

11-12-2015

Convergence Properties of the Randomized Extended Gauss-Seidel and Kaczmarz Methods

Anna Ma

Claremont Graduate University

Deanna Needell

Claremont McKenna College

Aaditya Ramdas

University of California - Berkeley

Recommended Citation

A. Ma, D. Needell and A. Ramdas. "Convergence Properties of the Randomized Extended Gauss-Seidel and Kaczmarz Methods." *SIAM Journal on Matrix Analysis and Applications*, vol. 36, iss. 4, 1590–1604, 2015.

This Article - postprint is brought to you for free and open access by the CMC Faculty Scholarship at Scholarship @ Claremont. It has been accepted for inclusion in CMC Faculty Publications and Research by an authorized administrator of Scholarship @ Claremont. For more information, please contact scholarship@cuc.claremont.edu.

CONVERGENCE PROPERTIES OF THE RANDOMIZED EXTENDED GAUSS-SEIDEL AND KACZMARZ METHODS

ANNA MA, DEANNA NEEDELL, AADITYA RAMDAS

Abstract. The Kaczmarz and Gauss-Seidel methods both solve a linear system $\mathbf{X}\boldsymbol{\beta} = \mathbf{y}$ by iteratively refining the solution estimate. Recent interest in these methods has been sparked by a proof of Strohmer and Vershynin which shows the *randomized* Kaczmarz method converges linearly in expectation to the solution. Lewis and Leventhal then proved a similar result for the randomized Gauss-Seidel algorithm. However, the behavior of both methods depends heavily on whether the system is under or overdetermined, and whether it is consistent or not. Here we provide a unified theory of both methods, their variants for these different settings, and draw connections between both approaches. In doing so, we also provide a proof that an extended version of randomized Gauss-Seidel converges linearly to the least norm solution in the underdetermined case (where the usual randomized Gauss Seidel fails to converge). We detail analytically and empirically the convergence properties of both methods and their extended variants in all possible system settings. With this result, a complete and rigorous theory of both methods is furnished.

1. Introduction. We consider solving a linear system of equations

$$\mathbf{X}\boldsymbol{\beta} = \mathbf{y}, \quad (1.1)$$

for a (real or complex) $m \times n$ matrix \mathbf{X} , in various problem settings. Recent interest in the topic was reignited when Strohmer and Vershynin [26] proved the linear¹ convergence rate of the Randomized Kaczmarz (RK) algorithm that works on the rows of \mathbf{X} (data points). Following that, Leventhal and Lewis [15] proved the linear convergence of a Randomized Gauss-Seidel (RGS), i.e. Randomized Coordinate Descent, algorithm that works on the columns of \mathbf{X} (features).

When the system of equations is inconsistent (i.e. has no exact solution), as is typically the case when $m > n$ in real-world overconstrained systems, RK is known to not converge to the ordinary least squares solution

$$\boldsymbol{\beta}_{LS} := \arg \min_{\boldsymbol{\beta}} \frac{1}{2} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 \quad (1.2)$$

as studied by Needell [18]. Zouzias and Freris [30] extended the RK method with the modified Randomized Extended Kaczmarz (REK) algorithm, which linearly converges to $\boldsymbol{\beta}_{LS}$. Interestingly, in this setting, we will argue in Section 3.3 that RGS does converge to $\boldsymbol{\beta}_{LS}$ without any special extensions.

1.1. Motivation and contribution. The above introduction represents only half the story. When $m < n$, there are fewer constraints than variables, and the system has infinitely many solutions. In this case, especially if we have no prior reason to believe any additional sparsity in the signal structure, we are often interested in finding the least Euclidean norm solution:

$$\boldsymbol{\beta}_{LN} := \arg \min_{\boldsymbol{\beta}} \|\boldsymbol{\beta}\|_2 \quad \text{s.t.} \quad \mathbf{y} = \mathbf{X}\boldsymbol{\beta}. \quad (1.3)$$

While RGS converges to $\boldsymbol{\beta}_{LS}$ in the overcomplete setting, we shall argue in Section 3.3 that in the undercomplete setting it does not converge to $\boldsymbol{\beta}_{LN}$. We will also argue that RK does converge to $\boldsymbol{\beta}_{LN}$ without any extensions in this setting.

¹Mathematicians often refer to linear convergence as exponential convergence.

The main contribution of our paper is to provide a unified theory of these related iterative methods. We will also construct an extension to RGS that parallels REK, which unlike RGS does converge to β_{LN} (just as REK, unlike RK, converges to β_{LS}). Some desired properties for this algorithm include that it should *also* converge linearly, not require much extra computation, and work well in simulations. We shall see that our Randomized Extended Gauss-Seidel (REGS) method does indeed possess these desired properties. A summary of this unified theory is provided in Table 1.1.

Method	Overconstrained, consistent : convergence to β^* ?	Overconstrained, inconsistent : convergence to β_{LS} ?	Underconstrained : convergence to β_{LN} ?
RK	Yes [26]	No [18, Thm. 2.1]	Yes (Sec. 3.3)
REK	Yes [30]	Yes [30]	Yes (Sec. 3.3)
RGS	Yes [15]	Yes [15]	No (Sec. 3.3)
REGS	Yes (Remark 1)	Yes (Sec. 4.3)	Yes (Thm. 4.1)

TABLE 1.1

Summary of convergence properties for the overdetermined and consistent setting, overdetermined and inconsistent setting, and underdetermined settings. We write β^ to denote the solution to (1.1) in the overdetermined consistent setting, with β_{LS} and β_{LN} being defined in (1.2) and (1.3) for the other two settings.*

1.2. Paper Outline. In Section 2 we recap the three main existing algorithms mentioned in the introduction (RK, RGS, REK). We discuss the performance of these algorithms in the three natural settings described in Table 1.1 in Section 3. Section 4 introduces our proposed algorithm (REGS) and proves its linear convergence to the least norm solution, completing the theoretical framework. Lastly, we end with some simulation experiments in Section 5 to demonstrate the tightness and usefulness of our theory, and conclude in Section 6.

2. Existing Algorithms and Related Work. In this section, we will summarize the algorithms mentioned in the introduction, i.e. RK, RGS and REK. We will describe their iterative update rules and mention their convergence guarantees, leaving the details of convergence to the next section. Throughout the paper we will use the notation \mathbf{X}^i to represent the i th row of \mathbf{X} (or i th entry in the case of a vector) and $\mathbf{X}_{(j)}$ to denote the j th column of a matrix \mathbf{X} . We will write the estimation β as a column vector. We write vectors and matrices in boldface, and constants in standard font.

2.1. Randomized Kaczmarz (RK). The Kaczmarz method was first introduced in the notable work of Kaczmarz [14]. It has gained recent interest in tomography research where it is known as the *Algebraic Reconstruction Technique* (ART) [8, 17, 1, 13]. Although in its original form the method selects rows in a deterministic fashion (often simply cyclically), it has been well observed that a random selection scheme reduces the possibility of a poor choice of row ordering [9, 12]. Earlier convergence analysis of the randomized variant were obtained (e.g. [29]), but yielded bounds with expressions that were difficult to evaluate. Strohmer and Vershynin [26] showed that the RK method described above has an expected linear convergence rate to the solution β^* of (1.1), and are the first to provide an explicit convergence rate in expectation which depends only on the geometric properties of the system. This work was extended by Needell [18] to the inconsistent case, analyzed almost surely

by Chen and Powell [3], accelerated in several ways [7, 6, 24, 21, 20], and extended to more general settings [15, 25, 19].

We describe here the randomized variant of the Kaczmarz method put forth by Strohmer and Vershynin [26]. Taking \mathbf{X}, \mathbf{y} as input and starting from an arbitrary initial estimate for $\boldsymbol{\beta}$ (for example $\boldsymbol{\beta}_0 = \mathbf{0}$), RK repeats the following in each iteration. First, a random row $i \in \{1, \dots, m\}$ is selected with probability proportional to its Euclidean norm, i.e.

$$\Pr(\text{row} = i) = \frac{\|\mathbf{X}^i\|_2^2}{\|\mathbf{X}\|_F^2},$$

where $\|\mathbf{X}\|_F$ denotes the Frobenius norm of \mathbf{X} . Then, project the current iterate onto that row, i.e.

$$\boldsymbol{\beta}_{t+1} := \boldsymbol{\beta}_t + \frac{(y^i - \mathbf{X}^i \boldsymbol{\beta}_t)}{\|\mathbf{X}^i\|_2^2} (\mathbf{X}^i)^*, \quad (2.1)$$

where here and throughout \mathbf{X}^* denotes the (conjugate) transpose of \mathbf{X} .

