Revisiting Compressed Sensing: Exploiting the Efficiency of Simplex and Sparsification Methods

Robert Vanderbei

Kevin Lin

Han Liu

Lie Wang

Abstract We propose two approaches to solve large-scale compressed sensing problems. The first approach uses the parametric simplex method to recover very sparse signals by taking a small number of simplex pivots, while the second approach reformulates the problem using Kronecker products to achieve faster computation via a sparser problem formulation. In particular, we focus on the computational aspects of these methods in compressed sensing.

For the first approach, if the true signal is very sparse and we initialize our solution to be the zero vector, then a customized parametric simplex method usually takes a small number of iterations to converge. Our numerical studies show that this approach is 10 times faster than state-of-the-art methods for recovering very sparse signals.

The second approach can be used when the sensing matrix is the Kronecker product of two smaller matrices. We show that the best-known sufficient condition for the Kronecker compressed sensing (KCS) strategy to obtain a perfect recovery is more restrictive than the corresponding condition if using the first approach. However, KCS can be formulated as a linear program (LP) with a very sparse constraint matrix, whereas the first approach involves a completely dense constraint matrix. Hence, algorithms that benefit from sparse problem

Robert J. Vanderbei

Department of Ops. Res. and Fin. Eng., Princeton University, Princeton, NJ 08544.

Tel.: +609-258-2345

E-mail: rvdb@princeton.edu

representation, such as interior point methods (IPMs), are expected to have computational advantages for the KCS problem. We numerically demonstrate that KCS combined with IPMs is up to 10 times faster than vanilla IPMs and state-of-the-art methods such as $\ell_1\_\ell_s$ and Mirror Prox regardless of the sparsity level or problem size.

Keywords Linear programming · compressed sensing · parametric simplex method · sparse signals · interior-point methods

Mathematics Subject Classification (2000) MSC 65K05 · 62P99

1 Introduction and Contribution Overview

Compressed sensing (CS) aims to recover a sparse signal from a small number of measurements. The theoretical foundation of compressed sensing was first laid out by Donoho (2006) and Candès et al. (2006) and can be traced further back to the sparse recovery work of Donoho and Stark (1989); Donoho and Huo (2001); Donoho and Elad (2003). More recent progress in the area of compressed sensing is summarized in Elad (2010) and Kutyniok (2012).

Let $\mathbf{x}^* := (x_1^*, \dots, x_n^*)^T \in \mathbb{R}^n$ denote a signal to be recovered. We assume n is large and $\mathbf{x}^*$ is sparse (i.e., many entries of x^* are zero). Let $\mathbf{A}$ be a given (or chosen) $m \times n$ matrix with $m < n$. Let a_{ij} denote the $(i, j)^{\text{th}}$ element of $\mathbf{A}$. The *compressed sensing problem* aims to recover $\mathbf{x}^*$ from the compressed, noise-free signal $\mathbf{y} := \mathbf{A}\mathbf{x}^* \in \mathbb{R}^m$.

Specifically, we wish to find the sparsest solution to an underdetermined linear system by solving

$$\min_{\mathbf{x}} \|\mathbf{x}\|_0 \text{ subject to } \mathbf{A}\mathbf{x} = \mathbf{y}, \tag{P_0}$$

where $\|\mathbf{x}\|_0 := \sum_{i=1}^n \mathbf{1}(x_i \neq 0)$. This problem is NP-hard because of the nonconvexity of the ℓ_0 pseudo-norm. To handle this challenge, Chen et al. (1998) proposed the *basis pursuit* approach in which $\|\mathbf{x}\|_0$ is replaced by $\|\mathbf{x}\|_1 := \sum_{i=1}^n |x_i|$ to obtain a convex optimization problem

$$\min_{\mathbf{x}} \|\mathbf{x}\|_1 \text{ subject to } \mathbf{A}\mathbf{x} = \mathbf{y}. \tag{P_1}$$

Donoho and Elad (2003) and Cohen et al. (2009) provide conditions under which the solutions to Problems (P_0) and (P_1) are unique in their respective problems.

One key question is to understand what conditions guarantee that the solutions to (P_0) and (P_1) are equal. Various sufficient conditions have been discovered. For example, letting $\mathbf{A}_{*S}$

denote the submatrix of $\mathbf{A}$ with columns indexed by a subset $S \subset \{1, \ldots, n\}$, we say that $\mathbf{A}$ has the *k-restricted isometry property (k-RIP)* with constant δ_k if for any S with cardinality k,

$$(1 - \delta_k)\|\mathbf{v}\|_2^2 \leq \|\mathbf{A}_{*S}\mathbf{v}\|_2^2 \leq (1 + \delta_k)\|\mathbf{v}\|_2^2 \quad \text{for any } \mathbf{v} \in \mathbb{R}^k, \tag{1}$$

where $\|\mathbf{v}\|_2 := \sqrt{\sum_{j=1}^n v_j^2}$. We define $\delta_k(\mathbf{A})$ to be the smallest value of δ_k for which the matrix $\mathbf{A}$ has the k-RIP property. This property was first introduced by Candès et al. (2006). Under the assumption that $k := \|\mathbf{x}^*\|_0 \ll n$ and $\mathbf{A}$ satisfies the k-RIP condition, Cai and Zhang (2012) prove that whenever $\delta_k(\mathbf{A}) < 1/3$, the solutions to (P$_0$) and (P$_1$) are the same. The works in Donoho and Tanner (2005a,b, 2009) take a different approach by studying the convex geometric properties based on $\mathbf{A}$ instead of the RIP conditions to derive conditions when the solutions of (P$_0$) and (P$_1$) are the same.

Existing algorithms for solving the convex program (P$_1$) include interior-point methods (Candès et al., 2006; Kim et al., 2007), projected gradient methods (Figueiredo et al., 2008), first-order methods (Juditsky et al., 2014) and Bregman iterations (Yin et al., 2008). Besides algorithms that solve the convex program (P$_1$), several greedy algorithms have been proposed, including matching pursuit (Mallat and Zhang, 1993) and its many variants (Tropp, 2004; Donoho et al., 2006; Needell and Vershynin, 2009; Needell and Tropp, 2010; Donoho et al., 2009; Foucart, 2011). To achieve more scalability, combinatorial algorithms such as the so-called "Heavy Hitters on Steroids" (HHS) pursuit algorithm (Gilbert et al., 2007) and sub-linear Fourier transform (Iwen, 2010) have also been developed.

In this paper, we revisit the optimization aspects of the classical compressed sensing formulation (P$_1$) and one of its extensions, Kronecker compressed sensing (Duarte and Baraniuk, 2012). We consider two ideas for accelerating iterative algorithms. One reduces the total number of iterations, and the other reduces the computation required to do an iteration. We demonstrate the effectiveness of these ideas by thorough numerical simulations.

Our first idea, an optimization algorithm, is motivated by the fact that if the desired solution is very sparse, it should be reached after a relatively small number of simplex iterations starting from the zero vector. Such an idea motivates the usage of an optimization algorithm that can exploit the solution sparsity, e.g, the *parametric simplex method* (see, e.g., Vanderbei (2007); Dantzig (1998)). In this paper, we propose a customized parametric simplex method. Compared to standard simplex methods, our algorithm exploits a new pivot rule tailored for compressed

sensing problems. It is well known that slightly altering the pivot rules and formulation of the simplex method can result in a significant increase in computational speed (Forrest and Goldfarb, 1992; Pan, 2008).

While the simplex method has an exponential computational complexity in the worse case (Klee and Minty, 1972), we emphasize that the parametric simplex method is a suitable optimization method for our compressed sensing problem. This matches existing theoretical results on the simplex method that state the "average complexity" or "smoothed complexity" of the simplex method is polynomial. See Adler et al. (1987); Spielman and Teng (2004); Post and Ye (2015) and the references within.

