# Superfast Tikhonov Regularization of Toeplitz Systems 

Christopher K. Turnes, Doru Balcan, Justin Romberg*

June 25, 2018


#### Abstract

Toeplitz-structured linear systems arise often in practical engineering problems. Correspondingly, a number of algorithms have been developed that exploit Toeplitz structure to gain computational efficiency when solving these systems. The earliest "fast" algorithms for Toeplitz systems required $\mathcal{O}\left(n^{2}\right)$ operations, while more recent "superfast" algorithms reduce the cost to $\mathcal{O}\left(n \log ^{2} n\right)$ or below.

In this work, we present a superfast algorithm for Tikhonov regularization of Toeplitz systems. Using an "extension-and-transformation" technique, our algorithm translates a Tikhonov-regularized Toeplitz system into a type of specialized polynomial problem known as tangential interpolation. Under this formulation, we can compute the solution in only $\mathcal{O}\left(n \log ^{2} n\right)$ operations. We use numerical simulations to demonstrate our algorithm's complexity and verify that it returns stable solutions.


## 1 Introduction

This paper develops a computationally efficient and numerically stable algorithm for solving systems of equations with Toeplitz structure. Toeplitz matrices, which arise in problems involving temporally- or spatially-invariant systems, have constant diagonal coefficients $T=\left[a_{i-j}\right]$. This type of structure can be exploited to accelerate the calculation of least-squares solutions, a strategy first adopted in the mid 1900s when Levinson [1], Durbin [2], and others [3] [4] [5] developed a series of inversion algorithms that reduced the cost of solving an $n \times n$ Toeplitz system from $\mathcal{O}\left(n^{3}\right)$ to only $\mathcal{O}\left(n^{2}\right)$ operations. Algorithms like Levinson's are called fast, and they have proven exceedingly useful in a wide variety of signal-processing applications. A concise but comprehensive review of fast Toeplitz solvers is available in [6].

While these algorithms have small overhead costs and are well-suited for small- to medium-size problems, they require a number of operations an order higher than the number of free parameters in the systems they solve. Over the past few decades, new algorithms have bridged this gap by inverting Toeplitz matrices in strictly less than $\mathcal{O}\left(n^{2}\right)$ operations (we note in particular [7], [8], and [9]), earning the designation "superfast." While usually more complicated than their predecessors, superfast algorithms can provide an enormous reduction in computation time for large or very large matrices.

Unfortunately, inversion is rarely practical; the matrices that arise in applications are seldom square and nonsingular. As a result, many Toeplitz inversion methods have been modified to solve more general

[^0]least-squares problems. For example, a superfast pseudoinversion algorithm was given in [10], where the authors extended the algorithm of [8] to solve rectangular Toeplitz systems with full column rank. Their extended algorithm applies the pseudoinverse in $\mathcal{O}\left(N \log ^{2} N\right)$ operations, where $N$ is the number of free parameters defining the system. Of course, this is not an isolated development; other notable Toeplitz least-squares solvers include [11], [12], [13], and [14], to name a few.

In this work, we further adapt the approach of [10] for Tikhonov regularization of Toeplitz systems. Our results are not unprecedented; a superfast Tikhonov solver based on displacement structure (first introduced in [15]) was presented in [16]. However, ours takes a fundamentally different approach: we reformulate the system as an interpolation problem instead of using its structure to accelerate matrix factorization. This strategy allows our algorithm to be more general in scope. Whereas the approach of [16] applies only when the regularization penalizes solutions of large Euclidean norm, our algorithm is applicable for any Toeplitz regularizer, easily extends to include a number of different regularization terms, and applies when the Gramian $G_{T}=T^{H} T$, but not the matrix $T$ itself, is Toeplitz.

The remainder of the paper is organized as follows. In Section 2, we formulate the Tikhonovregularization problem in a general sense and discuss the specific cases our algorithm addresses. To conclude the section, we outline our two-part "extension-and-transformation" approach. Subsequently, we develop the first portion of this approach by reframing the regularization problems as partial-circulantblock systems in Section 3. In Section 4, we use Fourier operators to transform these systems into polynomial-interpolation problems and explain how such problems may be solved efficiently. We then summarize our algorithm in Section 5 and discuss implementation issues in Section 6. In Section 7, we give the results of numerical simulations verifying the cost and utility of our algorithm. Finally, we summarize our developments and detail potential future extensions in Section 8.

We have implemented our algorithm as part of a comprehensive code base written in MATLAB and $\mathrm{C}++$. The code package can be obtained online ${ }^{1}$, along with scripts that reproduce all of the numerical experiments in Section 7.

## 2 Problem formulation

When a system of equations $T x=b$ is ill-posed, either because the solution $x$ is non-unique or because it is non-existent, a least-squares solution $\hat{x}$ is usually calculated instead. The solution $\hat{x}$ is designed to minimize the norm of the residual vector $r=b-T \hat{x}$, but might have undesirable properties depending on the characteristics of the matrix $T$. A common technique to avoid this problem is to regularize the system by including a term that penalizes solutions that are not well-behaved. This approach is known as Tikhonov regularization ${ }^{2}$, and is formulated by expressing $\hat{x}$ as the minimizer of the optimization problem

$$
\begin{equation*}
\hat{x}=\underset{x}{\arg \min }\|T x-b\|^{2}+\|L x\|^{2} . \tag{1}
\end{equation*}
$$

The matrix $L$, which provides the regularization, is referred to as the Tikhonov matrix or the regularizer. If a solution to (1) exists, it is expressed in closed form as

$$
\begin{align*}
\hat{x} & =\left(T^{H} T+L^{H} L\right)^{-1} T^{H} b \\
& =\left(G_{T}+G_{L}\right)^{-1} T^{H} b, \tag{2}
\end{align*}
$$

[^1]where $G_{T}$ and $G_{L}$ are the Gramians of $T$ and $L$, respectively.
In this work, we exploit Toeplitz structure to reduce the number of operations necessary to compute (2). Throughout the exposition, we consider three specific scenarios:

- General problem: $T$ and $L$ are $m \times n$ and $p \times n$ Toeplitz.
- $\ell_{2}$-norm penalization: $T$ is $m \times n$ Toeplitz while $L$ is a scaled identity matrix.
- Toeplitz-Gramian problem: $G_{T}=T^{H} T$ is $n \times n$ Hermitian Toeplitz while $L$ is $p \times n$ Toeplitz.

In fact, our approach can be applied to a much larger class of problems as well; these three simply represent those cases that appear to be most interesting.