Intuitively, this update can be seen as greedily satisfying the i th equation in the linear system. Indeed, it is easy to see that after the update,

$$\mathbf{X}^i \boldsymbol{\beta}_{t+1} = y^i. \quad (2.2)$$

Referring to (1.2) and defining

$$L(\boldsymbol{\beta}) = \frac{1}{2} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|^2 = \frac{1}{2} \sum_{i=1}^m (y^i - \mathbf{X}^i \boldsymbol{\beta})^2,$$

we can alternatively interpret this update as stochastic gradient descent (choosing a random data-point on which to update), where the step size is the inverse Lipschitz constant of the stochastic gradient

$$\nabla^2 \frac{1}{2} (y^i - \mathbf{X}^i \boldsymbol{\beta})^2 = \|\mathbf{X}^i\|_2^2.$$

2.2. Randomized Extended Kaczmarz (REK). For inconsistent systems, the RK method does not converge to the least-squares solution as one might desire. This fact is clear since the method at each iteration projects completely onto a selected solution space, being unable to break the so-called *convergence horizon*. One approach to overcome this is to use relaxation parameters, so that the estimates are not projected completely onto the subspace at each iteration [28, 27, 2, 10]. Recently, Zouzias and Freris [30] proposed a variant of the RK method motivated by the work of Popa [23] which instead includes a random projection to iteratively reduce the component of \mathbf{y} which is orthogonal to the range of \mathbf{X} . This method, named Randomized Extended Kaczmarz (REK) can be described by the following iteration updates, which can be initialized with $\boldsymbol{\beta}_0 = \mathbf{0}$ and $\mathbf{z}_0 = \mathbf{y}$:

$$\boldsymbol{\beta}_{t+1} := \boldsymbol{\beta}_t + \frac{(y^i - z_t^i - \mathbf{X}^i \boldsymbol{\beta}_t)}{\|\mathbf{X}^i\|_2^2} (\mathbf{X}^i)^*, \quad \mathbf{z}_{t+1} = \mathbf{z}_t - \frac{\langle \mathbf{X}_{(j)}, \mathbf{z}_t \rangle}{\|\mathbf{X}_{(j)}\|_2^2} \mathbf{X}_{(j)}. \quad (2.3)$$

Here, a column $j \in \{1, \dots, n\}$ is also selected at random with probability proportional to its Euclidean norm:

$$\Pr(\text{column} = j) = \frac{\|\mathbf{X}_{(j)}\|_2^2}{\|\mathbf{X}\|_F^2}, \quad (2.4)$$

and again $\mathbf{X}_{(j)}$ denotes the j th column of \mathbf{X} . Here, \mathbf{z}_t approximates the component of \mathbf{y} which is orthogonal to the range of \mathbf{X} , allowing for the iterates β_t to converge to the true least-squares solution of the system. Zouzias and Freris [30] prove that REK converges linearly in expectation to this solution β_{LS} .

2.3. Randomized Gauss-Seidel (RGS). Again taking \mathbf{X}, \mathbf{y} as input and starting from an arbitrary β_0 , the Randomized Gauss-Seidel (RGS) method (or the Randomized Coordinate Descent method) repeats the following in each iteration. First, a random column $j \in \{1, \dots, n\}$ is selected as in (2.4). We then minimize the objective $L(\beta) = \frac{1}{2} \|\mathbf{y} - \mathbf{X}\beta\|_2^2$ with respect to this coordinate to get

$$\beta_{t+1} := \beta_t + \frac{\mathbf{X}_{(j)}^*(\mathbf{y} - \mathbf{X}\beta_t)}{\|\mathbf{X}_{(j)}\|_2^2} \mathbf{e}_{(j)} \quad (2.5)$$

where $\mathbf{e}_{(j)}$ is the j th coordinate basis column vector (all zeros with a 1 in the j th position). It can be seen as greedily minimizing the objective with respect to the j th coordinate. Indeed, letting $\mathbf{X}_{(-j)}, \beta^{-j}$ represent \mathbf{X} without its j th column and β without its j th coordinate,

$$\frac{\partial L}{\partial \beta^j} = -\mathbf{X}_{(j)}^*(\mathbf{y} - \mathbf{X}\beta) = -\mathbf{X}_{(j)}^*(\mathbf{y} - \mathbf{X}_{(-j)}\beta^{-j} - \mathbf{X}_{(j)}\beta^j). \quad (2.6)$$

Setting this equal to zero for the coordinate-wise minimization, we get the aforementioned update (2.5) for β^j . Alternatively, since $[\nabla L(\beta)]^j = -\mathbf{X}_{(j)}^*(\mathbf{y} - \mathbf{X}\beta)$, the above update can intuitively be seen as a univariate descent step where the step size is the inverse Lipschitz constant of the gradient along the j th coordinate, since the (j, j) entry of the Hessian is

$$[\nabla^2 L(\beta)]_{j,j} = (\mathbf{X}^* \mathbf{X})_{j,j} = \|\mathbf{X}_{(j)}\|_2^2.$$

Leventhal and Lewis [15] showed that this algorithm has an expected linear convergence rate. We will detail the convergence properties of this algorithm and the others in the next section.

3. Problem Variations. We first examine the differences in behavior of the two algorithms RGS and RK in three distinct but related settings. This will highlight the opposite behaviors of these two similar algorithms.

When the system of equations (1.1) has a unique solution, we represent this by β^* . This happens when $m \geq n$, and the system is consistent. Assuming that \mathbf{X} has full column rank,

$$\beta^* = (\mathbf{X}^* \mathbf{X})^{-1} \mathbf{X}^* \mathbf{y}, \quad (3.1)$$

and then $\mathbf{X}\beta^* = \mathbf{y}$.

When (1.1) does not have any consistent solution, we refer to the least-squares solution of (1.2) as β_{LS} . This could happen in the overconstrained case, when $m > n$. Again, assuming that \mathbf{X} has full column rank, we have

$$\beta_{LS} = (\mathbf{X}^* \mathbf{X})^{-1} \mathbf{X}^* \mathbf{y}, \quad (3.2)$$

and we can write $\mathbf{r} := \mathbf{y} - \mathbf{X}\beta_{LS}$ as the residual vector.

When (1.1) has infinitely many solutions, we call the minimum Euclidean norm solution given by (1.3), β_{LN} . This could happen in the underconstrained case, when $m < n$. Assuming that \mathbf{X} has full row rank, we have

$$\beta_{LN} = \mathbf{X}^*(\mathbf{X}\mathbf{X}^*)^{-1}\mathbf{y}. \quad (3.3)$$

In the above notation, the *LS* stands for Least Squares and *LN* for Least Norm. We shall return to each of these three situations in that order in future sections.

One of our main contributions is to achieve a unified understanding of the behavior of RK and RGS in these different situations. The literature for RK deals mainly with the first two settings only (see [26], [18], [30]). In the third setting, one readily obtains convergence to an *arbitrary* solution (see e.g. (3) of [16]), but the convergence to the least norm solution is not often studied (likely for practical reasons). The literature for RGS typically focuses on more general setups than our specific quadratic least squares loss function $L(\beta)$ (see Nesterov [22] or Richtárik and Takáč [25]). However, for both the purposes of completeness, and for a more thorough understanding of the relationship between RK and RGS, it turns out to be crucial to analyze all three settings (for equations (1.1)-(1.3)).