Our second idea, a problem reformulation, requires the sensing matrix $\mathbf{A}$ to be the Kronecker product of two smaller matrices, $\mathbf{B}$ and $\mathbf{C}$. Since we are typically allowed to design the sensing matrix $\mathbf{A}$ ourselves, this requirement does not impose any practical limitations. This formulation results in a *Kronecker compressed sensing* (KCS) problem that has been considered before (Duarte and Baraniuk, 2012). In our paper, we reformulate the linear program to ensure the constraint matrix is very sparse so that the problem can be solved efficiently. The computational advantage of using sparse constraint matrices has been well established in the linear programming literature (Vanderbei, 1991; Lustig et al., 1991; Gill et al., 1991; Vanderbei, 1993). While most optimization research in compressed sensing focuses on creating customized algorithms, our approach uses existing algorithms but a sparser problem formulation to speed up computation. To the best of our knowledge, such an idea has not been exploited in the compressed sensing literature yet.

Theoretically, KCS involves a tradeoff between computation and statistics: it gains computational advantages (as will be shown in the numerical section) at the price of requiring more measurements (i.e., larger m). More specifically, using sub-Gaussian random sensing matrices (to be defined later), whenever $m = O(k^2 \log^2(\sqrt{n}/k))$, KCS recovers the true signal with probability at least $1 - 4\exp(-C\sqrt{m})$ for some constant $C > 0$. The sample requirement m for KCS is grows quadratically with k, as compared to the linear rate of $m = O(k \log(n/k))$ in standard compressed sensing. We provide more details in later sections.

The rest of this paper is organized as follows. In the next section, we describe how to solve the standard compressed sensing version (P_1) of the problem using the parametric simplex method. Then in Section 3 we describe the statistical foundation behind Kronecker compressed sensing (KCS). In Section 4 we present the sparse formulation of KCS that dramatically speeds

up existing optimization algorithms (such as interior point methods). In Section 5 we provide numerical comparisons against state-of-the-art methods to show the advantage of our methods. The problem formulations considered in Sections 2 through 5 involve no noise—the measurements are assumed to be exact. In Section 6 we conclude and discuss how our methods can be used to handle noisy or approximately-sparse settings of compressed sensing.

2 Compressed Sensing via the Parametric Simplex Method

Consider the following parametric perturbation to (P$_1$) for a specified μ:

$$\{\widehat{\mathbf{x}}, \widehat{\boldsymbol{\epsilon}}\} = \underset{\mathbf{x}, \boldsymbol{\epsilon}}{\operatorname{argmin}}\ \mu\|\mathbf{x}\|_1 + \|\boldsymbol{\epsilon}\|_1 \tag{P$_2$}$$

$$\text{subject to} \quad \mathbf{Ax} + \boldsymbol{\epsilon} = \mathbf{y},$$

where $\boldsymbol{\epsilon} \in \mathbb{R}^m$ denotes the residual. Since μ is a tuning parameter that affects the solution $\widehat{\mathbf{x}}$, we should denote the solution to (P$_2$) as $\widehat{\mathbf{x}}_\mu$, but for notational simplicity, we write $\widehat{\mathbf{x}}$ as shown in (P$_2$) instead. For large values of μ, the optimal solution is $\widehat{\mathbf{x}} = \mathbf{0}$ and $\widehat{\boldsymbol{\epsilon}} = \mathbf{y}$. For small, strictly-positive values of μ, the situation reverses: $\widehat{\boldsymbol{\epsilon}} = \mathbf{0}$ and we've solved the original problem (P$_1$). Belloni and Chernozhukov (2011) considered the above formulation, and they provide statistical guarantees on the solution for a particular magnitude of μ.

In this section, we require that the sensing matrix $\mathbf{A}$ satisfies a suitable k-RIP property specified by the following lemma.

Lemma 1 (Cai and Zhang (2012)) *For a small enough μ, let $\{\widehat{\mathbf{x}}, \mathbf{0}\}$ be the optimal solution of (P$_2$), and let $k = \|\mathbf{x}^*\|_0$ be the sparsity of vector $\mathbf{x}^*$. If $\delta_k(\mathbf{A}) < 1/3$, then $\widehat{\mathbf{x}} = \mathbf{x}^*$.*

Problem (P$_2$) is referred to as ℓ_1-penalized quantile regression. It can be solved by the parametric simplex method (Vanderbei, 2007; Dantzig, 1998), which is a homotopy method used for sensitivity and perturbation analysis. In particular, we start at a large value of μ and successively reduce it to form a solution path. That is, the solution to Problem (P$_2$) for a particular value of μ serves as the initialization to solve the same problem for a smaller value of μ. It can be shown that such a solution path is piecewise linear and all the transition points can be easily calculated. Algorithmically, the parametric simplex algorithm calculates the full solution path until we arrive at a value of μ for which the optimal solution has $\widehat{\boldsymbol{\epsilon}} = \mathbf{0}$, at which point we have solved the original problem (P$_1$). The advantage of the parametric simplex

method over the standard simplex method is that it solves the entire regularization path in terms of μ. Specifically, each basic solution produced by the algorithm is optimal for some interval of μ-values. Once the solution path indexed by μ is given, we could pick the desired μ as specified in Belloni and Chernozhukov (2011) to satisfy certain statistical properties. However, we focus on picking the first μ for which $\widehat{\boldsymbol{\epsilon}} = \mathbf{0}$. More details of the parametric simplex algorithm are illustrated in Figure 1.

Fig. 1 Illustration of the parametric simplex method. The horizontal line (black) corresponds to varying values of μ. We explicitly use superscripts to denote the iteration counter for clarity (the solution path segment between two consecutive transition points is called an iteration). The horizontal line is partitioned into a finite number of intervals such that each interval corresponds to a solution $\{\mathbf{x}, \boldsymbol{\epsilon}\}$ that is optimal for any value of μ within that interval. μ is initialized to be $\|\mathbf{A}\|_1$, which ensures the initialization $\{\mathbf{x}^{(0)}, \boldsymbol{\epsilon}^{(0)}\}$ is optimal. The algorithm decreases μ toward 0 until it reaches a solution $\{\mathbf{x}^{(T)}, \boldsymbol{\epsilon}^{(T)}\}$ where $\boldsymbol{\epsilon}^{(T)} = \mathbf{0}$. Then $\mathbf{x}^{(T)} = \widehat{\mathbf{x}}$ is our desired optimal solution to Problem (P$_1$). If the solution to (P$_1$) is unique, the interval corresponding to $\{\mathbf{x}^{(T)}, \boldsymbol{\epsilon}^{(T)}\}$ will contain $\mu = 0$. Since we obtain the entire solution path, other methods such as Belloni and Chernozhukov (2011) can be used to pick the solution $\mathbf{x}^{(T-1)}$ corresponding to μ^*, though it is not the focus of our paper.