To solve these problems, we present an algorithm that calculates (2) with superfast complexity $\mathcal{O}\left(N \log ^{2} N\right)$, where $N$ is the number of free parameters. Our algorithm can be divided into three stages. First, we decouple the Gramian matrices $G_{T}$ and $G_{L}$ to express the problem as a Toeplitz-block system. ${ }^{3}$ Next, we follow the "extension-and-transformation" approach of [17] to translate our linear-algebraic problems into the context of a polynomial problem known as tangential interpolation [18]. Finally, using strategies similar to those of [10], we solve the tangential-interpolation problems in $\mathcal{O}\left(N \log ^{2} N\right)$ to compute the solutions $\hat{x}$.

## 3 Extensions of Tikhonov systems

Our algorithm reformulates linear-algebraic problems as polynomial problems. The first step in achieving this change of context is to expand the linear systems by adding block rows and columns. We motivate this process by demonstrating how a Toeplitz system may be extended to form an equivalent partialcirculant system. We then use this technique to extend the Tikhonov systems of the three problems, expressing them as partial-circulant-block systems. As a point of notation, we will indicate block row and column sizes outside of matrix brackets to better communicate submatrix sizes.

### 3.1 Circulant extensions of Toeplitz matrices

In this section, we present a simple extension of Toeplitz matrices into circulant submatrices. This technique can be used to replace a Toeplitz problem with a larger circulant-like system, but will not add or remove any information. Instead, it will yield a form that is more easily manipulable.

Consider the linear system $T x=b$, where $T=\left[a_{i-j}\right]$ is $m \times n$ Toeplitz. Our objective is to find a system with circulant structure that we might solve in place of the original problem. For any $k \geq 0$, we define an integer $N_{k}:=m+n+k-1$. If $\bar{T}$ is the $(n+k-1) \times n$ Toeplitz matrix

$$
\bar{T}=\left[\begin{array}{cccc}
a_{-n-k} & a_{m-1} & \cdots & a_{m-n+1} \\
a_{-n-k+1} & a_{-n-k} & \cdots & a_{m-n+2} \\
\vdots & \vdots & \ddots & \vdots \\
a_{-1} & a_{-2} & \cdots & a_{-n}
\end{array}\right]
$$

[^2]where the coefficients $a_{i}$ are arbitrary for $i \leq-n$, then the $N_{k} \times n$ matrix
\[

C_{k}(T)=\left[$$
\begin{array}{c}
\frac{\mathrm{n}}{\bar{T}}  \tag{3}\\
T
\end{array}
$$\right] $$
\begin{aligned}
& \mathrm{n}+\mathrm{k}-1 \\
& \mathrm{~m}
\end{aligned}
$$
\]

consists of the first $n$ columns of the $N_{k} \times N_{k}$ circulant matrix

$$
\operatorname{circ}_{k}(T)=\left[\begin{array}{ccc}
a_{-n-k} & \cdots & a_{-n-k+1} \\
\vdots & \ddots & \vdots \\
a_{m-1} & \cdots & a_{-n-k}
\end{array}\right]
$$

We call $\bar{T}$ the extension matrix of $T$, while $C_{k}(T)$ is the $k$-circulant extension of $T$.
The matrix $T$ in the original linear system cannot simply be replaced by its circulant extension $C_{k}(T)$, as this would introduce extra rows to the system without accounting for them on the right-hand side of the equation. Instead, we must also add an equal number of columns, which we generate by defining an artificial unknown $\gamma=-\bar{T} x .{ }^{4}$ Using $\gamma$, if $\mathcal{I}_{n+k-1}$ is the $(n+k-1) \times(n+k-1)$ identity matrix, the linear system $T x=b$ is equivalent to

$$
\left.\begin{array}{cc} 
 \tag{4}\\
\mathrm{n}+\mathrm{k}-1 \\
\mathrm{~m}
\end{array} \quad \begin{array}{cc}
\mathrm{n} & \mathrm{n}+\mathrm{k}-1 \\
\bar{T} & \mathcal{I}_{n+k-1} \\
T & \mathbf{0}
\end{array}\right]\left[\begin{array}{l}
x \\
\gamma
\end{array}\right]=\left[\begin{array}{l}
\mathbf{0} \\
b
\end{array}\right] .
$$

The expansion adds just as many columns as rows, so the new system is as over- or underdetermined as the original.

The blocks of (4) can be grouped into partial-circulant matrices. Let the matrix subscript $\Gamma_{\ell}$ represent the submatrix formed from the first $\ell$ columns of a matrix. For example, the $k$-circulant extension $C_{k}(T)$ can be written as $\left(\operatorname{circ}_{k}(T)\right)_{\Gamma_{n}}$. Using this notation, and grouping the block rows of (4) together, each block column is a circulant submatrix:

$$
\left[\begin{array}{ll}
C_{k}(T) & \left(\mathcal{I}_{N_{k}}\right)_{\Gamma_{n+k-1}}
\end{array}\right]\left[\begin{array}{l}
x \\
\gamma
\end{array}\right]=\left[\begin{array}{l}
\mathbf{0} \\
b
\end{array}\right] .
$$

This grouping is advantageous; circulant matrices are easier to manipulate than Toeplitz matrices, as they are diagonalized by the Fourier matrix. Namely, if $\mathcal{F}_{N_{k}}$ is the $N_{k} \times N_{k}$ Fourier matrix, the matrix $C_{k}(T)$ may be decomposed as

$$
\begin{equation*}
C_{k}(T)=\left(\mathcal{F}_{N_{k}}^{H} \Lambda \mathcal{F}_{N_{k}}\right)_{\Gamma_{n}}=\mathcal{F}_{N_{k}}^{H} \Lambda\left(\mathcal{F}_{N_{k}}\right)_{\Gamma_{n}} \tag{5}
\end{equation*}
$$

for some $N_{k} \times N_{k}$ diagonal matrix $\Lambda$. Of course, there exists a similar decomposition for the block $\left(\mathcal{I}_{N_{k}}\right)_{\Gamma_{n+k}}$ as well. To avoid compounding subscripts as in (5), we will omit those indicating the size of the Fourier matrices, as they can be easily inferred from context. For instance, we express (5) as

$$
C_{k}(T)=\mathcal{F}^{H} \Lambda \mathcal{F}_{\Gamma_{n}} .
$$

[^3]
### 3.2 Extension for the general problem

We now apply the technique of Section 3.1 to the Tikhonov system in (2) for the general problem. First, we define two artificial variables $\sigma_{1}:=T \hat{x}$ and $\sigma_{2}:=L \hat{x}$, which allow us to decouple the Gramian matrices:

$$
\begin{aligned}
T^{H} b & =T^{H}(T x)+L^{H}(L x) \\
& =T^{H} \sigma_{1}+L^{H} \sigma_{2}
\end{aligned}
$$

Using these variable definitions, (2) may be replaced by a system containing only Toeplitz blocks:

$$
\left[\begin{array}{ccc}
\mathbf{0} & T^{H} & L^{H}  \tag{6}\\
-T & \mathcal{I}_{m} & \mathbf{0} \\
-L & \mathbf{0} & \mathcal{I}_{p}
\end{array}\right]\left[\begin{array}{c}
\hat{x} \\
\sigma_{1} \\
\sigma_{2}
\end{array}\right]=\left[\begin{array}{c}
T^{H} b \\
\mathbf{0} \\
\mathbf{0}
\end{array}\right]
$$

Since all of the blocks of this system are Toeplitz, we can extend them as in (4), using 1-circulant extensions for simplicity. Since the extension size must be large enough that all of the blocks become circulant submatrices, the first block row of the extended system will have $q+n$ total rows, where $q=\max (m, p)$.