1. When β^* is a unique consistent solution, we present proofs of the linear convergence of both algorithms - the results are known from papers by [26] and [15] but are presented here in a novel manner so that their relationship becomes clearer and direct comparison is easily possible.
2. When β_{LS} is the (inconsistent) least squares solution, we show why RGS iterates converge linearly to β_{LS} , but RK iterates do not - making RGS preferable. These facts are not hard to see, but we make it more intuitively and mathematically clear why this should be the case.
3. When β_{LN} is the minimum norm consistent solution, we explain why RK converges linearly to it, but RGS iterates do not (both so far seemingly undocumented observations) - making RK preferable.

Together, the above three points complete the picture (with solid accompanying intuition) of the opposing behavior of RK and RGS. Later, we will present our variant of the RGS method, the Randomized Extended Gauss-Seidel (REGS), and compare with the corresponding variant of RK (REK). This new analysis will complete the unified framework for these methods.

3.1. Overconstrained System, Consistent. Here we will assume that $m > n$, \mathbf{X} has full column rank, and the system is consistent, so $\mathbf{y} = \mathbf{X}\beta^*$. First, let us write the updates used by both algorithms in a revealing fashion. If RK and RGS select row i and column j at step $t + 1$, and \mathbf{e}^i (resp. $\mathbf{e}_{(j)}$) is the i th coordinate basis row (resp. column) vector, then the updates can be rewritten as:

$$\text{(RK)} \quad \beta_{t+1} := \beta_t + \frac{\mathbf{e}^i \mathbf{r}_t}{\|\mathbf{X}^i\|_2^2} (\mathbf{X}^i)^* \quad (3.4)$$

$$\text{(RGS)} \quad \beta_{t+1} := \beta_t + \frac{\mathbf{X}_{(j)}^* \mathbf{r}_t}{\|\mathbf{X}_{(j)}\|_2^2} \mathbf{e}_j \quad (3.5)$$

where $\mathbf{r}_t = \mathbf{y} - \mathbf{X}\boldsymbol{\beta}_t = \mathbf{X}\boldsymbol{\beta}^* - \mathbf{X}\boldsymbol{\beta}_t$ is the residual vector at iteration t . Then multiplying both equations by \mathbf{X} gives

$$(RK) \quad \mathbf{X}\boldsymbol{\beta}_{t+1} := \mathbf{X}\boldsymbol{\beta}_t + \frac{\mathbf{X}^i(\boldsymbol{\beta}^* - \boldsymbol{\beta}_t)}{\|\mathbf{X}^i\|_2^2} \mathbf{X}(\mathbf{X}^i)^* \quad (3.6)$$

$$(RGS) \quad \mathbf{X}\boldsymbol{\beta}_{t+1} := \mathbf{X}\boldsymbol{\beta}_t + \frac{\mathbf{X}_{(j)}^* \mathbf{X}(\boldsymbol{\beta}^* - \boldsymbol{\beta}_t)}{\|\mathbf{X}_{(j)}\|_2^2} \mathbf{X}_{(j)}. \quad (3.7)$$

We now come to an important difference, which is the key update equation for RK and RGS.

First, from the update (3.4) for RK, we have that $\boldsymbol{\beta}_{t+1} - \boldsymbol{\beta}_t$ is parallel to \mathbf{X}^i . Also, $\boldsymbol{\beta}_{t+1} - \boldsymbol{\beta}^*$ is orthogonal to \mathbf{X}^i (since $\mathbf{X}^i(\boldsymbol{\beta}_{t+1} - \boldsymbol{\beta}^*) = \mathbf{y}^i - \mathbf{y}^i = 0$). Then by the Pythagorean theorem,

$$\|\boldsymbol{\beta}_{t+1} - \boldsymbol{\beta}^*\|_2^2 = \|\boldsymbol{\beta}_t - \boldsymbol{\beta}^*\|_2^2 - \|\boldsymbol{\beta}_{t+1} - \boldsymbol{\beta}_t\|_2^2. \quad (3.8)$$

Note that from the update (3.7), we have that $\mathbf{X}\boldsymbol{\beta}_{t+1} - \mathbf{X}\boldsymbol{\beta}_t$ is parallel to $\mathbf{X}_{(j)}$. Also, $\mathbf{X}\boldsymbol{\beta}_{t+1} - \mathbf{X}\boldsymbol{\beta}^*$ is orthogonal to $\mathbf{X}_{(j)}$ (since $\mathbf{X}_{(j)}^*(\mathbf{X}\boldsymbol{\beta}_{t+1} - \mathbf{X}\boldsymbol{\beta}^*) = \mathbf{X}_{(j)}^*(\mathbf{X}\boldsymbol{\beta}_{t+1} - \mathbf{y}) = 0$ by the optimality condition $\partial L / \partial \boldsymbol{\beta}^j = 0$). Then again by the Pythagorean theorem,

$$\|\mathbf{X}\boldsymbol{\beta}_{t+1} - \mathbf{X}\boldsymbol{\beta}^*\|_2^2 = \|\mathbf{X}\boldsymbol{\beta}_t - \mathbf{X}\boldsymbol{\beta}^*\|_2^2 - \|\mathbf{X}\boldsymbol{\beta}_{t+1} - \mathbf{X}\boldsymbol{\beta}_t\|_2^2. \quad (3.9)$$

The rest of the proof follows by simply substituting for the last term in the above two equations, and is presented in the following table for easy comparison. Note $\boldsymbol{\Sigma} = \mathbf{X}^* \mathbf{X}$ is the full-rank covariance matrix and we first take expectations with respect to the randomness at the $(t+1)$ st step, conditioning on all randomness up to the t th step. We later iterate this expectation.

Randomized Kaczmarz:	Randomized Gauss-Seidel:
$\mathbb{E}_t \ \boldsymbol{\beta}_{t+1} - \boldsymbol{\beta}^*\ _2^2$	$\mathbb{E}_t \ \mathbf{X}\boldsymbol{\beta}_{t+1} - \mathbf{X}\boldsymbol{\beta}^*\ _2^2$
$= \ \boldsymbol{\beta}_t - \boldsymbol{\beta}^*\ _2^2 - \mathbb{E} \ \boldsymbol{\beta}_{t+1} - \boldsymbol{\beta}_t\ _2^2$	$= \ \mathbf{X}\boldsymbol{\beta}_t - \mathbf{X}\boldsymbol{\beta}^*\ _2^2 - \mathbb{E} \ \mathbf{X}\boldsymbol{\beta}_{t+1} - \mathbf{X}\boldsymbol{\beta}_t\ _2^2$
$= \ \boldsymbol{\beta}_t - \boldsymbol{\beta}^*\ _2^2$	$= \ \mathbf{X}\boldsymbol{\beta}_t - \mathbf{X}\boldsymbol{\beta}^*\ _2^2$
$- \sum_i \frac{\ \mathbf{X}^i\ _2^2}{\ \mathbf{X}\ _F^2} \frac{(\mathbf{X}^i(\boldsymbol{\beta}_t - \boldsymbol{\beta}^*))^2}{(\ \mathbf{X}^i\ _2^2)^2} \ \mathbf{X}^i\ _2^2$	$- \sum_j \frac{\ \mathbf{X}_{(j)}\ _2^2}{\ \mathbf{X}\ _F^2} \frac{(\mathbf{X}_{(j)}^* \mathbf{X}(\boldsymbol{\beta}_t - \boldsymbol{\beta}^*))^2}{(\ \mathbf{X}_{(j)}\ _2^2)^2} \ \mathbf{X}_{(j)}\ _2^2$
$= \ \boldsymbol{\beta}_t - \boldsymbol{\beta}^*\ _2^2 \left(1 - \frac{\ \mathbf{X}(\boldsymbol{\beta}_t - \boldsymbol{\beta}^*)\ _2^2}{\ \mathbf{X}\ _F^2 \ \boldsymbol{\beta}_t - \boldsymbol{\beta}^*\ _2^2} \right)$	$= \ \mathbf{X}\boldsymbol{\beta}_t - \mathbf{X}\boldsymbol{\beta}^*\ _2^2 \left(1 - \frac{\ \mathbf{X}^* \mathbf{X}(\boldsymbol{\beta}_t - \boldsymbol{\beta}^*)\ _2^2}{\ \mathbf{X}\ _F^2 \ \mathbf{X}\boldsymbol{\beta}_t - \mathbf{X}\boldsymbol{\beta}^*\ _2^2} \right)$
$\leq \ \boldsymbol{\beta}_t - \boldsymbol{\beta}^*\ _2^2 \left(1 - \frac{\lambda_{\min}(\boldsymbol{\Sigma})}{Tr(\boldsymbol{\Sigma})} \right)$	$\leq \ \mathbf{X}\boldsymbol{\beta}_t - \mathbf{X}\boldsymbol{\beta}^*\ _2^2 \left(1 - \frac{\lambda_{\min}(\boldsymbol{\Sigma})}{Tr(\boldsymbol{\Sigma})} \right)$

