(x^{*}, v, j\right)>b\left(x^{*}, v, x_{v}^{*}\right)+\frac{\delta^{2}|T|}{216|D|^{2}}$ for all $j \neq x_{v}^{*}$. Therefore, the probability that $\hat{b}\left(v, x_{v}\right)$ is not the smallest $\hat{b}(v, j)$ is bounded by $|D|$ times the probability that a particular $\hat{b}(v, j)$ differs from its mean by at least $\delta^{2}|T| /\left(432|D|^{2}\right)$. Therefore $\operatorname{Pr}\left[x_{v}^{(1)} \neq x_{v}^{*}\right] \leq|D|_{720|D|^{3}}=$ $\frac{\delta}{720|D|^{2}}$. Therefore, by Markov bound, with probability $1-$ $1 /(10|D|)$ the number of corrupted $\delta^{2}|T| /\left(216|D|^{2}\right)$-clear variables is at most $\delta|T| /(72|D|)$.

There are two types of bad events: the additive error algorithm failing and our own random samples failing. We choose constants so that each of these events has probability at most $1 /(10|D|)$. The principle path in the recursion tree has length at most $|D|$, so the overall probability of a bad event is at most $2 / 10$. We hereafter assume no bad events occur.

Lemma 28. The number of $\delta c / 3$-unclear variables in clusters of size at least $c+1$ is at most $\frac{6 O P T}{\delta c}$.

Proof. Let confusing variable refer to a $\delta c / 3$-unclear variable in a cluster of size at least $c+1$. Let $v$ be such a variable, in cluster $\mathcal{C}$ in OPT. By unclearness,

$$
b\left(x^{*}, v, x_{v}^{*}\right) \geq b\left(x^{*}, v, j\right)-\delta c / 3
$$

for appropriate $j \neq x_{v}^{*}$ and by rigidity

$$
b\left(x^{*}, v, x_{v}^{*}\right)+b\left(x^{*}, v, j\right) \geq \delta|\mathcal{C} \backslash\{v\}| \geq \delta c .
$$

Adding these inequalities we see $b\left(x^{*}, v, x_{v}^{*}\right) \geq \delta c / 3$. Therefore

$$
\begin{aligned}
O P T & =\frac{1}{2} \sum_{v} b\left(x^{*}, v, x_{v}^{*}\right) \geq \frac{1}{2} \sum_{v \text { confusing }} \delta c / 3 \\
& =\mid\{v \in T: v \text { confusing }\} \mid \delta c / 6
\end{aligned}
$$

so $\mid\{v \in T: v$ confusing $\} \left\lvert\, \leq \frac{6 O P T}{\delta c}\right.$.
Lemma 29. For all $v, i,\left|b\left(x^{(1)}, v, i\right)-b\left(x^{*}, v, i\right)\right| \leq \frac{\delta|T|}{24|D|}$
Proof. First we show bounds on three classes of corrupted variables:

1. Lemma 27 bounds the number of $\frac{\delta^{2}|T|}{216|D|^{2}}$-clear corrupted variables by $\frac{\delta|T|}{72|D|}$.
2. The number of variables in clusters of size at most $\frac{\delta|T|}{72|D|^{2}}$ is bounded by $\frac{\delta|T|}{72|D|}$.
3. The number of $\frac{\delta^{2}|T|}{216|D|^{2}}$-unclear corrupted variables in clusters of size ${ }^{4}$ at least $\frac{\delta|T|}{72|D|^{2}}+1$ is bounded by, using Lemma 28 and (3), $\frac{6 O P T}{\delta} \frac{72|D|^{2}}{\delta|T|} \leq \frac{\delta^{3}|T|^{2}}{6 \cdot 72 \cdot 75|D|^{3}}$. $\frac{6 \cdot 72|D|^{2}}{\delta^{2}|T|}=\frac{\delta|T|}{72|D|}$.

Therefore the total number of corrupted variables in $x^{(1)}$ is at most $\frac{\delta|T|}{72|D|}+\frac{\delta|T|}{72|D|}+\frac{\delta|T|}{72|D|}=\frac{\delta|T|}{24|D|}$. The easy observation that $\left|b\left(x^{(1)}, v, i\right)-b\left(x^{*}, v, i\right)\right|$ is bounded by the number of corrupted variables in $x^{(1)}$ proves the Lemma.

Lemma 30. There exists an obvious variable in $T$ that is in a cluster of size at least $|T| / 2|D|$.