If we are to add the extensions $\overline{T^{H}}, \overline{L^{H}}, \bar{T}$, and $\bar{L}$ to our system, we must introduce three corresponding artificial variables to compensate for the extra rows:

$$
\gamma_{1}:=-\overline{T^{H}} \sigma_{1}-\overline{L^{H}} \sigma_{2}, \quad \quad \gamma_{2}:=\bar{T} \hat{x}, \quad \gamma_{3}:=\bar{L} \hat{x}
$$

With these variables defined, and letting $\hat{\boldsymbol{b}}=-T^{H} b$, the system of (6) has the equivalent form

$$
\begin{align*}
& \begin{array}{l}
\mathrm{q} \\
\mathrm{n} \\
\mathrm{n} \\
\mathrm{~m} \\
\mathrm{~m} \\
\mathrm{n} \\
\mathrm{p}
\end{array}\left[\begin{array}{ccccccc}
\mathrm{n} & \mathrm{~m} & \mathrm{p} & \mathrm{q} & \mathrm{n} & \mathrm{n} & 1 \\
\mathbf{0} & \overline{T^{H}} & \overline{L^{H}} & \mathcal{I}_{q} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\
\mathbf{0} & T^{H} & L^{H} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \hat{\boldsymbol{b}} \\
-\bar{T} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathcal{I}_{n} & \mathbf{0} & \mathbf{0} \\
-T & \mathcal{I}_{m} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\
-\bar{L} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathcal{I}_{n} & \mathbf{0} \\
-L & \mathbf{0} & \mathcal{I}_{p} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0}
\end{array}\right]\left[\begin{array}{c}
\hat{x} \\
\sigma_{1} \\
\sigma_{2} \\
\gamma_{1} \\
\gamma_{2} \\
\gamma_{3} \\
1
\end{array}\right]=\mathbf{0}  \tag{7}\\
& C \quad p^{*}=0,
\end{align*}
$$

where we have moved the right-hand side vector into the matrix to generate a homogeneous system of equations. We have also expressed the unknown variables as a single block column vector $\boldsymbol{p}^{*}$ to simplify future expressions.

Grouping the blocks of the matrix $\boldsymbol{C}$ together to form circulant submatrices, we may replace the original system of (2) with the much larger system

$$
\left.\begin{array}{c} 
 \tag{8}\\
\mathrm{q}+\mathrm{n} \\
\mathrm{~m}+\mathrm{n} \\
\mathrm{p}+\mathrm{n}
\end{array} \begin{array}{ccccccc}
\mathrm{n} & \mathrm{~m} & \mathrm{p} & \mathrm{q} & \mathrm{n} & \mathrm{n} & 1 \\
\mathbf{0} & C_{12} & C_{13} & C_{14} & \mathbf{0} & \mathbf{0} & C_{17} \\
C_{21} & C_{22} & \mathbf{0} & \mathbf{0} & C_{25} & \mathbf{0} & \mathbf{0} \\
C_{31} & \mathbf{0} & C_{33} & \mathbf{0} & \mathbf{0} & C_{36} & \mathbf{0}
\end{array}\right] \boldsymbol{p}^{*}=\mathbf{0} .
$$

The $\left\{C_{i j}\right\}$ are circulant submatrices, and can be factored with Fourier matrices of the appropriate sizes.

### 3.3 Extension for $\ell_{2}$-norm penalization

The system in (8) may be simplified if the Tikhonov matrix is a scaled identity matrix (i.e., if $L=\beta \mathcal{I}_{n}$ for some $\beta \in \mathbb{C}$ ). Specifically, it may be drastically reduced in size by observing that the extension matrix for $L$ is $\mathbf{0}$ if all arbitrary entries are set to zero. Using this fact, the extended system for $\ell_{2}$-norm penalization is given by

| q |
| :--- |
| n |
| n |
| n |
| m |
| n |
| n |
| n |\(\left[\begin{array}{ccccccc}\mathrm{n} \& \mathrm{m} \& \mathrm{n} \& \mathrm{q} \& \mathrm{n} \& \mathrm{n} \& 1 <br>

\mathbf{0} \& \overline{T^{H}} \& \mathbf{0} \& \mathcal{I}_{q} \& \mathbf{0} \& \mathbf{0} \& \mathbf{0} <br>
\mathbf{0} \& T^{H} \& \beta^{*} \mathcal{I}_{n} \& \mathbf{0} \& \mathbf{0} \& \mathbf{0} \& \hat{\boldsymbol{b}} <br>
-\bar{T} \& \mathbf{0} \& \mathbf{0} \& \mathbf{0} \& \mathcal{I}_{n} \& \mathbf{0} \& \mathbf{0} <br>
-T \& \mathcal{I}_{m} \& \mathbf{0} \& \mathbf{0} \& \mathbf{0} \& \mathbf{0} \& \mathbf{0} <br>
\mathbf{0} \& \mathbf{0} \& \mathbf{0} \& \mathbf{0} \& \mathbf{0} \& \mathcal{I}_{n} \& \mathbf{0} <br>
-\beta \mathcal{I}_{n} \& \mathbf{0} \& \mathcal{I}_{n} \& \mathbf{0} \& \mathbf{0} \& \mathbf{0} \& \mathbf{0}\end{array}\right]\left[$$
\begin{array}{c}\hat{x} \\
\sigma_{1} \\
\sigma_{2} \\
\gamma_{1} \\
\gamma_{2} \\
\gamma_{3} \\
1\end{array}
$$\right]=\mathbf{0}\),
where $\beta^{*}$ is the complex conjugate of $\beta$.
The fifth block row of (9) implies $\gamma_{3}=\mathbf{0}$, while the sixth implies $\beta \hat{x}=\sigma_{2}$. As a result, we can eliminate both $\gamma_{3}$ and $\sigma_{2}$ as variables. Since $\sigma_{2}$ has been removed, the only matrix with a non-zero extension in the first block row is $T^{H}$, and therefore we set $q=m$ regardless of whether or not $m<n$.