Here, $\lambda_{\min}(\boldsymbol{\Sigma}) \|\boldsymbol{\beta}_t - \boldsymbol{\beta}^*\|_2^2 \leq \|\mathbf{X}(\boldsymbol{\beta}_t - \boldsymbol{\beta}^*)\|_2^2$ i.e. $\lambda_{\min}(\boldsymbol{\Sigma})$ is the smallest eigenvalue of $\boldsymbol{\Sigma}$ (singular value of \mathbf{X}). It follows that

$$(RK) \quad \mathbb{E} \|\boldsymbol{\beta}_t - \boldsymbol{\beta}^*\|_2^2 \leq \left(1 - \frac{\lambda_{\min}(\boldsymbol{\Sigma})}{Tr(\boldsymbol{\Sigma})} \right)^t \|\boldsymbol{\beta}_0 - \boldsymbol{\beta}^*\|_2^2 \quad (3.10)$$

$$(RGS) \quad \mathbb{E} \|\boldsymbol{\beta}_t - \boldsymbol{\beta}^*\|_{\boldsymbol{\Sigma}}^2 \leq \left(1 - \frac{\lambda_{\min}(\boldsymbol{\Sigma})}{Tr(\boldsymbol{\Sigma})} \right)^t \|\boldsymbol{\beta}_0 - \boldsymbol{\beta}^*\|_{\boldsymbol{\Sigma}}^2,$$

where $\|\mathbf{w}\|_{\boldsymbol{\Sigma}}^2 = \mathbf{w}^* \boldsymbol{\Sigma} \mathbf{w} = \|\mathbf{X}\mathbf{w}\|_2^2$ is the norm induced by $\boldsymbol{\Sigma}$. Since $\boldsymbol{\Sigma}$ is invertible when $m > n$ and \mathbf{X} has full column rank, the last equation also implies linear convergence

of $\mathbb{E}\|\beta_t - \beta^*\|_2^2$. The final results exist in Strohmer and Vershynin [26], Leventhal and Lewis [15] but there is utility in seeing the two proofs in a form that differs from their original presentation, side by side. In this setting, both RK and RGS are essentially equivalent (without computational considerations).

3.2. Overconstrained System, Inconsistent. Here, we will assume that $m > n$, \mathbf{X} is full column rank, and the system is inconsistent, so $\mathbf{y} = \mathbf{X}\beta_{LS} + \mathbf{r}$, where \mathbf{r} is such that $\mathbf{X}^*\mathbf{r} = 0$. It is easy to see this condition, because as mentioned earlier,

$$\beta_{LS} = (\mathbf{X}^*\mathbf{X})^{-1}\mathbf{X}^*\mathbf{y},$$

implying that $\mathbf{X}^*\mathbf{X}\beta_{LS} = \mathbf{X}^*\mathbf{y}$. Substituting $\mathbf{y} = \mathbf{X}\beta_{LS} + \mathbf{r}$ gives that $\mathbf{X}^*\mathbf{r} = 0$.

In this setting, RK is known to not converge to the least squares solution, as is easily verified experimentally and geometrically. The tightest convergence upper bounds known are by [18] and [30] who show that

$$\begin{aligned} \mathbb{E}\|\beta_t - \beta_{LS}\|_2^2 &\leq \left(1 - \frac{\lambda_{\min}(\mathbf{\Sigma})}{\text{Tr}(\mathbf{\Sigma})}\right)^t \|\beta_0 - \beta_{LS}\|_2^2 + \frac{\|\mathbf{r}\|_2^2}{\lambda_{\min}(\mathbf{\Sigma})} \\ &= \left(1 - \frac{\sigma_{\min}^2(\mathbf{X})}{\|\mathbf{X}\|_F^2}\right)^t \|\beta_0 - \beta_{LS}\|_2^2 + \frac{\|\mathbf{r}\|_2^2}{\sigma_{\min}^2(\mathbf{X})}, \end{aligned}$$

where we write $\sigma_{\min}(\mathbf{X})$ to denote the smallest (non-zero) singular value of \mathbf{X} and again $\|\mathbf{X}\|_F$ its Frobenius norm. Attempting the previous proof, (3.8) no longer holds – the Pythagorean theorem fails because $\beta_{t+1} - \beta_{LS}$ is no longer orthogonal to \mathbf{X}^i since $\mathbf{X}^i(\beta_{t+1} - \beta_{LS}) = y^i - \mathbf{X}^i\beta_{LS} \neq 0$. Intuitively, the reason RK does not converge is that every update of RK (say of row i) is a projection onto the “wrong” hyperplane that has constant y^i (where the “right” hyperplane would involve projecting onto a parallel hyperplane with constant $y^i - r^i$ where \mathbf{r} was defined above). An alternative intuition is that all RK updates are in the span of the rows, but β_{LS} is not in the row span. These intuitive explanations are easily confirmed by experiments seen in [30, pp. 787–788], [18, pp. 402]. Zouzias and Freris [30] alleviate this issue with the REK algorithm, whose convergence obeys

$$\mathbb{E}\|\beta_t - \beta_{LS}\|_2^2 \leq \left(1 - \frac{\sigma_{\min}^2(\mathbf{X})}{\|\mathbf{X}\|_F^2}\right)^{\lfloor t/2 \rfloor} \left(1 + 2\frac{\sigma_{\min}^2(\mathbf{X})}{\sigma_{\max}^2(\mathbf{X})}\|\beta_{LS}\|_2^2\right). \quad (3.11)$$

However, the fate of RK doesn’t hold for RGS. Almost magically, in the previous proof, the Pythagorean theorem still holds in equation (3.9) because

$$\mathbf{X}_{(j)}^*(\mathbf{X}\beta_{t+1} - \mathbf{X}\beta_{LS}) = \mathbf{X}_{(j)}^*(\mathbf{X}\beta_{t+1} - \mathbf{y}) + \mathbf{X}_{(j)}^*(\mathbf{y} - \mathbf{X}\beta_{LS}) = 0. \quad (3.12)$$

The first term is 0 by the optimality condition for β_{t+1} i.e. $\mathbf{X}_{(j)}^*(\mathbf{X}\beta_{t+1} - \mathbf{y}) = \partial L / \partial \beta^j = 0$. The second term is zero by the global optimality of β_{LS} , i.e. $\mathbf{X}^*(\mathbf{y} - \mathbf{X}\beta_{LS}) = \nabla L = 0$. Also, $\mathbf{\Sigma}$ is full rank as before. Indeed, RGS works in the space of fitted values $\mathbf{X}\beta$ and not the iterates β .

In summary, RK does not converge to the LS solution, but RGS does at the same linear rate. This is what motivated the development of Randomized Extended Kaczmarz (REK) by Zouzias and Freris [30] which, as discussed earlier, is a modification of RK designed to converge to β_{LS} by randomly projecting out r . An independent paper by Dumitrescu [5] argues however that in this setting RGS is preferable to REK in terms of computational convergence.

3.3. Underconstrained System, Infinite Solutions. Here, we will assume that $m < n$, \mathbf{X} is full row rank and the system is consistent with infinitely many solutions. As mentioned earlier, it is easy to show that

$$\beta_{LN} = \mathbf{X}^*(\mathbf{X}\mathbf{X}^*)^{-1}\mathbf{y}$$

(which clearly satisfies $\mathbf{X}\beta_{LN} = \mathbf{y}$). Every other consistent solution can be expressed as

$$\beta = \beta_{LN} + \mathbf{z} \quad \text{where} \quad \mathbf{X}\mathbf{z} = 0.$$

Clearly any such β would also satisfy $\mathbf{X}\beta = \mathbf{X}\beta_{LN} = \mathbf{y}$. Since $\mathbf{X}\mathbf{z} = 0$, $\mathbf{z} \perp \beta_{LN}$ implying $\|\beta\|^2 = \|\beta_{LN}\|^2 + \|\mathbf{z}\|^2$, showing that β_{LN} is indeed the minimum norm solution as claimed.