Proof. We say a variable $v$ is confusing' if it is nonobvious and its cluster in OPT has size at least $|T| / 2|D|$. By Lemma 28 and (3)

$$
\begin{aligned}
\mid\{v \in T: v \text { confusing' }\} \mid & \leq \frac{12|D|}{\delta|T|} O P T \\
& \leq \frac{12|D|}{\delta|T|} \frac{\delta^{3}|T|^{2}}{6 \cdot 72^{2}|D|^{3}}<|T| / 2
\end{aligned}
$$

Simple counting shows there are at most $|T| / 2$ variables of $T$ in clusters of size less than $|T| / 2|D|$, hence there must be an obvious variable in a big cluster.

Lemma 31. The number of finished clusters w.r.t. $T^{\prime}$ strictly exceeds the number of finished clusters w.r.t. T.

Proof. Let $v$ be the variable promised by Lemma 30 and $\mathcal{C}_{i}$ its cluster in OPT. For any obvious variable $u$ in $\mathcal{C}_{i}$ note that $u$ is $\delta\left|\mathcal{C}_{i}\right| / 3 \geq \delta|T| / 6|D|$-clear, so Lemma 29 implies

$$
\begin{aligned}
& b\left(x^{(1)}, u, i\right) \\
& \quad \leq b\left(x^{*}, u, i\right)+\frac{\delta|T|}{24|D|}<b\left(x^{*}, u, j\right)+\frac{\delta|T|}{24|D|}-\frac{\delta|T|}{6|D|} \\
& \quad \leq b\left(x^{(1)}, u, j\right)+2 \frac{\delta|T|}{24|D|}-\frac{\delta|T|}{6|D|}=b\left(x^{(1)}, u, j\right)-\frac{\delta|T|}{12|D|}
\end{aligned}
$$

hence $u$ is in the set $C$ of clear variables defined in Algorithm 2. Therefore, no obvious variables in $\mathcal{C}_{i}$ are in $T^{\prime}$ so $\mathcal{C}_{i}$ is finished w.r.t. $T^{\prime}$. The existence of $v$ implies $\mathcal{C}_{i}$ is not finished w.r.t. $T$, so $\mathcal{C}_{i}$ is newly finished. To complete the proof note that $T^{\prime} \subseteq T$ so finished is a monotonic property.

Lemma 32. $\left(T^{\prime}, y^{\prime}\right)$ satisfy the invariant $v \in V \backslash T^{\prime} \rightarrow$ $y_{v}^{\prime}=x_{v}^{*}$.

Proof. Fix $v \in V \backslash T^{\prime}$. If $v \in T$ the conclusion follows from the inductive hypothesis. If $v \in T \backslash T^{\prime}=C$ we need to show $y_{v}^{\prime}=x_{v}^{*}$.
${ }^{4}$ To simplify the exposition we assume that $\frac{\delta|T|}{72|D|^{2}}$ is an integer. Sketch of general case: if $\frac{\delta|T|}{72|D|^{2}}$ is large fiddle with the constants and if $\frac{\delta|T|}{72|D|^{2}}$ is small $T$ is small so one can afford to guess all possible assignments of the variables in $T$.

Let $i=y_{v}^{\prime}$. For any $j \neq i$, use Lemma 29 to obtain

$$
\begin{aligned}
b\left(x^{*}, v, i\right) & \leq b\left(x^{(1)}, v, i\right)+\frac{\delta|T|}{24|D|} \\
& <b\left(x^{(1)}, v, j\right)+\frac{\delta|T|}{24|D|}-\frac{\delta|T|}{12|D|} \\
& \leq b\left(x^{*}, v, j\right)+2 \frac{\delta|T|}{24|D|}-\frac{\delta|T|}{12|D|}=b\left(x^{*}, v, j\right)
\end{aligned}
$$

so by optimality of $x^{*}$ we have the Lemma.
Lemmas 31 and 32 complete the inductive proof of Lemma 26 .

## Acknowledgements

We would like to thank Claire Mathieu and Joel Spencer for raising a question on approximability status of the GaleBerlekamp game and Alex Samorodintsky for interesting discussions.