After reducing the number of unknowns, we are left with the following simplified system:

$$
\begin{gathered}
\\
\mathrm{m} \\
\mathrm{n} \\
\mathrm{n} \\
\mathrm{~m}
\end{gathered}\left[\begin{array}{ccccc}
\mathrm{n} & \mathrm{~m} & \mathrm{~m} & \mathrm{n} & 1 \\
\mathbf{0} & \overline{T^{H}} & \mathcal{I}_{m} & \mathbf{0} & \mathbf{0} \\
|\beta|^{2} \mathcal{I}_{n} & T^{H} & \mathbf{0} & \mathbf{0} & \hat{\boldsymbol{b}} \\
-\bar{T} & \mathbf{0} & \mathbf{0} & \mathcal{I}_{n} & \mathbf{0} \\
-T & \mathcal{I}_{m} & \mathbf{0} & \mathbf{0} & \mathbf{0}
\end{array}\right]\left[\begin{array}{c}
\hat{x} \\
\sigma_{1} \\
\gamma_{1} \\
\gamma_{2} \\
1
\end{array}\right]=\boldsymbol{C} \boldsymbol{p}^{*}=\mathbf{0} .
$$

Again grouping the Toeplitz components together with their extension matrices, we can express this as a system with partial-circulant blocks:

$$
\begin{gather*}
\mathrm{n}  \tag{10}\\
\mathrm{~m}+\mathrm{n} \\
\mathrm{~m}+\mathrm{n}
\end{gather*}\left[\begin{array}{ccccc}
C_{11} & C_{12} & C_{13} & \mathrm{n} & 1 \\
C_{21} & C_{22} & \mathbf{0} & C_{24} & C_{15}
\end{array}\right] \boldsymbol{p}^{*}=\mathbf{0}
$$

This system is significantly smaller than that of (8), and it can be solved more efficiently.

### 3.4 Extension for the Toeplitz-Gramian problem

If the Gramian matrix $G_{T}=T^{H} T$ is Toeplitz, the Tikhonov system can be extended in a manner similar to the developments of Section 3.2. However, this is not always a wise approach; if $L$ is a scaled identity matrix, the matrix $\left(G_{T}+L^{H} L\right)$ is also Toeplitz, meaning the problem may be solved by the directinversion algorithm of [8] (or numerous alternatives). Therefore, when we consider the Toeplitz-Gramian problem, we assume without loss of generality that $L$ is a generic $p \times n$ Toeplitz matrix.

For this problem, we need only introduce a single artificial variable $\sigma=L x$. Using the two equations

$$
\begin{aligned}
G_{T} x+L^{H} \sigma & =T^{H} b \\
-L x+\sigma & =0
\end{aligned}
$$

we can perform 1-circulant extensions on $G_{T}, L^{H}$, and $L$ to obtain the block system

$$
\begin{gathered}
\mathrm{q} \\
\mathrm{n} \\
\mathrm{n} \\
\mathrm{p}
\end{gathered}\left[\begin{array}{ccccc}
\mathrm{n} & \mathrm{p} & \mathrm{q} & \mathrm{n} & 1 \\
\mathrm{p} & \overline{L^{H}} & \mathcal{I}_{q} & \mathbf{0} & \mathbf{0} \\
G & L^{H} & \mathbf{0} & \mathbf{0} & \hat{\boldsymbol{b}} \\
-\bar{L} & \mathbf{0} & \mathbf{0} & \mathcal{I}_{n} & \mathbf{0} \\
-L & \mathcal{I}_{p} & \mathbf{0} & \mathbf{0} & \mathbf{0}
\end{array}\right]\left[\begin{array}{c}
x \\
\sigma \\
\gamma_{1} \\
\gamma_{2} \\
1
\end{array}\right]=\boldsymbol{C} \boldsymbol{p}^{*}=\mathbf{0} .
$$

Since we are working with the Gramian of $T$ rather than $T$ itself, $q=\max (n, p)$.
Grouping the component blocks together yields the partial-circulant-block system:

$$
\begin{array}{ccccc}
\mathrm{q}+\mathrm{n}  \tag{11}\\
\mathrm{p}+\mathrm{n}
\end{array}\left[\begin{array}{ccccc}
\mathrm{p} & \mathrm{q} & \mathrm{n} & 1 \\
C_{11} & C_{12} & C_{13} & \mathbf{0} & C_{15} \\
C_{21} & C_{22} & \mathbf{0} & C_{24} & \mathbf{0}
\end{array}\right] \boldsymbol{p}^{*}=\mathbf{0} .
$$

This is an identical formulation to the $\ell_{2}$-norm-penalization problem, but with the circulant blocks $\left\{C_{i j}\right\}$ defined differently (and with different matrix sizes). Again, compared to the system of (8), the system of (11) is significantly smaller and simpler to solve.

## 4 Transforming and solving block circulant systems

In this section, we detail the "transformation" part of our approach, using Fourier operators to translate the systems of (8), (10), and (11) into the language of polynomials. This new formulation allows us to solve interpolation problems rather than linear-algebraic problems, and is based on the framework of [19]. Our approach is not confined to the examples we consider; any nonsingular system having only partial-circulant blocks may be solved in a similar way.

We consider only the general problem for the rest of the analysis, as the requisite adaptations for $\ell_{2}$-norm penalization and the Toeplitz-Gramian problem are straightforward. While there are fewer interpolation conditions for the latter two problems, their solution methods are virtually identical and all theoretical results extend easily. We also assume that $m=p=q$ for simplicity, though this assumption is not fundamental to our results.

### 4.1 Transformation

After extending the systems, we place them in the context of polynomials by applying Fourier operators. Let $N=m+n$ be the number of rows in each block row of (8). Using the diagonal decomposition of (5), we can left-multiply by the $3 N \times 3 N$ Fourier-block operator

$$
\mathcal{I}_{3} \otimes \mathcal{F}=\left[\begin{array}{ccc}
\mathcal{F} & \mathbf{0} & \mathbf{0} \\
\mathbf{0} & \mathcal{F} & \mathbf{0} \\
\mathbf{0} & \mathbf{0} & \mathcal{F}
\end{array}\right]
$$

transforming the homogeneous system into the form

$$
\begin{gather*}
 \tag{12}\\
\\
\mathrm{N} \\
\mathrm{~N}
\end{gather*}\left[\begin{array}{ccccccc}
\mathrm{N} & \mathrm{~N} & \mathrm{~N} & \mathrm{~N} & \mathrm{~N} & \mathrm{~N} & \mathrm{~N} \\
\mathrm{~N} & \Lambda_{12} & \Lambda_{13} & \Lambda_{14} & \mathbf{0} & \mathbf{0} & \Lambda_{17} \\
\mathrm{~N}
\end{array}\left[\begin{array}{cccccc}
\Lambda_{21} & \Lambda_{22} & \mathbf{0} & \mathbf{0} & \Lambda_{25} & \mathbf{0} \\
\Lambda_{31} & \mathbf{0} & \Lambda_{33} & \mathbf{0} & \mathbf{0} & \Lambda_{36} \\
\mathbf{0}
\end{array}\right] \boldsymbol{F}^{*}=\mathbf{0},\right.
$$

where $\Lambda_{i j}=\operatorname{diag}\left(\lambda_{i j}^{(0)}, \ldots, \lambda_{i j}^{(N-1)}\right)$ is the diagonal matrix in the factorization of $C_{i j}$ and $\boldsymbol{F}$ is the blockdiagonal matrix

$$
\boldsymbol{F}=\operatorname{diag}\left(\mathcal{F}_{\Gamma_{n}}, \mathcal{F}_{\Gamma_{m}}, \mathcal{F}_{\Gamma_{m}}, \mathcal{F}_{\Gamma_{m}}, \mathcal{F}_{\Gamma_{n}}, \mathcal{F}_{\Gamma_{n}}, \mathcal{F}_{\Gamma_{1}}\right)
$$