In this case, RK has good behavior, and starting from $\beta_0 = 0$, it does converge linearly to β_{LN} . Intuitively, $\beta_{LN} = \mathbf{X}^*\alpha$ (for $\alpha = (\mathbf{X}\mathbf{X}^*)^{-1}\mathbf{y}$) and hence is in the row span of \mathbf{X} . Starting from $\beta_0 = 0$, RK only adds multiples of rows to its iterates, and hence will never have any component orthogonal to the row span of \mathbf{X} . There is exactly one solution with no component orthogonal to the row span of \mathbf{X} , and that is β_{LN} , and hence RK converges linearly to the required point, where the rate can be bounded in exactly the same way as (3.10). It is important not to start from an arbitrary β_0 since the RK updates can never eliminate any component of β_0 that is perpendicular to the row span of \mathbf{X} . Of course, the same properties are shared by REK for this case as well. It is noted in Zouzias and Freris [30, Sec. 2.1] that the REK converges at the same rate for underdetermined systems as it does overdetermined systems.

Mathematically, the previous earlier proof works because the Pythagorean theorem holds since it is a consistent system. Now, Σ is not full rank but note that since both β_{LN} and β_t are in the row span, $\beta_t - \beta_{LN}$ has no component orthogonal to \mathbf{X} (unless it equals zero in which case the algorithm has already converged). Hence $\lambda_{\min}(\Sigma)\|\beta_t - \beta_{LN}\|^2 \leq \|\mathbf{X}(\beta_t - \beta_{LN})\|^2$ holds, where $\lambda_{\min}(\Sigma)$ is now understood to be the smallest positive eigenvalue of Σ . To summarize, the exact same bound (3.10) still holds in this case, with the appropriate understanding of $\lambda_{\min}(\Sigma)$ and under the assumption that the initialized β_0 is in the row span of \mathbf{X} .

RGS unfortunately suffers the opposite fate. The iterates do not converge to β_{LN} , even though $\mathbf{X}\beta_t$ does converge to $\mathbf{X}\beta_{LN}$. Mathematically, the convergence proof still carries forward as before, but in the last step when $\mathbf{X}^*\mathbf{X}$ cannot be inverted because it is not full rank. Hence we get convergence of the residual to zero, without getting convergence of the iterates to the least norm solution. Intuitively, the iterates of RGS add components to the estimates that are orthogonal to the row span of \mathbf{X} . These components are never eliminated because in minimizing the residual, they are ignored. Therefore, RGS is able to minimize the residual without finding the least norm solution.

Unfortunately, when each update is cheaper for RK than RGS (due to matrix size), RGS is preferred for reasons of convergence and when it is cheaper for RGS than RK, RK is preferred.

4. REGS. We next introduce an extension of RGS, analogous to the extension REK of RK. The purpose of extending RK was to allow for convergence to the least squares solution. Now, the purpose of extending RGS is to allow for convergence to the least norm solution. We view this method as a completion to the unified analysis of these approaches, and it may also possess advantages in its own right.

4.1. The algorithm. Consider the linear system (1.1) with $m < n$. Let β_{LN} denote the least norm solution of the underdetermined system as described in (1.3). The REGS algorithm is described by the following pseudo-code. Analogous to the role z plays in REK, z iteratively approximates the component in β orthogonal to the row-span of \mathbf{X} . By iteratively removing this component, we converge to the least norm solution. Note that outputting β_t instead of $\beta_t^{LN} = \beta_t - z_t$ in Algorithm 1 recovers the RGS algorithm. This may be preferable in the overdetermined setting.

Algorithm 1 Randomized Extended Gauss-Seidel (REGS)

```

1: procedure ( $\mathbf{X}$ ,  $\mathbf{y}$ , maxIter)  $\triangleright m \times n$  matrix  $\mathbf{X}$ ,  $\mathbf{y} \in \mathbb{C}^m$ , maximum iterations  $T$ 
2:   Initialize  $\beta_0 = \mathbf{0}$ ,  $z_0 = \mathbf{0}$ 
3:   for  $t = 1, 2, \dots, T$  do
4:     Choose column  $\mathbf{X}_{(j)}$  with probability  $\frac{\|\mathbf{X}_{(j)}\|_2^2}{\|\mathbf{X}\|_F^2}$ 
5:     Choose row  $\mathbf{X}^i$  with probability  $\frac{\|\mathbf{X}^i\|_2^2}{\|\mathbf{X}\|_F^2}$ 
6:     Set  $\gamma_t = \frac{\mathbf{X}_{(j)}^*(\mathbf{X}\beta_{t-1} - \mathbf{y})}{\|\mathbf{X}_{(j)}\|_2^2} \mathbf{e}_{(j)}$ 
7:     Set  $\beta_t = \beta_{t-1} + \gamma_t$ 
8:     Set  $\mathbf{P}_i = \mathbf{I}d_n - \frac{(\mathbf{X}^i)^* \mathbf{X}^i}{\|\mathbf{X}^i\|_2^2} \triangleright \mathbf{I}d_n$  denotes the  $n \times n$  identity matrix
9:     Update  $z_t = \mathbf{P}_i(z_{t-1} + \gamma_t)$ 
10:    Update  $\beta_t^{LN} = \beta_t - z_t$ 
11:  end for
12:  Output  $\beta_t^{LN}$ 
13: end procedure

```

4.2. Main result. Our main result for the REGS method shows linear convergence to the least norm solution.

THEOREM 4.1. *The REGS algorithm outputs an estimate β_T^{LN} such that*

$$\mathbb{E}\|\beta_T^{LN} - \beta_{LN}\|_2^2 \leq \alpha^T \|\beta^{LN}\|_2^2 + 2\alpha^{\lfloor T/2 \rfloor} \frac{B}{1 - \alpha} \quad (4.1)$$

where $B = \frac{\|\mathbf{X}\beta_{LN}\|_2^2}{\|\mathbf{X}\|_F^2}$ and $\alpha = \left(1 - \frac{\sigma_{m \wedge n}^2(\mathbf{X})}{\|\mathbf{X}\|_F^2}\right)$.

Proof: We devote the remainder of this section to the proof of Theorem 4.1.

Let \mathbb{E}_{t-1} denote the expected value conditional on the first $t-1$ iterations, and in-state the notation of the theorem. That is, $\mathbb{E}_{t-1}[\cdot] = \mathbb{E}[\cdot | i_1, j_1, i_2, j_2, \dots, i_{t-1}, j_{t-1}]$ where i_{t^*} is the t^{th} row chosen and j_{t^*} is the t^{th} column chosen. We denote conditional expectation with respect to the choice of column as $\mathbb{E}_{t-1}^j[\cdot] = \mathbb{E}[\cdot | i_1, j_1, \dots, i_{t-1}, j_{t-1}, i_t]$. Similarly, we denote conditional expectation with respect to the choice of row as $\mathbb{E}_{t-1}^i[\cdot] = \mathbb{E}[\cdot | i_1, j_1, \dots, i_{t-1}, j_{t-1}, j_t]$. Then note by the law of total expectation we have that $\mathbb{E}_{t-1}[\cdot] = \mathbb{E}_{t-1}^i[\mathbb{E}_{t-1}^j[\cdot]]$. We will use the following elementary facts and lemmas.