To formulate (P$_2$) as a linear program, we reparametrize it using nonnegative variables and equality constraints. To this end, we split each variable into the difference between two nonnegative variables:

$$\mathbf{x} = \mathbf{x}^+ - \mathbf{x}^- \quad \text{and} \quad \boldsymbol{\epsilon} = \boldsymbol{\epsilon}^+ - \boldsymbol{\epsilon}^-,$$

where the entries of $\mathbf{x}^+, \mathbf{x}^-, \boldsymbol{\epsilon}^+, \boldsymbol{\epsilon}^-$ are all nonnegative. The next step is to replace $\|\mathbf{x}\|_1$ with $\mathbf{1}^T(\mathbf{x}^+ + \mathbf{x}^-)$ and to make a similar substitution for $\|\boldsymbol{\epsilon}\|_1$. In general, the sum $x_j^+ + x_j^-$ does not equal the absolute value $|x_j|$ but it is easy to see that equality holds at optimality. This is a well-known and standard technique for rewriting problems involving ℓ_1-norms as linear

programs. The resulting linear program becomes

$$\textbf{Simplex CS:} \qquad \min_{\mathbf{x}^+,\mathbf{x}^-,\boldsymbol{\epsilon}^+,\boldsymbol{\epsilon}^-} \mu \mathbf{1}^T(\mathbf{x}^+ + \mathbf{x}^-) + \mathbf{1}^T(\boldsymbol{\epsilon}^+ + \boldsymbol{\epsilon}^-) \qquad (\text{P}_3)$$

$$\text{subject to} \qquad \mathbf{A}(\mathbf{x}^+ - \mathbf{x}^-) + (\boldsymbol{\epsilon}^+ - \boldsymbol{\epsilon}^-) = \mathbf{y}$$

$$\mathbf{x}^+, \mathbf{x}^-, \boldsymbol{\epsilon}^+, \boldsymbol{\epsilon}^- \geq \mathbf{0}.$$

For μ large enough, the optimal solution must be $\mathbf{x}^+ = \mathbf{x}^- = \mathbf{0}$, and $\boldsymbol{\epsilon}^+ - \boldsymbol{\epsilon}^- = \mathbf{y}$. Given that our variables are required to be nonnegative, we initialize $\boldsymbol{\epsilon}^+$ and $\boldsymbol{\epsilon}^-$ according to

$$y_i > 0 \implies \epsilon_i^+ = y_i > 0 \text{ and } \epsilon_i^- = 0, \qquad (2)$$

$$y_i < 0 \implies \epsilon_i^- = -y_i > 0 \text{ and } \epsilon_i^+ = 0. \qquad (3)$$

The equality case can be decided either way. Such an initialization is feasible for arbitrary μ. Furthermore, declaring the nonzero variables to be *basic* variables and the zero variables to be *nonbasic*, we use this as our initial basic solution for the parametric simplex method. We note that this initial solution is optimal for $\mu \geq \|\mathbf{A}\|_1 := \max_{1 \leq j \leq n} \sum_{i=1}^m |a_{ij}|$, the largest column-wise ℓ_1 norm of $\mathbf{A}$. That is, for any solution $\{\mathbf{x}, \boldsymbol{\epsilon}\}$ such that $\mathbf{y} = \mathbf{A}\mathbf{x} + \boldsymbol{\epsilon}$, our initial solution setting $\mathbf{x} = \mathbf{0}$ and $\boldsymbol{\epsilon} = \mathbf{y}$ will be the global optima: $\|\mathbf{y}\|_1 = \|\mathbf{A}\mathbf{x} + \boldsymbol{\epsilon}\|_1 \leq \|\mathbf{A}\|_1\|\mathbf{x}\|_1 + \|\boldsymbol{\epsilon}\|_1 \leq \mu\|\mathbf{x}\|_1 + \|\boldsymbol{\epsilon}\|_1$. The pseudo-code to initialize and determine μ is presented in Algorithm 1, where Step 3 and Step 4 refer to the pivot step described in Chapters 7 and 8 of Vanderbei (2001).

Algorithm 1 Pseudo-Code for Parametric Simplex Method

Require: Inputs $\mathbf{A}$ and $\mathbf{y}$ as in Problem (P$_3$).

1: Set $\mu = \|\mathbf{A}\|_{1,1}$. Set initial optimal solution to be $\mathbf{x}^+ = \mathbf{x}^- = \mathbf{0}$ and $\boldsymbol{\epsilon}^+$ and $\boldsymbol{\epsilon}^-$ to follow (2) and (3) respectively, as illustrated in Figure 1.

2: **while** $\boldsymbol{\epsilon}^+ + \boldsymbol{\epsilon}^- \neq \mathbf{0}$ **do**

3:      Determine the smallest value of μ such that the current solution for (P$_3$) is optimal.

4:      For the current value of μ, apply a simplex pivot step to determine a new feasible solution.

5: **end while**

6: **return** The optimal solution to (P$_1$), $\widehat{\mathbf{x}} = \mathbf{x}^+ - \mathbf{x}^-$.

---

3 Kronecker Compressed Sensing Formulation

In this section, we introduce the Kronecker compressed sensing idea (Duarte and Baraniuk, 2012). While Duarte and Baraniuk (2012) discusses Kronecker sensing to handle multi-dimensional

signals, we show that representing one-dimensional signals as multi-dimensional signals will bear computational benefits. Unlike the previous section where we required the sensing matrix $\mathbf{A}$ to satisfy the $k-$RIP condition, in this section we impose a structural requirement on the sensing matrix. In particular, for given matrices $\mathbf{B}$ and $\mathbf{C}$, we consider the problem formulation that has the following representation:

$$\textbf{IPM CS:} \qquad \min \|\mathbf{x}\|_1 \ \text{ subject to } \ (\mathbf{B} \otimes \mathbf{C})\mathbf{x} = \mathbf{y}. \tag{P_4}$$

The matrices $\mathbf{B}$ and $\mathbf{C}$ are of size $m_1 \times n_1$ and $m_2 \times n_2$ respectively, and $\mathbf{A}$, our sensing matrix, is $(m_1 m_2) \times (n_1 n_2)$ and defined as the Kronecker product of $\mathbf{B}$ and $\mathbf{C}$:

$$\mathbf{A} := \mathbf{B} \otimes \mathbf{C} = \begin{bmatrix} \mathbf{C}b_{11} & \cdots & \mathbf{C}b_{1n_1} \\ \vdots & \ddots & \vdots \\ \mathbf{C}b_{m_1 1} & \cdots & \mathbf{C}b_{m_1 n_1} \end{bmatrix}.$$

This Kronecker structural assumption on the new sensing matrix $\mathbf{A}$ motivates a new matrix-based sensing strategy based on a left and right sensing matrix. To see this, assuming we want to recover a signal vector $\mathbf{x}^* \in \mathbb{R}^n$, we first reparameterize $\mathbf{x}^*$ into a matrix $\mathbf{X}^* \in \mathbb{R}^{n_2 \times n_1}$ by assigning each consecutive length-n_2 sub-vector of $\mathbf{x}^*$ to a column of $\mathbf{X}^*$. Here, without loss of generality, we assume $n = n_1 \times n_2$. Then, the observed matrix $\mathbf{Y}$ is given by $\mathbf{Y} := \mathbf{C}\mathbf{X}^*\mathbf{B}^T$. We define $\|\mathbf{X}\|_0 := \sum_{i,j} \mathbf{1}(x_{ij} \neq 0)$ and $\|\mathbf{X}\|_{1,1} := \sum_{i,j} |x_{ij}|$. Let the vec$(\cdot)$ operator take a matrix and concatenate its elements column-by-column to build one large column-vector containing all the elements of the matrix.