The transformation allows us to replace (8) with a set of polynomial interpolation conditions. The key to such a shift in perspective is to equate a vector $\rho=\left[\rho_{i}\right]$ with a polynomial $\rho(z)=\sum_{i} \rho_{i} z^{i-1}$. Under this equivalence, the product

$$
\mathcal{F} \rho=\left[\begin{array}{ccc}
\omega_{0}^{0} & \cdots & \omega_{0}^{N-1}  \tag{13}\\
\vdots & \ddots & \vdots \\
\omega_{N-1}^{0} & \cdots & \omega_{N-1}^{N-1}
\end{array}\right]\left[\begin{array}{c}
\rho_{0} \\
\vdots \\
\rho_{N-1}
\end{array}\right]=\left[\begin{array}{c}
\rho\left(\omega_{0}\right) \\
\vdots \\
\rho\left(\omega_{N-1}\right)
\end{array}\right]
$$

where $\omega_{k}=\mathbf{e}^{\mathbf{j} 2 \pi k / N}$, produces a vector of the polynomial evaluations $\left[\rho\left(\omega_{i}\right)\right]$.
The components of $\boldsymbol{p}^{*}$ in (12) correspond to polynomials $\left\{\hat{x}(z), \sigma_{1}(z), \sigma_{2}(z), \gamma_{1}(z), \gamma_{2}(z), \gamma_{3}(z), 1\right\}$. Applying the equivalence of (13), the product $\boldsymbol{F} \boldsymbol{p}^{*}$ may be treated as an unknown vector containing the values of these polynomials at the nodes $\omega_{k}$. Since the diagonal matrices $\Lambda_{i j}$ scale these values, (12) is equivalent to a set of $3 N$ interpolation conditions:

$$
\begin{align*}
& \lambda_{12}^{(k)} \sigma_{1}\left(\omega_{k}\right)+\lambda_{13}^{(k)} \sigma_{2}\left(\omega_{k}\right)+\lambda_{14}^{(k)} \gamma_{1}\left(\omega_{k}\right)+\lambda_{17}^{(k)}=0, \\
& \lambda_{21}^{(k)} \hat{x}\left(\omega_{k}\right)+\lambda_{22}^{(k)} \sigma_{1}\left(\omega_{k}\right)+\lambda_{25}^{(k)} \gamma_{2}\left(\omega_{k}\right)=0, \text { and }  \tag{14}\\
& \lambda_{31}^{(k)} \hat{x}\left(\omega_{k}\right)+\lambda_{33}^{(k)} \sigma_{2}\left(\omega_{k}\right)+\lambda_{36}^{(k)} \gamma_{3}\left(\omega_{k}\right)=0
\end{align*}
$$

for $k=0, \ldots, N-1$. These equations are markedly different from those of a standard interpolation problem. Rather than prescribing individual polynomial values at each node, the conditions define values of weighted sums of polynomials.

To solve this type of problem, we gather the unknown polynomials into a vector polynomial $p^{*}(z) \in$ $\mathbb{C}[z]^{7 \times 1}$. Defining the interpolation conditions

$$
\begin{aligned}
F_{k} & =\left[\begin{array}{ccccccc}
0 & \lambda_{1,2}^{(k)} & \lambda_{1,3}^{(k)} & \lambda_{1,4}^{(k)} & 0 & 0 & \lambda_{1,7}^{(k)}
\end{array}\right] \\
G_{k} & =\left[\begin{array}{lllllll}
\lambda_{2,1}^{(k)} & \lambda_{2,2}^{(k)} & 0 & 0 & \lambda_{2,5}^{(k)} & 0 & 0
\end{array}\right], \text { and } \\
H_{k} & =\left[\begin{array}{lllllll}
\lambda_{3,1}^{(k)} & 0 & \lambda_{3,3}^{(k)} & 0 & 0 & \lambda_{3,6}^{(k)} & 0
\end{array}\right]
\end{aligned}
$$

the solution $p^{*}(z)$ has (component-wise) degree

$$
\operatorname{deg}\left(p^{*}\right)<\left[\begin{array}{lllllll}
n & m & m & m & n & n & 1 \tag{15}
\end{array}\right]^{T}
$$

and satisfies

$$
\begin{equation*}
F_{k} p^{*}\left(\omega_{k}\right)=G_{k} p^{*}\left(\omega_{k}\right)=H_{k} p^{*}\left(\omega_{k}\right)=0 \tag{16}
\end{equation*}
$$

for $k=0, \ldots, N-1$. Equations in the form of (16) are known as tangential-interpolation conditions.

### 4.2 Tangential interpolation

After transformation, we can compute (2) by tangentially interpolating a polynomial vector from the conditions in (16). The interpolation can be calculated efficiently with the algorithm first described
in [20] and later improved in [21]. To introduce this method, we present several components of the theoretical framework of [19]. As our goal for this section is to provide an overview, we omit many technical details.

We begin by defining an algebraic context for the problem. The solution $p^{*}(z)$ is an element of $\mathbb{C}[z]^{7 \times 1}$, the space of all $7 \times 1$ vector polynomials with complex coefficients. The set $\mathbb{C}[z]^{7 \times 1}$ is a module, a more abstract form of a vector space. It is similar in nature to the vector space $\mathbb{C}^{7 \times 1}$, but its elements are vectors of complex-valued polynomials.

The vector polynomial $p^{*}(z)$ is not an arbitrary element of $\mathbb{C}[z]^{7 \times 1}$; it also satisfies the interpolation conditions in (14). We therefore limit our search for $p^{*}(z)$ to the much smaller set of elements with this property. To formally establish this set, we define the tangential-interpolation residual.