FACT 1. ([30, Fact 3]) *For any \mathbf{P}_i as in the algorithm, $\mathbb{E}\|\mathbf{P}_i \mathbf{w}\|_2^2 \leq \alpha \|\mathbf{w}\|_2^2$ for any \mathbf{w} .*

LEMMA 4.2. ([15, Thm. 3.6]) *We have that*

$$\mathbb{E}_{t-1} \|\mathbf{X}\beta_t - \mathbf{X}\beta_{LN}\|_2^2 \leq \alpha \|\mathbf{X}\beta_{t-1} - \mathbf{X}\beta_{LN}\|_2^2$$

and that

$$\mathbb{E}\|\mathbf{X}\boldsymbol{\beta}_t - \mathbf{X}\boldsymbol{\beta}_{LN}\|_2^2 \leq \alpha^t \|\mathbf{X}\boldsymbol{\beta}_0 - \mathbf{X}\boldsymbol{\beta}_{LN}\|_2^2.$$

Now we first consider $\|\boldsymbol{\beta}_t^{LN} - \boldsymbol{\beta}_{LN}\|_2^2$:

$$\begin{aligned} \|\boldsymbol{\beta}_t^{LN} - \boldsymbol{\beta}_{LN}\|_2^2 &= \|\boldsymbol{\beta}_t - \mathbf{z}_t - \boldsymbol{\beta}_{LN}\|_2^2 \\ &= \|\boldsymbol{\beta}_t - \mathbf{P}_i(\mathbf{z}_{t-1} + \boldsymbol{\gamma}_t) - \mathbf{P}_i\boldsymbol{\beta}_{LN} - (\mathbf{Id}_n - \mathbf{P}_i)\boldsymbol{\beta}_{LN}\|_2^2 \\ &= \|\boldsymbol{\beta}_t - \mathbf{P}_i(\mathbf{z}_{t-1} + \boldsymbol{\beta}_t - \boldsymbol{\beta}_{t-1}) - \mathbf{P}_i\boldsymbol{\beta}_{LN} - (\mathbf{Id}_n - \mathbf{P}_i)\boldsymbol{\beta}_{LN}\|_2^2 \\ &= \|(\mathbf{Id}_n - \mathbf{P}_i)\boldsymbol{\beta}_t + \mathbf{P}_i(\boldsymbol{\beta}_{t-1} - \mathbf{z}_{t-1}) - \mathbf{P}_i\boldsymbol{\beta}_{LN} - (\mathbf{Id}_n - \mathbf{P}_i)\boldsymbol{\beta}_{LN}\|_2^2 \\ &= \|(\mathbf{Id}_n - \mathbf{P}_i)\boldsymbol{\beta}_t + \mathbf{P}_i\boldsymbol{\beta}_{t-1}^{LN} - \mathbf{P}_i\boldsymbol{\beta}_{LN} - (\mathbf{Id}_n - \mathbf{P}_i)\boldsymbol{\beta}_{LN}\|_2^2 \\ &= \|\mathbf{P}_i(\boldsymbol{\beta}_{t-1}^{LN} - \boldsymbol{\beta}_{LN}) + (\mathbf{Id}_n - \mathbf{P}_i)(\boldsymbol{\beta}_t - \boldsymbol{\beta}_{LN})\|_2^2 \\ &= \|\mathbf{P}_i(\boldsymbol{\beta}_{t-1}^{LN} - \boldsymbol{\beta}_{LN})\|_2^2 + \|(\mathbf{Id}_n - \mathbf{P}_i)(\boldsymbol{\beta}_t - \boldsymbol{\beta}_{LN})\|_2^2. \end{aligned} \quad (4.2)$$

So far, we have only used substitution of variables as defined for the algorithm and that $\boldsymbol{\beta}_{LN} = \mathbf{P}_i\boldsymbol{\beta}_{LN} + (\mathbf{Id}_n - \mathbf{P}_i)\boldsymbol{\beta}_{LN}$ is an orthogonal decomposition. We first focus on the expected value of the second term.

LEMMA 4.3. *We also have that*

$$\mathbb{E}_{t-1}\|(\mathbf{Id}_n - \mathbf{P}_i)(\boldsymbol{\beta}_t - \boldsymbol{\beta}_{LN})\|_2^2 \leq \frac{\alpha\|\mathbf{X}(\boldsymbol{\beta}_{t-1} - \boldsymbol{\beta}_{LN})\|_2^2}{\|\mathbf{X}\|_F^2}.$$

Proof:

$$\begin{aligned} &\mathbb{E}_{t-1}\|(\mathbf{Id}_n - \mathbf{P}_i)(\boldsymbol{\beta}_t - \boldsymbol{\beta}_{LN})\|_2^2 \\ &= \mathbb{E}_{t-1}[(\boldsymbol{\beta}_t - \boldsymbol{\beta}_{LN})^*(\mathbf{Id}_n - \mathbf{P}_i)^*(\mathbf{Id}_n - \mathbf{P}_i)(\boldsymbol{\beta}_t - \boldsymbol{\beta}_{LN})] \\ &= \mathbb{E}_{t-1}[(\boldsymbol{\beta}_t - \boldsymbol{\beta}_{LN})^*(\mathbf{Id}_n - \mathbf{P}_i)(\boldsymbol{\beta}_t - \boldsymbol{\beta}_{LN})] \\ &= \mathbb{E}_{t-1}\left[(\boldsymbol{\beta}_t - \boldsymbol{\beta}_{LN})^* \left(\frac{\mathbf{X}^i \mathbf{X}^i}{\|\mathbf{X}^i\|_2^2}\right) (\boldsymbol{\beta}_t - \boldsymbol{\beta}_{LN})\right] \\ &= \mathbb{E}_{t-1}\left[\frac{\|\mathbf{X}^i(\boldsymbol{\beta}_t - \boldsymbol{\beta}_{LN})\|_2^2}{\|\mathbf{X}^i\|_2^2}\right] \\ &= \mathbb{E}_{t-1}^j\left[\mathbb{E}_{t-1}^i\frac{\|\mathbf{X}^i(\boldsymbol{\beta}_t - \boldsymbol{\beta}_{LN})\|_2^2}{\|\mathbf{X}^i\|_2^2}\right] \\ &= \mathbb{E}_{t-1}^j\left[\sum_{i=1}^m \frac{\|\mathbf{X}^i(\boldsymbol{\beta}_t - \boldsymbol{\beta}_{LN})\|_2^2}{\|\mathbf{X}^i\|_2^2} \cdot \frac{\|\mathbf{X}^i\|_2^2}{\|\mathbf{X}\|_F^2}\right] \\ &= \mathbb{E}_{t-1}^j\left[\frac{\|\mathbf{X}(\boldsymbol{\beta}_t - \boldsymbol{\beta}_{LN})\|_2^2}{\|\mathbf{X}\|_F^2}\right] \\ &\leq \frac{\alpha\|\mathbf{X}(\boldsymbol{\beta}_{t-1} - \boldsymbol{\beta}_{LN})\|_2^2}{\|\mathbf{X}\|_F^2}. \end{aligned}$$

The first line follows by expanding the norm, the second line since $(\mathbf{Id}_n - \mathbf{P}_i)$ is a projection matrix, the third line from the definition of \mathbf{P}_i , the fourth line is

computation, the fifth line follows from the law of total expectation, the next two lines are computation, and finally the last line follows by Lemma 4.2. Notice that in the seventh line, $\mathbb{E}_{t-1}^j = \mathbb{E}_{t-1}$ because the random variable β_t only depends on the choice of columns. \square

We want to control the term $r_t = \mathbb{E}\|(\mathbf{I}d_n - \mathbf{P}_i)(\beta_t - \beta_{\mathbf{LN}})\|_2^2$ by bounding it by some α and B such that $r_t \leq \alpha^t B$. We calculate this here:

$$\begin{aligned} \mathbb{E}\|(\mathbf{I}d_n - \mathbf{P}_i)(\beta_t - \beta_{\mathbf{LN}})\|_2^2 &= \mathbb{E}[\mathbb{E}_{t-1}\|(\mathbf{I}d_n - \mathbf{P}_i)(\beta_t - \beta_{\mathbf{LN}})\|_2^2] \\ &\leq \frac{\alpha \mathbb{E}\|\mathbf{X}(\beta_{t-1} - \beta_{\mathbf{LN}})\|_2^2}{\|\mathbf{X}\|_F^2} \\ &\leq \alpha^t \frac{\|\mathbf{X}\beta_0 - \mathbf{X}\beta_{\mathbf{LN}}\|_2^2}{\|\mathbf{X}\|_F^2}. \end{aligned}$$

The first line follows by definition, the second is by Lemma 4.3, and the third by Lemma 4.2.