Given $\mathbf{Y} \in \mathbb{R}^{m_2 \times m_1}$ and the sensing matrices $\mathbf{B}$ and $\mathbf{C}$, we make one important observation. Recall the constraint in Problem (P_4). If $\mathbf{y} := \mathbf{A}\mathbf{x}^*$ where $\mathbf{x}^* = \text{vec}(\mathbf{X}^*)$, we have $\mathbf{y} = \text{vec}(\mathbf{Y})$. This means that Problem (P_4) is equivalent to

$$\min \|\mathbf{X}\|_{1,1} \ \text{ subject to } \ \mathbf{C}\mathbf{X}\mathbf{B}^T = \mathbf{Y}. \tag{P_5}$$

In other words, if $\widehat{\mathbf{x}}$ is the solution to (P_4) and $\widehat{\mathbf{X}}$ is the solution to (P_5), then $\widehat{\mathbf{x}} = \text{vec}(\widehat{\mathbf{X}})$. Hence, we can interpret our Kronecker compressed sensing scheme as either having a Kronecker structural assumption on $\mathbf{A}$ or having two separate (left and right) sensing matrices $\mathbf{B}$ and $\mathbf{C}$.

Recall the definition of RIP given in Equation (1). To understand the statistical theory of imposing the Kronecker structure on the sensing matrix $\mathbf{A}$, we use Lemma 2 from Duarte and Baraniuk (2012), which establishes the relationships between the k-RIP constants of $\delta_k(\mathbf{B})$, $\delta_k(\mathbf{C})$, and $\delta_k(\mathbf{A})$.

Lemma 2 (Duarte and Baraniuk (2012)) *Suppose* $\mathbf{A} = \mathbf{B} \otimes \mathbf{C}$. *Then*

$$1 + \delta_k(\mathbf{A}) \leq (1 + \delta_k(\mathbf{B}))(1 + \delta_k(\mathbf{C})). \tag{4}$$

In addition, we define a sub-Gaussian distribution as follows.

Definition 1 (Sub-Gaussian Distribution) We say a mean-zero random variable X follows a *sub-Gaussian distribution* if there exists some $\sigma \in \mathbb{R}_+$ such that

$$\mathbb{E} \exp(tX) \leq \exp\left(\frac{\sigma^2 t^2}{2}\right) \ \text{ for all } t \in \mathbb{R}.$$

It is clear that the Gaussian distribution with mean 0 and variance σ^2 satisfies the above definition. The next theorem provides sufficient conditions for perfect recovery of KCS.

Theorem 1 *Suppose the entries of matrices* $\mathbf{B}$ *and* $\mathbf{C}$ *follow a sub-Gaussian distribution with parameter* σ. *Then there exists a constant* $C > 0$ *(depending on* σ*) such that whenever*

$$m_1 \geq C \cdot k \log(n_1/k) \ \text{ and } \ m_2 \geq C \cdot k \log(n_2/k),$$

the convex program (P$_5$) *attains perfect recovery with probability*

$$\mathbb{P}\left(\widehat{\mathbf{X}} = \mathbf{X}^*\right) \geq 1 - \underbrace{\left(2e^{-\frac{m_1}{2C}} + 2e^{-\frac{m_2}{2C}}\right)}_{\rho(m_1, m_2)}.$$

Proof (for Theorem 1) We use the equivalence between Problem (P$_5$) and Problem (P$_4$). From Lemma 1 and Lemma 2, it suffices to show that $\delta_k(\mathbf{B})$ and $\delta_k(\mathbf{C})$ are both less than $2/\sqrt{3} - 1$. Let $\tau := 2/\sqrt{3} - 1$. From Theorem 9.2 of Foucart and Rauhut (2013), there exist constants C_1 and C_2 (depending on σ) such that if $m_1 \geq 2C_1 \tau^{-2} k \log(n_1/k)$ and $m_2 \geq 2C_2 \tau^{-2} k \log(n_2/k)$, then

$$
\begin{aligned}
\mathbb{P}\left(\delta_k(\mathbf{B}) < \frac{2}{\sqrt{3}} - 1 \text{ and } \delta_k(\mathbf{C}) < \frac{2}{\sqrt{3}} - 1\right) &= 1 - \left(\mathbb{P}\left(\delta_k(\mathbf{B}) \geq \tau\right) + \mathbb{P}\left(\delta_k(\mathbf{C}) \geq \tau\right)\right) \\
&\geq 1 - \left(2e^{-\frac{\tau^2 m_1}{2C_1}} + 2e^{-\frac{\tau^2 m_2}{2C_2}}\right) \\
&\geq 1 - \rho(m_1, m_2).
\end{aligned}
$$

$\square$

An analogous result can be derived from the results in Duarte and Baraniuk (2012) which uses the wavelet basis. From the above theorem, we see that for $m_1 = m_2 = \sqrt{m}$ and $n_1 = n_2 = \sqrt{n}$, whenever the number of measurements satisfies

$$m = O\left(k^2 \log^2\left(\frac{\sqrt{n}}{k}\right)\right), \tag{5}$$

we have $\widehat{\mathbf{X}} = \mathbf{X}^*$ with probability at least $1 - 4\exp(-C\sqrt{m})$ for some constant C.

Here we compare the above result to that of (standard) compressed sensing problem (P_1). Following the same argument as in Theorem 1, whenever

$$m = O\left(k \log\left(\frac{n}{k}\right)\right), \tag{6}$$

we have $\widehat{\mathbf{x}} = \mathbf{x}^*$ with probability at least $1 - 2\exp(-Cm)$. Comparing (6) to (5), we see that KCS needs more stringent conditions for perfect recovery. Specifically, for a fixed n, as k (the unknown sparsity level) increases, the required number of samples for KCS will grow quadratically with k rate as opposed to linearly. However, in the next section and in the numerical results, we will see that KCS enjoys a tremendous improvement in computation time. This will illustrate our tradeoff between computational performance and statistical recovery.

4 Sparsifying the Constraint Matrix for Efficient Computation

We can use standard LP algorithms such as an interior-point algorithm to solve (P_4). However, to achieve improved computational performance, we carefully formulate the problem so that our algorithm can explicitly exploit the Kronecker structure of $\mathbf{A}$. The key to efficiently solving the linear programming problem associated with the Kronecker sensing problem lies in noting that the dense, Kronecker product $\mathbf{A}$ can be factored into a product of two sparse matrices:

$$\mathbf{A} = \begin{bmatrix} \mathbf{C}b_{11} & \cdots & \mathbf{C}b_{1n_1} \\ \vdots & \ddots & \vdots \\ \mathbf{C}b_{m_11} & \cdots & \mathbf{C}b_{m_1n_1} \end{bmatrix} = \begin{bmatrix} \mathbf{C} & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{C} & \cdots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \cdots & \mathbf{C} \end{bmatrix} \begin{bmatrix} b_{11}\mathbf{I}_{n_2} & b_{12}\mathbf{I}_{n_2} & \cdots & b_{1n_1}\mathbf{I}_{n_2} \\ b_{21}\mathbf{I}_{n_2} & b_{22}\mathbf{I}_{n_2} & \cdots & b_{2n_1}\mathbf{I}_{n_2} \\ \vdots & \vdots & \ddots & \vdots \\ b_{m_11}\mathbf{I}_{n_2} & b_{m_11}\mathbf{I}_{n_2} & \cdots & b_{m_1n_1}\mathbf{I}_{n_2} \end{bmatrix} =: \mathbf{V}\mathbf{W},$$

where $\mathbf{I}_{n_2}$ denotes an $n_2 \times n_2$ identity matrix and $\mathbf{0}$ denotes an $m_2 \times m_2$ zero matrix. Notice that while the matrix $\mathbf{A}$ is usually completely dense, it is a product of two very sparse matrices: $\mathbf{V} = \mathbf{I}_{m_1} \otimes \mathbf{C} \in \mathbb{R}^{m \times m_1 n_2}$ and $\mathbf{W} = \mathbf{B} \otimes \mathbf{I}_{n_2} \in \mathbb{R}^{m_1 n_2 \times n}$. Hence, if we introduce new variables $\mathbf{z}$, we can rewrite (P_4) equivalently as