Definition 1 (Tangential-interpolation residual). The residual of a vector polynomial $q(z) \in \mathbb{C}[z]^{7 \times 1}$ relative to the interpolation conditions $\left\{F_{k}, G_{k}, H_{k}\right\}$ is

$$
(r(q))_{k} \equiv\left[\begin{array}{l}
F_{k} \\
G_{k} \\
H_{k}
\end{array}\right] q\left(\omega_{k}\right) .
$$

With this definition, the set of all vector polynomials satisfying the interpolation conditions is given as

$$
\mathcal{S}:=\left\{p(z) \in \mathbb{C}[z]^{7 \times 1}:(r(p))_{k}=\mathbf{0} \quad \forall k\right\} .
$$

This set is a submodule of $\mathbb{C}[z]^{7 \times 1}$, and characterizes the null space of the matrix in (12).
The linear equations of (12) form an underdetermined homogeneous system, and have an infinite number of solutions. Correspondingly, the submodule $\mathcal{S}$ is infinite, and not all of its elements are related to the solution $p^{*}(z)$. For example, $\left(z^{N}-1\right) q(z) \in \mathcal{S}$ for any element $q(z) \in \mathbb{C}[z]^{7 \times 1}$. To compute the interpolation, then, we need to differentiate $p^{*}(z)$ from the irrelevant elements of $\mathcal{S}$.

Fortunately, the known degree structure of $p^{*}(z)$ sets it apart. Since the vector $\boldsymbol{p}^{*}$ is the unique solution to (7), the only elements of $\mathcal{S}$ that satisfy (15) are those of the form $\alpha p^{*}(z)$, where $\alpha \in \mathbb{C}$. As a result, if we can find an element of $\mathcal{S}$ with the proper degree structure, the solution $p^{*}(z)$ may be calculated with simple scaling.

To analyze the degree structure of the elements of $\mathcal{S}$, we use a tool known as the $\tau$-degree [19]. ${ }^{5}$
Definition 2 ( $\tau$-degree). Let $\tau \in \mathbb{Z}^{7}$. The $\tau$-degree of a vector polynomial $p(z) \in \mathbb{C}[z]^{7 \times 1}$ is the integer $\delta \in \mathbb{Z}$ given by

$$
\begin{equation*}
\delta=\tau-\operatorname{deg}(p)=\max _{i}\left(\operatorname{deg}\left(p_{i}\right)-\tau_{i}\right), \tag{17}
\end{equation*}
$$

where $\operatorname{deg}(0)=-\infty$.
The $\tau$-degree is the maximum polynomial degree in a vector polynomial after each of its components have been "shifted" by some set amount. It is represented visually in Fig. 1. For different choices of $\tau$, the $\tau$-degree may be different (and the components determining the $\tau$-degree may vary as well).

As Fig. 1 illustrates, the $\tau$-degree is parametric; its value for a fixed $p(z)$ depends on the parameters $\left\{\tau_{i}\right\}$. This property allows the $\tau$-degree to reflect how closely an element of $\mathcal{S}$ matches $p^{*}(z)$ in degree

[^4]

Figure 1: A visual representation of the $\tau$-degree: (a) polynomial degrees of each component $p_{i}(z)$ of a $7 \times 1$ polynomial vector $p(z)$; (b) individual components of a sample $\tau$ vector; (c) the degrees of $p_{i}(z)$ shifted by the $\tau$ values. In (c), the $\tau$-degree is the maximum of the shifted degree values, and is represented by a dashed line.
structure. To illustrate, we can consider the $\tau$-degree in standard polynomial interpolation. Suppose we wish to determine the polynomial $u(z)$ of minimal degree that satisfies

$$
u\left(\omega_{k}\right)=f_{k} \quad \text { for } \quad k=1, \ldots, K
$$

Since there are $K$ interpolation points, $\operatorname{deg}(u)=K-1$. For a given polynomial $v(z)$ that satisfies the conditions, if $\tau=K-1$ and $\tau-\operatorname{deg}(v)=k$, there are three possibilities:

- $k<0: v(z)$ will not obey all conditions in general;
- $k=0: v(z)=u(z)$; or
- $k>0: v(z)$ is not of minimal degree.

Next, consider the vector case. Letting $\tau=\left[\tau_{i}\right]$, where

$$
\tau_{i}= \begin{cases}m-1 & i=2,3,4  \tag{18}\\ n-1 & i=1,5,6 \\ 0 & i=7\end{cases}
$$

it follows that $\tau-\operatorname{deg}\left(p^{*}\right)=0$, indicating that $p^{*}(z)$ has the desired degree structure. Note that it is not the specific values of $\left\{\tau_{i}\right\}$ that matter, but the relative values. For instance, if $\tilde{\tau}=\tau+\mathbf{1}$, then a $\tilde{\tau}$-degree $\delta=-1$ corresponds to the proper degree structure.

Since $p^{*}(z)$ is the only element of $\mathcal{S}$ that satisfies (15), computing the Tikhonov-regularized solution amounts to finding an element of $\mathcal{S}$ with $\tau$-degree $\delta=0$ for $\tau$ as in (18). To find such an element, we make use of a special algebraic property of $\mathcal{S}$ : it is free, meaning it has a basis.

Submodule bases play a role similar to their linear-algebraic counterparts, allowing elements of the submodule to be described through expansions. Namely, a set $B(z)=\left\{B^{(1)}(z), \ldots, B^{(k)}(z)\right\}$ is a basis
for $\mathcal{S}$ if, for every element $p(z) \in \mathcal{S}$, there are unique polynomials $\alpha_{i}(z) \in \mathbb{C}[z]$ such that $p$ can be written

$$
\begin{equation*}
p(z)=\sum_{i=1}^{k} \alpha_{i}(z) B^{(i)}(z) \tag{19}
\end{equation*}
$$

The $\left\{\alpha_{i}(z)\right\}$ serve as "expansion polynomials" and (19) as a basis expansion. As is the case for linear subspaces, any element of $\mathcal{S}$ is entirely (and uniquely) described by its $\alpha_{i}(z)$ for a chosen basis, and the number of bases is infinite.

By Theorem 3.1 of [19], a basis for $\mathcal{S}$ has exactly seven elements $B^{(j)}(z) \in \mathbb{C}[z]^{7 \times 1}$, which may be gathered into the basis matrix polynomial $B(z)=\left[B^{(1)}(z) \cdots B^{(7)}(z)\right]$. While we do not explore their theoretical properties, submodule bases are an essential component of our algorithm; we will determine $p^{*}(z)$ by constructing a basis for $\mathcal{S}$ that has an element with $\tau$-degree $\delta=0$. Once we build such a basis, the solution $p^{*}(z)$ will be immediate.