Finally, we take the expected value of $\|\beta_t^{LN} - \beta_{\mathbf{LN}}\|_2^2$. From equation (4.2) and using Fact 1 we obtain:

$$\begin{aligned} \mathbb{E}\|\beta_t^{LN} - \beta_{\mathbf{LN}}\|_2^2 &= \mathbb{E}\|\mathbf{P}_i(\beta_{t-1}^{LN} - \beta_{\mathbf{LN}})\|_2^2 + \mathbb{E}\|(\mathbf{I}d_n - \mathbf{P}_i)(\beta_t - \beta_{\mathbf{LN}})\|_2^2 \\ &\leq \alpha \mathbb{E}\|\beta_{t-1}^{LN} - \beta_{\mathbf{LN}}\|_2^2 + \mathbb{E}\|(\mathbf{I}d_n - \mathbf{P}_i)(\beta_t - \beta_{\mathbf{LN}})\|_2^2. \end{aligned}$$

We complete the proof using the following lemma from [30]:

LEMMA 4.4. ([30, Thm. 8]) *Suppose that for some $\alpha, \bar{\alpha} < 1$, the following bounds hold for all $t^* \geq 0$:*

$$\mathbb{E}\|\beta_{t^*}^{LN} - \beta_{\mathbf{LN}}\|_2^2 \leq \alpha \mathbb{E}\|\beta_{t^*-1}^{LN} - \beta_{\mathbf{LN}}\|_2^2 + r_{t^*} \text{ and } r_{t^*} \leq \bar{\alpha}^{t^*} B.$$

Then for any $T > 0$,

$$\mathbb{E}\|\beta_T^{LN} - \beta_{\mathbf{LN}}\|_2^2 \leq \alpha^T \|\beta_0^{LN} - \beta_{\mathbf{LN}}\|_2^2 + (\alpha^{\lfloor T/2 \rfloor} + \bar{\alpha}^{\lfloor T/2 \rfloor}) \frac{B}{1 - \alpha}.$$

Letting $\alpha = \bar{\alpha} = \alpha$, $r_t^* = \mathbb{E}\|(\mathbf{I}d_n - \mathbf{P}_i)(\beta_t - \beta_{\mathbf{LN}})\|_2^2$, $B = \frac{\|\mathbf{X}\beta_0 - \mathbf{X}\beta_{\mathbf{LN}}\|_2^2}{\|\mathbf{X}\|_F^2}$, and noting that $\beta_0^{LN} = \beta_0 = 0$, we complete the proof of Theorem 4.1. \square

REMARK 1. *Here we note that the same proof works for overdetermined systems. In particular, this works because Lemma 4.2 holds for $\beta_{\mathbf{LS}}$ and β^* also (see Thm. 3.6 in [15]). Also, Lemma 4.3 follows for both overdetermined consistent systems (see table in Section 3.1) as well as overdetermined inconsistent systems (from (3.12) and subsequent arguments).*

4.3. Comparison. Theorem 4.1 shows that, like the RK and REK methods, REGS converges linearly to the least-norm solution in the underdetermined case. We believe it serves to complement existing analysis and completes the theory of these iterative methods in all three cases of interest. For that reason, we compare the three approaches for the underdetermined setting here. For ease of comparison, set α as in Theorem 4.1, and write $\kappa = \sigma_{\max}(\mathbf{X})/\sigma_{\min}(\mathbf{X})$ for the condition number of \mathbf{X} . From the convergence rate bounds for RK [26] and REK [30] given in Section 3, and after applying elementary bounds to (4.1) of Theorem 4.1, we have:

$$\text{(RK)} \quad \mathbb{E}\|\beta_t - \beta_{LN}\|_2^2 \leq \alpha^t \|\beta_{LN}\|_2^2 \quad (4.3)$$

$$\text{(REK)} \quad \mathbb{E}\|\beta_{2t} - \beta_{LN}\|_2^2 \leq \alpha^t (1 + 2\kappa^2) \|\beta_{LN}\|_2^2 \quad (4.4)$$

$$\text{(REGS)} \quad \mathbb{E}\|\beta_{2t} - \beta_{LN}\|_2^2 \leq \alpha^t (1 + 2\kappa^2) \|\beta_{LN}\|_2^2. \quad (4.5)$$

We find similar results in the overdetermined, inconsistent setting. Using the convergence rate bounds for RGS [15], REK [30], and REGS (Theorem 4.1), also given in section 3, we have:

$$\text{(RGS)} \quad \mathbb{E}\|\beta_t - \beta_{LS}\|_2^2 \leq \alpha^t \|\beta_{LS}\|_2^2 \quad (4.6)$$

$$\text{(REK)} \quad \mathbb{E}\|\beta_{2t} - \beta_{LS}\|_2^2 \leq \alpha^t (1 + 2\kappa^2) \|\beta_{LS}\|_2^2 \quad (4.7)$$

$$\text{(REGS)} \quad \mathbb{E}\|\beta_{2t} - \beta_{LS}\|_2^2 \leq \alpha^t (1 + 2\kappa^2) \|\beta_{LS}\|_2^2. \quad (4.8)$$

Thus, up to constant terms (which are likely artifacts of the proofs), the bounds provide the same convergence rate α , which is not surprising in light of the connections between the methods. In the next section, we compare these approaches experimentally.

5. Empirical Results. In this section we present our experimental results. The code used to run these experiments can be found at [4]. For each experiment, we initialize a matrix \mathbf{X} and vector β with independent standard normal entries and run 50 trials. The right hand side \mathbf{y} is taken to be $\mathbf{X}\beta$. At each iteration t , we keep track of the ℓ_2 -error $\|\beta_t^{LN} - \beta_{LN}\|_2^2$ and fix the stopping criterion to be $\|\beta_t^{LN} - \beta_{LN}\|_2^2 < 10^{-6}$ (of course in practice one chooses a more practical criterion). In each plot, the solid blue line represents the median ℓ_2 -error at iteration t , the light blue shaded region captures the range of error across trials, and the red line represents the theoretical upper bound at each iteration. In Figure 5.1, we show the convergence of β_t^{LN} for varying sized underdetermined linear systems. In Figure 5.2, we show the convergence of a matrix \mathbf{X} of size 700x1000 and its theoretical upper bound. As it turns out, the REGS algorithm often converges much faster than the theoretical worst-case bound.

We also tested REGS on tomography problems using the Regularization toolbox by Hansen [11] (<http://www.imm.dtu.dk/~pcha/Regutools/>). For the 2D tomography problem $\mathbf{X}\beta = \mathbf{y}$ with \mathbf{X} an $m \times n$ matrix where $n = dN^2$ and $m = N^2$, we use $N = 20$ and $d = 3$ for our experiments. Here, \mathbf{X} consists of samples of absorption along a random line on an $N \times N$ grid and d is the oversampling factor. The results from this experiment are shown in Figure 5.2.

We also compare the performance of all four algorithms (RK, REK, RGS, REGS) under the different settings discussed in this paper. Each line in each plot represents the median ℓ_2 -error at that iteration or CPU time over 50 trials using a stopping criterion of 10^{-6} . For the underdetermined case, \mathbf{X} is a 50×500 Gaussian matrix and a 500×50 Gaussian matrix for the overdetermined cases. In the overdetermined, inconsistent case, we set $\mathbf{y} = \mathbf{X}\beta + \mathbf{r}$ where $\mathbf{r} \in \text{null}(\mathbf{X}^*)$ (computed in Matlab using the `null()` function). Figure 5.3, Figure 5.4, and Figure 5.5 show the empirical results for the underdetermined, overdetermined inconsistent, and overdetermined consistent cases respectively. Note we only plot the methods which actually converge to the desired solution in each case. Looking at iterations to convergence, it seems that RK and RGS converge faster than their extended counterparts while REGS and REK converge to the desired solution at about the same rate.