$$\min_{\mathbf{x}, \mathbf{z}} \ \|\mathbf{x}\|_1 \tag{P_6}$$

$$\text{subject to} \quad \mathbf{W}\mathbf{x} - \mathbf{z} = 0$$

$$\mathbf{V}\mathbf{z} = \mathbf{y}.$$

Using our previous matrix notation, we can rewrite (P_6) as

$$\textbf{IPM KCS:} \qquad \min_{\mathbf{X},\mathbf{Z}} \ \|\mathbf{X}\|_{1,1} \qquad\qquad\qquad\qquad (\mathrm{P}_7)$$

$$\text{subject to} \quad \mathbf{CX} - \mathbf{Z} = 0$$

$$\mathbf{ZB}^T = \mathbf{Y}.$$

If we want to use a parametric simplex method to solve (P_6), we can, as before, split $\mathbf{x}$ and $\boldsymbol{\epsilon}$ into a difference between their positive and negative parts and enforce equality constraints to convert the problem into a linear program:

$$\textbf{Simplex KCS:} \qquad \min_{\mathbf{x}^+,\mathbf{x}^-,\boldsymbol{\epsilon}^+,\boldsymbol{\epsilon}^-,\mathbf{z}} \ \mu\mathbf{1}^T(\mathbf{x}^+ + \mathbf{x}^-) + \mathbf{1}^T(\boldsymbol{\epsilon}^+ + \boldsymbol{\epsilon}^-) \qquad\qquad (\mathrm{P}_8)$$

$$\text{subject to} \quad \mathbf{z} - \mathbf{W}(\mathbf{x}^+ - \mathbf{x}^-) \qquad\qquad = 0$$

$$\mathbf{Vz} \qquad\qquad + \ (\boldsymbol{\epsilon}^+ - \boldsymbol{\epsilon}^-) \ = \mathbf{y}$$

$$\mathbf{x}^+, \mathbf{x}^-, \boldsymbol{\epsilon}^+, \boldsymbol{\epsilon}^- \geq 0.$$

This formulation (P_8) has more variables and more constraints than (P_3), but now the constraint matrix is very sparse. We reiterate that while (P_8) and (P_3) are mathematically equivalent when $\mathbf{A}$ is a Kronecker product, the reparameterization in (P_8) is more amendable to fast computational performance. Specifically, for linear programming, the sparsity of the constraint matrix is a significant contributor towards computational efficiency (Vanderbei, 1991; Gill et al., 1991). In fact, we can view the decomposition in (P_7) as a sparsification technique analogous to one step of the fast-Fourier optimization idea described in (Vanderbei, 2012).

5 Numerical Results and Comparisons

Before showing our simulation results, we briefly describe three other popular methods to solve compressed sensing problems: $\ell_1\text{-}\ell_s$, *Mirror Prox*, and *Fast Hard Thresholding Pursuit*. We will compare the performance of our method against these three methods on noiseless compressed sensing problems.

Proposed Methods: In this paper, we have presented two ideas, the parametric simplex method (an optimization algorithm) and Kronecker compressed sensing (a problem reformulation). We have made implementations that use these two ideas either separately or jointly. We use "KCS" to refer to optimization problems that explicitly exploit sparsity structure in the

formulation, (P_7) and (P_8). We modify the parametric simplex algorithm described in Vanderbei (2007) implemented in C found at `http://www.orfe.princeton.edu/~rvdb/LPbook/src/index.html`, to solve both (P_3) and (P_8). We refer to these implementations as "Simplex" and "Simplex KCS" respectively in our simulation results. We also use an interior-point solver called LOQO (Vanderbei (1999)) to solve (P_4) and (P_7). We refer to these implementations as "IPM" and "IPM KCS" respectively in our simulation results.

(Specialized) Interior-Point Method: The ℓ_1-ℓ_s method (Kim et al., 2007) is a truncated Newton interior-point method using a preconditioner to solve

$$\widehat{\mathbf{x}} = \operatorname*{argmin}_{\mathbf{x}} \|\mathbf{A}\mathbf{x} - \mathbf{y}\|_2^2 + \lambda\|\mathbf{x}\|_1 \tag{P_9}$$

for a given regularization parameter λ. Since computing the Newton direction is computationally prohibitive for large-scale problems due to its Hessian, ℓ_1-ℓ_s circumvents this problem by approximating the Newton direction using preconditioned conjugate gradients. In general, interior-point methods are very efficient because they use second-order information. The MATLAB code for this algorithm can be found at `http://stanford.edu/~boyd/l1_ls/`.

Greedy Method: Fast Hard Thresholding Pursuit (FHTP) (Foucart, 2011) is a greedy algorithm that alternates between two steps to approximately solve

$$\widehat{\mathbf{x}} = \operatorname*{argmin}_{\mathbf{x}:\|\mathbf{x}\|_0=k} \|\mathbf{x}\|_1 : \mathbf{A}\mathbf{x} = \mathbf{y} \tag{P_{10}}$$

for a given sparsity level k. In the first step, it chooses the best k coordinates of $\mathbf{x}$ according to a certain criterion, and in the next step it optimizes $\mathbf{x}$ for only thse k coordinates while setting the remaining coordinates to 0. This algorithm is appealing because of its simplicity and its exact recovery as long as $\mathbf{A}$ satisfies an RIP condition and k is correctly chosen. The MATLAB code for this algorithm can be found at `http://www.math.drexel.edu/~foucart/HTP.zip`.

First-Order Method: The Mirror Prox algorithm (Juditsky et al., 2014) is a first-order algorithm that solves

$$\widehat{\mathbf{x}} = \operatorname*{argmin}_{\mathbf{x}} \|\mathbf{x}\|_1 : \|\mathbf{A}\mathbf{x} - \mathbf{y}\|_2 \leq \delta \tag{P_{11}}$$

for a given tolerance δ by reformulating the problem as a saddle-point problem and using a proximal method to solve its variational inequality. First-order algorithms are favored in compressed sensing literature for their computational speed since they avoid calculating the Hessian matrix, and saddle point formulations are advantageous because they naturally combine the

primal and dual of the problem via variational inequalities. The MATLAB code for this algorithm can be found at `http://www2.isye.gatech.edu/~nemirovs/MirrorProxJan10_2012.zip`.

Experimental Protocol: In the rest of this section, let $\mathbf{x}^*$ denote the true signal and $\widehat{\mathbf{x}}$ denote the estimated signal using one of the above algorithms. We measure the accuracy of the solution by

$$\text{Relative } \ell_1 \text{ error: } \frac{\|\mathbf{x}^* - \widehat{\mathbf{x}}\|_1}{\|\mathbf{x}^*\|_1} \qquad \text{and} \qquad \ell_\infty \text{ error: } \|\mathbf{x}^* - \widehat{\mathbf{x}}\|_\infty.$$

We compare five different algorithms, the parametric simplex method, the interior point method, $\ell_1\text{-}\ell_s$, FHTP and Mirror Prox at different sparsity levels k. FHTP requires two different modes, oracle and agnostic. In the former, FHTP is given the true sparsity of $\mathbf{x}^*$. In the latter, FHTP is always given a sparsity of 100 regardless of the true $\mathbf{x}^*$.