More specifically, to ensure that our basis has an element with $\tau$-degree $\delta=0$, we will construct what is known as a $\tau$-reduced basis. A $\tau$-reduced basis has elements that act "linearly independent" relative to the $\tau$-degree, meaning that linear combinations of the basis elements can not decrease the $\tau$ degree. In other words, if $\delta$ is the maximum $\tau$-degree among the elements of the basis, then there is no linear combination

$$
q(z)=\sum_{i} \alpha_{i}(z) B^{(i)}(z)
$$

such that $\tau-\operatorname{deg}(q)<\delta$ other than $\alpha_{i}(z)=0$.
Before describing how we construct a $\tau$-reduced basis, we close the section with two remarks. First, it is important to note that an arbitrary set of $\tau$-reduced elements of $\mathcal{S}$ is not necessarily a basis; it must also span $\mathcal{S}$. Second, we construct a $\tau$-reduced basis for $\mathcal{S}$ not out of necessity, but because we can devise an efficient algorithm to do so. Any efficient method for determining an element of $\mathcal{S}$ with $\tau$-degree $\delta=0$ would suffice.

### 4.3 Basis construction

In this section, we detail an algorithm for constructing a $\tau$-reduced basis corresponding to the tangentialinterpolation problem. We begin with the following theorem, which allows us to subdivide the process.

Theorem 1. Let $\sigma_{1}, \ldots, \sigma_{K}$ be interpolation nodes corresponding to the vectors $\phi_{1}, \ldots, \phi_{K} \in \mathbb{C}^{1 \times 7}$, where the conditions $\left\{\left(\sigma_{k}, \phi_{k}\right)\right\}$ are mutually distinct and $\phi_{k} \neq \mathbf{0}^{T} \forall k$. For some $1 \leq \kappa \leq K$ and $\tau_{K} \in \mathbb{Z}^{7 \times 1}$, let $B_{\kappa}(z) \in \mathbb{C}[z]^{7 \times 7}$ be a $\tau_{K}$-reduced basis corresponding to the interpolation data $\left\{\left(\overline{\sigma_{k}}, \phi_{k}\right): k=1, \ldots, \kappa\right\}$.

Denote $\delta_{i}=\tau_{K}-\operatorname{deg}\left(B_{\kappa}^{(i)}\right)$ for $i=1, \ldots, 7$, and define

$$
\tau_{\kappa \rightarrow K}=-\left[\begin{array}{lll}
\delta_{1} & \cdots & \delta_{7}
\end{array}\right]^{T}
$$

Let $B_{\kappa \rightarrow K}(z) \in \mathbb{C}[z]^{7 \times 7}$ be a $\tau_{\kappa \rightarrow K}$-reduced basis matrix corresponding to the interpolation data

$$
\left\{\left(\sigma_{k}, \phi_{k} B_{\kappa}\left(\sigma_{k}\right)\right): k=\kappa+1, \ldots, K\right\}
$$

Then $B_{K}(z)=B_{\kappa}(z) B_{\kappa \rightarrow K}(z)$ is a $\tau_{K}$-reduced basis matrix corresponding to the interpolation data

$$
\left\{\left(\sigma_{k}, \phi_{k}\right): k=1, \ldots, K\right\}
$$

Proof. See Van Barel and Bultheel [20], Theorem 3.
Theorem 1 provides a method of continually subdividing the interpolation problem into smaller subproblems. It is also the reason that we compute a $\tau$-reduced basis, as the main result does not hold without the bases being $\tau$-reduced. However, we still need to solve problems at the finest scale.

Consider a single interpolation condition, composed of a vector $\phi \in \mathbb{C}^{1 \times 7}$ and a node $\sigma$. Without loss of generality, let $\tau_{1}$ be the smallest value of $\tau$, and define $\mu_{i}=-\phi_{i} / \phi_{1}$. Then the polynomial matrix

$$
B(z)=\left[\begin{array}{cccc}
z-\sigma & \mu_{2} & \cdots & \mu_{7}  \tag{20}\\
0 & 1 & \cdots & 0 \\
\vdots & \vdots & \ddots & \vdots \\
0 & 0 & \cdots & 1
\end{array}\right]
$$

satisfies the tangential-interpolation condition

$$
\phi B(\sigma)=\mathbf{0}^{T}
$$

Since $\tau_{1}$ is the minimum $\tau$ value, $B(z)$ is a $\tau$-reduced basis for the submodule $\mathcal{S}_{1}$ defined by the single interpolation condition $(\sigma, \phi)$ (see [21]).

We can determine a basis for $\mathcal{S}$ in $\mathcal{O}\left(n^{2}\right)$ operations by processing the conditions serially with the single-point construction of (20) and by invoking Theorem 1. This process corresponds to the "fastonly" basis-construction algorithms of [17] and [21]. However, the $\mathcal{O}\left(n^{2}\right)$ complexity can be improved by exploiting the structure of the interpolation nodes. Since the nodes are roots of unity, we can reduce the number of required operations with the recursive "interleaving" data-splitting of Alg. 1. In this routine, TANINT is the serial basis-construction function, which for a single point amounts to computing (20).

```
Algorithm 1 Recursive tangential-interpolation algorithm for nodes \(\left\{\sigma_{k}\right\}\) and vectors \(\left\{\phi_{k}\right\}\).
    procedure \(B(z)=\operatorname{RecTANINT}\left(\left\{\sigma_{k}\right\},\left\{\phi_{k}\right\}\right)\)
        \(N=\operatorname{LENGTH}\left(\left\{\sigma_{k}\right\}\right)\)
        if \(N=1\) then
            \(B(z) \leftarrow \operatorname{TanINT}\left(\left\{\sigma_{k}, \phi_{k}\right\}\right)\)
        else
            \(B_{L}(z) \leftarrow \operatorname{REcTANINT}\left(\left\{\sigma_{2 k}, \phi_{2 k}\right\}\right)\)
            for \(k=1, \ldots, N / 2\) do
                    \(\phi_{2 k-1} \leftarrow \phi_{2 k-1} B_{L}\left(\sigma_{2 k-1}\right)\)
                end for
                \(B_{R}(z) \leftarrow \operatorname{REcTANINT}\left(\left\{\sigma_{2 k-1}, \phi_{2 k-1}\right\}\right)\)
                \(B(z) \leftarrow B_{L}(z) B_{R}(z)\)
        end if
    end procedure
```

To see how such a strategy is beneficial, we can estimate the total number of operations it requires. Let $C_{n}$ be the total cost of calling RectanInt on a set of $n>1$ points. Steps (1) and (3) require $C_{n / 2}$ operations by definition. In step (2), the evaluations $\left\{B_{L}\left(\sigma_{2 k-1}\right)\right\}$ must be computed and multiplied against the $\left\{\phi_{2 k-1}\right\}$. Since the $\left\{\sigma_{k}\right\}$ are roots of unity, $B_{L}(z)$ can be evaluated at the nodes $\left\{\sigma_{2 k-1}\right\}$ using an FFT of length $n / 2$. Once these values are obtained, there are a total of $p^{2} n / 2$ matrix-vector products
(where $p=7$ is the basis size). Therefore, the number of operations in step (2) is upper-bounded by $c_{1} n \log (n)$, where $c_{1}$ is a constant. Finally, step (4) requires a number of polynomial multiplications and additions that depends only on the basis size. Each multiplication can be computed with the FFT, but the FFT size depends on the degrees of the polynomials involved. By the nature of the basis construction, all of the polynomials in $B_{L}(z)$ and $B_{R}(z)$ must have degrees no greater than $n / 2$, since they are constructed from $n / 2$ conditions. Therefore, for another constant $c_{2}$, the number of operations needed to multiply $B_{L}(z) B_{R}(z)$ is bounded by $c_{2} n \log (n)$.