FIGURE 5.1. Left: ℓ_2 -error (log-scale) of REGS on a 150×500 matrix and its theoretical bound. Right: Comparison of ℓ_2 -error (log-scale) of REGS for $m \times 500$ sized matrices with $m = 50, 100, 150$.




FIGURE 5.2. Left: ℓ_2 -error (log-scale) of REGS on a 700×1000 matrix and its theoretical bound. Right: ℓ_2 -error (log-scale) of REGS on the tomography problem with a 400×1200 matrix.

6. Conclusion. The Kaczmarz and Gauss-Seidel methods operate in two different spaces (i.e. row versus column space), but share many parallels. In this paper we drew connections between these two methods, highlighting the similarities and differences in convergence analysis. The approaches possess conflicting convergence properties; RK converges to the desired solution in the underdetermined case but not the inconsistent overdetermined setting, while RGS does the exact opposite. The extended method REK in the Kaczmarz framework fixes this issue, converging to the solution in both scenarios. Here, we present the REGS method, a natural extension of RGS, which completes the overall picture. We hope that our unified analysis of all four methods will assist researchers working with these approaches.




FIGURE 5.3. Comparison of median ℓ_2 -error (log-scale) of RK, REK, and REGS for an underdetermined system.




FIGURE 5.4. Comparison of median ℓ_2 -error (log-scale) of RGS, REK, and REGS for an overdetermined, inconsistent system.

Acknowledgments. Needell was partially supported by NSF CAREER grant #1348721, and the Alfred P. Sloan Fellowship. Ma was supported in part by AFOSR MURI grant FA9550-10-1-0569. Ramdas was supported in part by ONR MURI grant N000140911052. The authors would also like to thank the Institute of Pure and Applied Mathematics (IPAM) at which this collaboration started, and the reviewers of this manuscript for their thoughtful suggestions.

References.

- [1] Byrne, C. L. [2008], *Applied iterative methods*, A K Peters Ltd., Wellesley, MA.
- [2] Censor, Y., Eggermont, P. P. B. and Gordon, D. [1983], ‘Strong underrelaxation in Kaczmarz’s method for inconsistent systems’, *Numerische Mathematik*




FIGURE 5.5. Comparison of median l_2 -error (log-scale) of RK, RGS, REK, and REGS for an overdetermined, consistent system

- 41(1), 83–92.
- [3] Chen, X. and Powell, A. [2012], ‘Almost sure convergence of the Kaczmarz algorithm with random measurements’, *J. Fourier Anal. Appl.* pp. 1–20. 10.1007/s00041-012-9237-2.
URL: <http://dx.doi.org/10.1007/s00041-012-9237-2>
- [4] *Code link* [n.d.]. <http://www.cmc.edu/pages/faculty/DNeedell/regs.zip>.
- [5] Dumitrescu, B. [2014], ‘On the relation between the randomized extended kaczmarz algorithm and coordinate descent’, *BIT Numerical Mathematics* pp. 1–11.
- [6] Eldar, Y. C. and Needell, D. [2011], ‘Acceleration of randomized Kaczmarz method via the Johnson-Lindenstrauss lemma’, *Numer. Algorithms* **58**(2), 163–177.
URL: <http://dx.doi.org/10.1007/s11075-011-9451-z>
- [7] Elfving, T. [1980], ‘Block-iterative methods for consistent and inconsistent linear equations’, *Numer. Math.* **35**(1), 1–12.
URL: <http://dx.doi.org/10.1007/BF01396365>
- [8] Gordon, R., Bender, R. and Herman, G. T. [1970], ‘Algebraic reconstruction techniques (ART) for three-dimensional electron microscopy and X-ray photography’, *J. Theoret. Biol.* **29**, 471–481.
- [9] Hamaker, C. and Solmon, D. C. [1978], ‘The angles between the null spaces of X-rays’, *J. Math. Anal. Appl.* **62**(1), 1–23.
- [10] Hanke, M. and Niethammer, W. [1990], ‘On the acceleration of Kaczmarz’s method for inconsistent linear systems’, *Linear Algebra and its Applications* **130**, 83–98.
- [11] Hansen, P. C. [1994], ‘Regularization tools: A matlab package for analysis and solution of discrete ill-posed problems’, *Numerical algorithms* **6**(1), 1–35.
- [12] Herman, G. and Meyer, L. [1993], ‘Algebraic reconstruction techniques can be made computationally efficient’, *IEEE Trans. Medical Imaging* **12**(3), 600–609.
- [13] Herman, G. T. [2009], *Fundamentals of computerized tomography: image reconstruction from projections*, Springer.
- [14] Kaczmarz, S. [1937], ‘Angenäherte auflösung von systemen linearer gleichungen’,

- Bull. Int. Acad. Polon. Sci. Lett. Ser. A* pp. 335–357.
- [15] Leventhal, D. and Lewis, A. S. [2010], ‘Randomized methods for linear constraints: convergence rates and conditioning’, *Math. Oper. Res.* **35**(3), 641–654.
URL: <http://dx.doi.org/10.1287/moor.1100.0456>
- [16] Liu, J., Wright, S. J. and Sridhar, S. [2014], ‘An asynchronous parallel randomized kaczmarz algorithm’, *arXiv preprint arXiv:1401.4780* .
- [17] Natterer, F. [2001], *The mathematics of computerized tomography*, Vol. 32 of *Classics in Applied Mathematics*, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA. Reprint of the 1986 original.
URL: <http://dx.doi.org/10.1137/1.9780898719284>
- [18] Needell, D. [2010], ‘Randomized Kaczmarz solver for noisy linear systems’, *BIT* **50**(2), 395–403.
URL: <http://dx.doi.org/10.1007/s10543-010-0265-5>
- [19] Needell, D., Sbrero, N. and Ward, R. [2014], ‘Stochastic gradient descent and the randomized kaczmarz algorithm’, *Math. Program. Series A* . to appear.
- [20] Needell, D. and Tropp, J. A. [2013], ‘Paved with good intentions: Analysis of a randomized block kaczmarz method’, *Linear Algebra and its Applications* .
- [21] Needell, D. and Ward, R. [2013], ‘Two-subspace projection method for coherent overdetermined linear systems’, *Journal of Fourier Analysis and Applications* **19**(2), 256–269.
- [22] Nesterov, Y. [2012], ‘Efficiency of coordinate descent methods on huge-scale optimization problems’, *SIAM J. Optimiz.* **22**(2), 341–362.
- [23] Popa, C. [1998], ‘Extensions of block-projections methods with relaxation parameters to inconsistent and rank-deficient least-squares problems’, *BIT* **38**(1), 151–176.
URL: <http://dx.doi.org/10.1007/BF02510922>
- [24] Popa, C., Preclik, T., Köstler, H. and Rüdiger, U. [2012], ‘On Kaczmarz’s projection iteration as a direct solver for linear least squares problems’, *Linear Algebra and Its Applications* **436**(2), 389–404.
- [25] Richtárik, P. and Takáč, M. [2012], ‘Iteration complexity of randomized block-coordinate descent methods for minimizing a composite function’, *Math. Program.* pp. 1–38.
- [26] Strohmer, T. and Vershynin, R. [2009], ‘A randomized Kaczmarz algorithm with exponential convergence’, *J. Fourier Anal. Appl.* **15**(2), 262–278.
URL: <http://dx.doi.org/10.1007/s00041-008-9030-4>
- [27] Tanabe, K. [1971], ‘Projection method for solving a singular system of linear equations and its applications’, *Numerische Mathematik* **17**(3), 203–214.
- [28] Whitney, T. M. and Meany, R. K. [1967], ‘Two algorithms related to the method of steepest descent’, *SIAM Journal on Numerical Analysis* **4**(1), 109–118.
- [29] Xu, J. and Zikatanov, L. [2002], ‘The method of alternating projections and the method of subspace corrections in Hilbert space’, *J. Amer. Math. Soc.* **15**(3), 573–597.
URL: <http://dx.doi.org/10.1090/S0894-0347-02-00398-3>
- [30] Zouzias, A. and Freris, N. M. [2013], ‘Randomized extended kaczmarz for solving least squares’, *SIAM Journal on Matrix Analysis and Applications* **34**(2), 773–793.