Since each algorithm is solving a slightly different optimization problem, we devise a methodology for fair comparison. An important quantity to achieve fairness is the *imprecision*, $\|\mathbf{A}\widehat{\mathbf{x}} - \mathbf{y}\|_2^2$, the degree to which the solution $\widehat{\mathbf{x}}$ satisfies the constraints. We first apply our proposed methods (Simplex KCS, IPM KCS, Simplex CS, IPM CS) and record their imprecision. Since simplex methods are exact algorithms, we cannot reasonably expect the same magnitude of constraint error with $\ell_1\text{-}\ell_s$ (P_9) and Mirror Prox (P_{11}). Hence, we require these last two to have imprecision of only up to two magnitudes more than the simplex-based algorithms. For each k, we found that $\lambda = 0.01$ in (P_9) achieves a comparable imprecision. Given the solution to $\ell_1\text{-}\ell_s$, we can easily set δ in (P_{11}) to match the precision. The parametric simplex method and oracle FHTP naturally achieve the highest precision in most cases.

To ensure that each optimization algorithm is solving the same problem, we sample $\mathbf{A}$ and $\mathbf{B}$ as Gaussian matrices where each entry is a standard Gaussian (mean 0, standard deviation 1). "Simplex KCS" and "IPM KCS" use the matrices $\mathbf{V} = \mathbf{I} \otimes \mathbf{C}$ and $\mathbf{W} = \mathbf{B} \otimes \mathbf{I}$, while all the other methods use the sensing matrix $\mathbf{A} = \mathbf{B} \otimes \mathbf{C}$. In the following, we perform two different simulation sets. In the first simulation set, we vary the sparsity level of $\mathbf{x}$. In the second, we vary the length of $\mathbf{x}$. In either simulation set, we simulate 10 trials for each sparsity level or length. Instructions for downloading and running the various codes/algorithms described in this section can be found at `http://www.orfe.princeton.edu/~rvdb/tex/CTS/kronecker_sim.html`.

Results: (Varying Sparsity) In the first simulation set, we vary the true sparsity level by generating random problems using $m = 1{,}122 = 33 \times 34$ and $n = 20{,}022 = 141 \times 142$ and vary the number of nonzeros k in signal $\mathbf{x}^*$ from 2 to 150.

Table 1 provides time measured in seconds, the relative ℓ_1 error and the ℓ_∞ error averaged across 10 trials for each level of sparsity. All the simulation times are shown in Figure 2. There are two observations. First, when the true sparsity is very small ($k \leq 70$), the parametric simplex method without the Kronecker structure (Simplex) outperforms most modern methods in terms of precision and time. Second, we see that previously slow methods (Simplex and IPM) are tremendously sped up once the problem is formulated using the Kronecker structure (Simplex KCS and IPM KCS). There is roughly a ten-fold improvement in speed. In fact, IPM KCS is uniformly faster than $\ell_1\text{-}\ell_s$ and Mirror Prox. Our results show that if our sensing matrix has Kronecker structure and our algorithm is adapted to exploit the this property, simplex and interior point methods are highly competitive.

Fig. 2 Solution times for a large number of trials having $m = 1,122$, $n = 20,022$, and various degrees of sparsity in the underlying signal. The horizontal axis shows the number of nonzeros in the signal. The vertical axis gives a semi-log scale of solution times. The error bars have lengths equal to one standard deviation based on the multiple trials. Between Simplex KCS and IPM KCS, we outperform $\ell_1\text{-}\ell_s$ and Mirror Prox.

In four of the trials for Simplex throughout the entire simulation set, the solver erroneously reported that the solution was unbounded. We suspect this is due to numerical imprecision after the hundreds of thousands of simplex iterations for large values of k. This is not unexpected since we coded the entire parametric simplex method ourselves in C instead of using commercialized functions in MATLAB.

(Varying Size) In the second simulation set, we vary the problem size by fixing the number of nonzeros k in signal $\mathbf{x}^*$ to 100 and constructing random problems using $m = 1{,}122 = 33 \times 34$ and varying $n = 141 \times 142$ to $n = 141 \times 402$. Let n_2 denote the varying dimension. Table 2 shows the same attributes as in the previous table. As shown in Figure 3, the relative order of the seven algorithms mostly persists throughout the simulation set. In particular, in all cases tried, Simplex KCS outperforms $\ell_1\text{-}\ell_s$ and Mirror Prox. IPM KCS outperforms most methods throughout the entire simulation set.

6 Discussion and Conclusions

We revisit compressed sensing from an optimization perspective. We advance the field of compressed sensing in two ways.

First, despite having an exponential worst-case complexity, the parametric simplex method is competitive for very sparse signals. It outperforms $\ell_1\text{-}\ell_s$ and Mirror Prox under this regime (and we suspect many other methods) in both time and precision. Also, by adopting a parametric simplex method, we solve the problem for all values of μ, thereby finding the entire solution path in one shot. This feature of the parametric simplex method allows the user to pick a particular value of μ or take the largest μ for which $\boldsymbol{\epsilon} = 0$. Our paper focused on the latter.

Second, if we use the Kronecker structure, both the parametric simplex and interior point methods speed up ten-fold, making them competitive with other modern CS algorithms. But, as explained earlier, the Kronecker sensing problem involves changing the underlying problem being solved. The sensing matrix is now viewed as the Kronecker product of two sub-Gaussian matrices. While we presented this idea using only the parametric simplex and interior point methods, we expect this idea to benefit most optimization methods. This is left for future investigation. The Kronecker idea is inspired by the fast-Fourier transform where dense matrices are split into products of sparse matrices (Vanderbei (1991), Vanderbei (2012)). Hence, any op-

	Sparsity $k = 2$			Sparsity $k = 20$		
	Time (Sec)	Rel. ℓ_1 Error	ℓ_∞ Error	Time (Sec)	Rel. ℓ_1 Error	ℓ_∞ Error
Simplex KCS	1e-8	0	0	2.000	0	0
IPM KCS	44.650	3.71e-10	9.41e-8	48.000	1.98e-8	1.36e-6
Simplex	2.000	0	0	3.500	0	0
IPM	730.350	3.78e-10	9.73e-8	758.700	2.54e-10	9.23e-8
l1ls	110.120	6.02e-6	1.08e-5	96.753	9.68e-6	0.00010
Mirror Prox	28.070	0.00025	0.00099	82.740	0.00143	0.01535
FHTP (Oracle)	16.912	1.30e-11	2.60e-11	14.270	2.67e-5	8.79e-5
FHTP (Agnostic)	191.575	0.25993	4.67440	196.25500	0.17329	3.49615

	Sparsity $k = 50$			Sparsity $k = 70$		
	Time (Sec)	Rel. ℓ_1 Error	ℓ_∞ Error	Time (Sec)	Rel. ℓ_1 Error	ℓ_∞ Error
Simplex KCS	12.000	0	0	25.500	0	0
IPM KCS	47.700	3.41e-8	7.62e-6	48.950	3.67e-8	6.95e-6
Simplex	19.500	0	0	66.000	0	0
IPM	758.700	3.49e-8	7.62e-6	821.250	4.23e-8	2.65e-6
l1ls	170.315	1.45e-5	0.00031	267.155	1.84e-5	0.00052
Mirror Prox	42.840	0.00011	0.06293	64.030	0.00015	0.11449
FHTP (Oracle)	13.273	0.00020	0.00160	15.575	0.00017	0.00144
FHTP (Agnostic)	146.215	0.00504	1.773	35.64750	0.00518	1.42405

	Sparsity $k = 100$			Sparsity $k = 150$		
	Time (Sec)	Rel. ℓ_1 Error	ℓ_∞ Error	Time (Sec)	Rel. ℓ_1 Error	ℓ_∞ Error
Simplex KCS	70.500	0	0	462.500	4.50e-6	0.00017
IPM KCS	49.950	5.2e-7	2.31e-5	56.500	1.40e-5	0.00136
Simplex	234.500	0	0	1587.500	2.00e-6	0.00010
IPM	783.450	5.31e-7	0.00035	794.500	1.50e-5	0.00150
l1ls	377.315	2.43e-5	0.00104	789.165	9.79e-5	0.00683
Mirror Prox	410.050	0.00011	0.42348	635.085	0.00036	2.43170
FHTP (Oracle)	16.231	0.00040	0.00471	79.677	0.01460	128.01000
FHTP (Agnostic)	20.439	0.00040	0.00471	148.945	0.01646	145.07500