## 6. REFERENCES

[1] N. Ailon and M. Charikar. Fitting tree metrics: Hierarchical clustering and phylogeny. In Proc. 46 th IEEE FOCS, pages 73-82, 2005.
[2] N. Alon, W. Fernandez de la Vega, R. Kannan, and M. Karpinski. Random Sampling and Approximation of MAX-CSP Problems. In Proc. 34th ACM STOC, pages 232-239, 2002. Journal version in J. Comput. System Sciences, 67(2):212-243, 2003.
[3] N. Alon, R. Panigrahy, and S. Yekhanin. Deterministic Approximation Algorithms for the Nearest Codeword Problem. Technical report, Elec. Coll. on Comp. Compl., ECCC TR08-065, 2008.
[4] S. Arora, L. Babai, J. Stern, and Z. Sweedyk. The Hardness of Approximate Optima in Lattices, Codes, and Systems of Linear Equations. In Proc. 34th IEEE FOCS, pages 724-733, Nov 1993. Journal version in J. Comput. System Sciences, 54(2):317-331, 1997.
[5] S. Arora, A. Frieze, and H. Kaplan. A New Rounding Procedure for the Assignment Problem with Applications to Dense Graph Arrangement Problems. In Proc. 37th IEEE FOCS, pages 21-30, Oct 1996. Journal version in Mathematical Programming, 92(1):1-36, 2002.
[6] S. Arora, D. Karger, and M. Karpinski. Polynomial Time Approximation Schemes for Dense Instances of NP-Hard Problems. In Proc. 27th ACM STOC, pages 284-293, 1995. Journal version in J. Comput. System Sciences, 58(1):193-210, 1999.
[7] C. Bazgan, W. Fernandez de la Vega, and M. Karpinski. Polynomial Time Approximation Schemes for Dense Instances of Minimum Constraint Satisfaction. Random Structures and Algorithms, 23(1):73-91, 2003.
[8] P. Berman and M. Karpinski. Approximating Minimum Unsatisfiability of Linear Equations. In Proc. 13th ACM-SIAM SODA, pages 514-516, 2002.
[9] P. Berman and M. Karpinski. Approximation Hardness of Bounded Degree MIN-CSP and MIN-BISECTION. In Proc. 29th ICALP, LNCS 2380, pages 623-632. Springer, 2002.
[10] J. Carlson and D. Stolarski. The Correct Solution to Berlekamp's Switching Game. Discrete Mathematics, 287(1-3):145-150, 2004.
[11] I. Dinur, G. Kindler, R. Raz, and S. Safra. Approximating CVP to Within Almost-Polynomial Factors is NP-Hard. Combinatorica, 23(2):205-243, 2003.
[12] U. Feige and L. Lovasz. Two prover one round proof systems: Their power and their problems. In Proc. 24th STOC, pages 733-741, 1992.
[13] W. Fernandez de la Vega. MAX-CUT has a Randomized Approximation Scheme in Dense Graphs. Random Struct. Algorithms, 8(3):187-198, 1996.
[14] W. Fernandez de la Vega, R. Kannan, and M. Karpinski. Approximation of Global MAX-CSP Problems. Technical Report TR06-124, Electronic Colloquim on Computation Complexity, 2006.
[15] P. C. Fishburn and N. J. Sloane. The Solution to Berlekamp's Switching Game. Discrete Math., 74(3):263-290, 1989.
[16] A. M. Frieze and R. Kannan. Quick Approximation to Matrices and Applications. Combinatorica, 19(2):175-220, 1999.
[17] I. Giotis and V. Guruswami. Correlation clustering with a fixed number of clusters. Theory of Computing, 2(1):249-266, 2006.
[18] A. Gupta and K. Talvar. Approximating unique games. In Proc. ${ }^{\text {17th }}$ ACM-SIAM SODA, pages 99-106, 2006.
[19] S. V. Lokam. Spectral Methods for Matrix Rigidity with Applications to Size-Depth Tradeoffs and Communication Complexity. In Proc. 36th IEEE FOCS, pages 6-15, 1995.
[20] C. Mathieu and W. Schudy. Yet Another Algorithm for Dense Max Cut: Go Greedy. In Proc. 19th ACM-SIAM SODA, pages 176-182, 2008.
[21] R. Roth and K. Viswanathan. On the Hardness of Decoding the Gale-Berlekamp Code. IEEE Transactions on Information Theory, 54(3):1050-1060, March 2008.
[22] M. Rudelson and R. Vershynin. Sampling from large matrices: An approach through geometric functional analysis. J. ACM, 54(4):21, 2007.
[23] J. Spencer. Ten Lectures on the Probabilistic Method. SIAM, Regional Conference Series, second edition, 1994.


[^0]:    *Parts of this work done while visiting Microsoft Research.
    ${ }^{\dagger}$ Parts of this work done while visiting University of Bonn.

[^1]:    Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
    (Corrected May 29, '09) STOC'09, May 31-June 2, 2009, Bethesda, Maryland, USA.
    Copyright 2009 ACM 978-1-60558-506-2/09/05 ...\$5.00.

[^2]:    ${ }^{1}$ Proof idea: consider a distribution of instances with cost zero $50 \%$ of the time and otherwise one random violated constraint and apply Yao's principle.

[^3]:    ${ }^{2}$ Alternatively Lemma 16 implies $|T| \leq \delta n /(24 k) \leq n / 2$

[^4]:    ${ }^{3}$ [1] chose $\{1,2, \ldots, M+1\}$ instead; the difference is merely notational.