Table 1 This table shows the time (seconds), relative ℓ_1 error and relative ℓ_∞ error for 6 selected sparsity levels averaged (median) over 10 trials. If the median value is smaller than 1e-4, we write out at least two significant digits in scientific notation. We write "0" to denote exact recovery (achieved by only the simplex method). The first two rows of each table represent our proposed methods. Parametric simplex method outperforms other methods for very sparse problems $k \leq 70$. Naturally, FHTP (oracle) is the fastest for $k \leq 100$, but we see that both our methods outperform FHTP in relative ℓ_1 error and uniform error. By incorporating Kronecker structure, we see that previously slow methods can experience a drastic speed-up (i.e., the difference between IPM and IPM KCS).
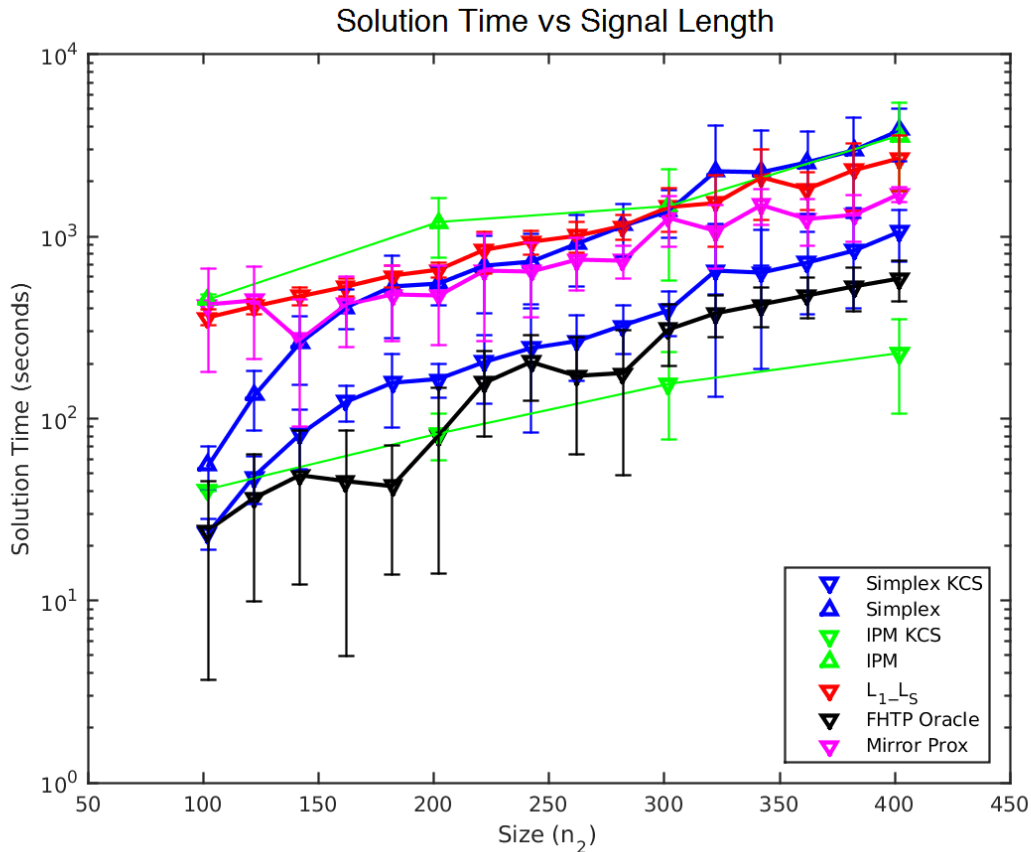
Fig. 3 Solution times for a large number of trials having $m = 1,122$ and $k = 100$. The number of dimensions for $\mathbf{x}^*$ in each trial is $n = 141 \times n_2$, and the horizontal axis shows the range of n_2. The vertical axis and error bars denote the same quantities as in Figure 2. IPM KCS outperforms most methods throughout the entire simulation set.

timization method that accommodates sparse matrix multiplication operations can potentially be altered to benefit further from a Kronecker compressed sensing scheme. The theoretical guarantees for using Kronecker compressed sensing are more stringent, however, which illustrates the tradeoff between computational efficiency and statistics.

In most applications, we expect some noise to corrupt the true signal. As long as the signal-to-noise ratio is high, we can adapt our methods to handle inexact constraints. Furthermore, the parametric simplex method computes solutions for all values of μ. As mentioned before, we can pick the solution associated with the value suggested in Belloni and Chernozhukov (2011) to achieve statistical properties for noisy cases. If we have a specific residual size $\|\boldsymbol{\epsilon}\|_1$ that we are willing to tolerate, we can pick the appropriate solution from the solution path. On the

	Size $n_2 = 102$			Size $n_2 = 202$		
	Time (Sec)	Rel. ℓ_1 Error	ℓ_∞ Error	Time (Sec)	Rel. ℓ_1 Error	ℓ_∞ Error
Simplex KCS	23.500	0	0	165.500	1.50e-6	1.50e-5
IPM KCS	40.732	3.54e-8	6.34e-8	92.351	2.56e-6	5.23e-5
Simplex	55.300	0	0	552.500	2.50e-6	2.50e-5
IPM	452.140	2.64e-8	4.34e-8	1409.206	6.35e-6	6.54e-5
l1ls	359.672	3.54e-5	1.21e-5	691.541	2.21e-5	1.62e-5
Mirror Prox	421.268	0.00985	0.00242	502.532	0.00983	0.00357
FHTP (Oracle)	24.327	0.30032	0.01964	115.465	0.51763	0.10540
	Size $n_2 = 302$			Size $n_2 = 402$		
	Time (Sec)	Rel. ℓ_1 Error	ℓ_∞ Error	Time (Sec)	Rel. ℓ_1 Error	ℓ_∞ Error
Simplex KCS	652.100	0.20000	0.00200	1750.100	0.50000	0.01023
IPM KCS	185.532	0.01729	0.40000	306.290	0.00966	0.40001
Simplex	2283.200	0.01063	0.01065	5555.700	0.01594	0.01597
IPM	3143.280	0.09238	0.01829	6503.130	0.12620	0.01810
l1ls	1541.290	0.20003	0.01731	2575.930	0.60002	0.09930
Mirror Prox	1043.022	0.21061	0.02327	1682.610	0.60639	0.12080
FHTP (Oracle)	382.556	1.16788	0.47230	727.805	1.19667	0.87309

Table 2 This table, similar format as Table 1, shows 4 selected dimension sizes n_2 averaged (median) over 10 trials. Simplex KCS and IPM KCS outperform most methods in terms of time and accuracy. The parametric simplex (without Kronecker structure) performs better than ℓ_1-ℓ_s for $k = 100$ even as n_2 grows. Note that the ℓ_∞ error for $n_2 = 402$ is large for ℓ_1-ℓ_s, Mirror Prox and FHTP, meaning at least one dimension is drastically incorrect.

other hand, if one were interested settings where $\mathbf{x}$ is approximately sparse and we wanted to recover the s-largest entries of $\mathbf{x}$, one could use either our "Simplex" or "IPM KCS" method. The relation between the solution and the true signal can then be determined through existing theory in Candès (2008). In future work, we plan to investigate noisy and approximately sparse settings more thoroughly and extend the proposed method to the setting of 1-bit compressed sensing.

Acknowledgment

References

ADLER, I., KARP, R. M. and SHAMIR, R. (1987). A simplex variant solving an $m \times d$ linear program in $O(\min(m_2, d_2))$ expected number of pivot steps. *Journal of Complexity* **3** 372–387.

BELLONI, A. and CHERNOZHUKOV, V. (2011). ℓ_1-penalized quantile regression in high-dimensional sparse models. *Annals of Statistics* **39** 82–130.

CAI, T. T. and ZHANG, A. (2012). Sharp RIP bound for sparse signal and low-rank matrix recovery. *Applied and Computational Harmonic Analysis* .

CANDÈS, E., ROMBERG, J. and TAO, T. (2006). Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE Transactions on Information Theory* **52** 489–509.

CANDÈS, E. J. (2008). The restricted isometry property and its implications for compressed sensing. *Comptes Rendus Mathematique* **346** 589–592.

CHEN, S. S., DONOHO, D. L. and SAUNDERS, M. A. (1998). Atomic decomposition by basis pursuit. *SIAM Journal on Scientific Computing* **20** 33–61.

COHEN, A., DAHMEN, W. and DEVORE, R. (2009). Compressed sensing and best k-term approximation. *J. Amer. Math. Soc* 211–231.

DANTZIG, G. B. (1998). *Linear Programming and Extensions.* Princeton University Press.

DONOHO, D. L. (2006). Compressed sensing. *IEEE Transactions on Information Theory* **52** 1289–1306.

DONOHO, D. L. and ELAD, M. (2003). Optimally sparse representation in general (nonorthogonal) dictionaries via ℓ_1-minimization. *Proc. Natl. Acad. Sci. USA* **100**.

DONOHO, D. L., ELAD, M. and TEMLYAKOV, V. N. (2006). Stable recovery of sparse overcomplete representations in the presence of noise. *IEEE Transactions on Information Theory* **52** 6–18.

DONOHO, D. L. and HUO, X. (2001). Uncertainty principles and ideal atomic decomposition. *IEEE Transactions on Information Theory* **47** 2845–2862.

DONOHO, D. L., MALEKI, A. and MONTANARI, A. (2009). Message passing algorithms for compressed sensing. *Proc. Natl. Acad. Sci. USA* **106** 18914–18919.

DONOHO, D. L. and STARK, P. B. (1989). Uncertainty principles and signal recovery. *SIAM J. Appl. Math.* **49** 906–931.

DONOHO, D. L. and TANNER, J. (2005a). Neighborliness of randomly projected simplices in high dimensions. *Proceedings of the National Academy of Sciences* **102** 9452–9457.

DONOHO, D. L. and TANNER, J. (2005b). Sparse nonnegative solutions of underdetermined linear equations by linear programming. *Proceedings of the National Academy of Sciences* 9446–9451.

DONOHO, D. L. and TANNER, J. (2009). Observed universality of phase transitions in high-dimensional geometry, with implications for modern data analysis and signal processing. *Philos. Trans. Roy. Soc. S.-A* **367** 4273–4293.

DUARTE, M. F. and BARANIUK, R. G. (2012). Kronecker compressive sensing. *IEEE Transactions on Image Processing* **21** 494–504.

ELAD, M. (2010). *Sparse and Redundant Representations - From Theory to Applications in Signal and Image Processing.* Springer.

FIGUEIREDO, M., NOWAK, R. and WRIGHT, S. (2008). Gradient projection for sparse reconstruction: Application to compressed sensing and other inverse problems. *IEEE Journal of Selected Topics in Signal Processing* **1** 586–597.

FORREST, J. J. and GOLDFARB, D. (1992). Steepest-edge simplex algorithms for linear programming. *Mathematical programming* **57** 341–374.

FOUCART, S. (2011). Hard thresholding pursuit: An algorithm for compressive sensing. *SIAM Journal on Numerical Analysis* **49** 2543–2563.

FOUCART, S. and RAUHUT, H. (2013). *A Mathematical Introduction to Compressive Sensing.* Springer.

GILBERT, A. C., STRAUSS, M. J., TROPP, J. A. and VERSHYNIN, R. (2007). One sketch for all: Fast algorithms for compressed sensing. In *STOC* (D. S. Johnson and U. Feige, eds.). ACM.

GILL, P. E., MURRAY, W., PONCELEON, D. B. and SAUNDERS, M. A. (1991). Solving reduced kkt systems in barrier methods for linear and quadratic programming. Tech. rep., DTIC Document.

IWEN, M. A. (2010). Combinatorial sublinear-time Fourier algorithms. *Foundations of Computational Mathematics* **10** 303–338.

JUDITSKY, A., KARZAN, F. K. and NEMIROVSKI, A. (2014). ℓ_1 minimization via randomized first order algorithms. Tech. rep., Université Joseph Fourier.

KIM, S., KOH, K., LUSTIG, M., BOYD, S. and GORINEVSKY, D. (2007). An interior-point method for large-scale l_1-regularized least squares. *IEEE Transactions on Selected Topics in Signal Processing* **1** 606–617.

KLEE, V. and MINTY, G. J. (1972). How good is the simplex method? *Inequalities-III* 159–175.

KUTYNIOK, G. (2012). Compressed sensing: Theory and applications. *CoRR* **abs/1203.3815**.

LUSTIG, I. J., MULVEY, J. M. and CARPENTER, T. J. (1991). Formulating two-stage stochastic programs for interior point methods. *Operations Research* **39** 757–770.

MALLAT, S. and ZHANG, Z. (1993). Matching pursuits with time-frequency dictionaries. *Signal Processing, IEEE Transactions on* **41** 3397–3415.

NEEDELL, D. and TROPP, J. A. (2010). CoSaMP: Iterative signal recovery from incomplete and inaccurate samples. *Commun. ACM* **53** 93–100.

NEEDELL, D. and VERSHYNIN, R. (2009). Uniform uncertainty principle and signal recovery via regularized orthogonal matching pursuit. *Foundations of Computational Mathematics* **9** 317–334.

PAN, P.-Q. (2008). A largest-distance pivot rule for the simplex algorithm. *European Journal of Operational Research* **187** 393–402.

POST, I. and YE, Y. (2015). The simplex method is strongly polynomial for deterministic markov decision processes. *Mathematics of Operations Research* **40** 859–868.

SPIELMAN, D. A. and TENG, S.-H. (2004). Smoothed analysis of algorithms: Why the simplex algorithm usually takes polynomial time. *Journal of the ACM (JACM)* **51** 385–463.

TROPP, J. A. (2004). Greed is good: Algorithmic results for sparse approximation. *IEEE Transactions on Information Theory* **50** 2231–2242.

VANDERBEI, R. (1991). Splitting dense columns in sparse linear systems. *Lin. Alg. and Appl.* **152** 107–117.

VANDERBEI, R. (1999). LOQO: An interior point code for quadratic programming. *Optimization Methods and Software* **12** 451–484.

VANDERBEI, R. (2007). *Linear Programming: Foundations and Extensions.* 3rd ed. Springer.

VANDERBEI, R. J. (1993). Alpo: Another linear program optimizer. *ORSA Journal on Computing* **5** 134–146.

VANDERBEI, R. J. (2001). Linear programming. *Foundations and extensions, International Series in Operations Research & Management Science* **37**.

VANDERBEI, R. J. (2012). Fast Fourier optimization. *Math. Prog. Comp.* **4** 1–17.

YIN, W., OSHER, S., GOLDFARB, D. and DARBON, J. (2008). Bregman iterative algorithms for ℓ_1-minimization with applications to compressed sensing. *SIAM J. Img. Sci.* **1** 143–168.