Factoring in these costs, we get the recursive cost formula

$$
C_{n}=2 C_{n / 2}+c n \log (n),
$$

where the constant $c$ is determined by the costs of steps (2) and (4). Since this expression is a recurrence, the overall cost of constructing the basis is given by the Master Theorem [22] as $\mathcal{O}\left(n \log ^{2} n\right)$. The key to replacing a factor of $n$ from the fast-only basis construction with a factor of $\log ^{2}(n)$ is our ability to evaluate $B_{L}\left(\sigma_{2 k-1}\right)$ in only $\mathcal{O}(n \log (n))$ operations with the FFT.

To conclude the section, we examine the degree structure of the resulting basis. Recall from Section 4.2 that we can obtain the solution to the Tikhonov-regularization problem if we can construct a $\tau$-reduced basis for $\mathcal{S}$ with a column having $\tau$-degree $\delta=0$. We now argue that this will indeed be the case.

For the general Tikhonov problem, there are a total of $3 N$ interpolation conditions. In each step of the algorithm, a single condition is used to increase the $\tau$-degree of exactly one column of the current basis - the column with lowest $\tau$-degree - by one. Beginning with no interpolation conditions, and specifying the starting basis as the identity matrix, the $\tau$-degree of the $j^{t h}$ column is $-\tau_{j}$. During each iteration, the column with lowest $\tau$-degree has its degree increased by one (while the remaining $\tau$-degrees are unchanged), and therefore the difference in $\tau$-degree between any two columns of the final basis can be at most one. Since the values of $\tau$ sum to $3 N-6$, the sum of the final $\tau$-degrees is $-(3 N-6)+3 N=6$. Since the maximum $\tau$-degree difference between any two columns is one, six of the columns will have $\delta_{j}=1$ while the seventh has $\delta_{j}=0$. Thus, our solution is guaranteed.

## 5 A superfast algorithm for Tikhonov regularization

Our Tikhonov-regularization algorithm can be summarized as follows. First, we decouple the Gramian matrices $G_{T}$ and $G_{L}$ by introducing the artificial variables $\sigma_{1}$ and $\sigma_{2}$, turning the original system into the Toeplitz-block system of (6). We then define the additional artificial variables $\left\{\gamma_{i}\right\}$, which allow us to replace the blocks of this system with partial-circulant matrices using extensions in the form of (3). The resulting partial-circulant-block system is given in (8).

Once the matrix has been extended, we transform it with a Fourier-block operator to obtain the diagonal-block system of (12). The coefficients of the diagonal matrices in this expression define tangentialinterpolation conditions, and can be calculated through $N$-point FFTs. Once these conditions are established, we use the divide-and-conquer basis-construction algorithm to build a $\tau$-reduced basis $B^{*}(z)$, with $\tau$ as in (18). A single column $j$ of the basis will have $\tau$-degree equal to zero, and the solution $\hat{x}$ is given by

$$
\hat{x}(z)=\frac{B_{1}^{*(j)}(z)}{B_{7}^{*(j)}}
$$

The basis-construction algorithm can be unstable if not implemented carefully. As in [8] and [10], we pivot the interpolation conditions to avoid multiplying by small $\mu_{i}$ values early in the construction. Without pivoting, numerical errors may propagate as the algorithm progresses.

In addition, we take precaution not to process any interpolation conditions for a subproblem that might cause precision errors in later subproblems. If processing an interpolation condition might result in a numerically unstable basis for a subproblem, that condition is skipped and included only after all remaining subproblems have been solved. The skipped conditions are marked as "difficult points," and our calculation is more robust if we treat them only at the conclusion of the algorithm. Since the difficult points are processed with the serial basis constructor, the algorithm's efficiency is a function of the number of difficult points we encounter. For this reason, it is important to avoid generating many difficult points; we shall further discuss this consideration in Section 6.

## 6 Implementation

In this section, we detail two previously unexplored implementation issues with the basis-construction algorithm. First, we discuss data partitioning and its potential to affect the number of points marked as difficult. Second, we consider the task of constructing a basis when the number of interpolation conditions is not a power of two, as this has the potential to drastically increase the required overhead. By addressing these issues, we can maintain a small overhead cost for all matrix sizes and further improve the accuracy of our algorithm.

### 6.1 Data partitioning

In the pseudoinversion algorithm of [10], the authors separate their basis construction into two stages, each of which uses a separate set of the interpolation conditions. Extending such an approach to the general Tikhonov problem amounts to computing the final basis as the product

$$
B^{*}(z)=B_{H}(z) B_{G}(z) B_{F}(z)
$$

where the bases $B_{H}(z), B_{G}(z)$, and $B_{F}(z)$ are constructed using the conditions $\left\{H_{k}\right\},\left\{G_{k}\right\}$, and $\left\{F_{k}\right\}$, respectively. This strategy has the potential to dramatically reduce the number of required calculations, as it allows for smaller basis sizes in the intermediate problems and some data-independent precomputation. Unfortunately, it also tends to generate many difficult points, making it inefficient in practice.

To illustrate, consider when the matrix $L$ is square, producing $N=2 n$ conditions $\left\{H_{k}\right\}$. In this scenario, we have the interpolation conditions

$$
\lambda_{3,3}^{(k)}=(-1)^{k} \quad \text { and } \quad \lambda_{3,6}^{(k)}=1
$$

When we subdivide the conditions during the construction of the basis $B_{H}(z)$, we find that

$$
\lambda_{3,3}^{(2 k)}=\lambda_{3,6}^{(2 k)}, \quad \text { and } \quad \lambda_{3,3}^{(2 k-1)}=-\lambda_{3,6}^{(2 k-1)}
$$

This collinearity is detrimental when we attempt to solve either of the two subproblems. Specifically, since $\sigma_{2}(z)$ and $\gamma_{3}(z)$ have the same degree structure and the same interpolation conditions, they become indistinguishable in the subproblems.

One way to ameliorate this difficulty is to modify the subdivision strategy. Rather than interleaving subdivisions, we propose the "paired-interleaving" of

$$
\begin{aligned}
H^{(1)} & =\left\{H_{1}, H_{2}, H_{5}, H_{6}, \cdots\right\} \\
H^{(2)} & =\left\{H_{3}, H_{4}, H_{7}, H_{8}, \cdots\right\}
\end{aligned}
$$

$$
\left\{\begin{array}{c}
F_{4 k}, G_{4 k}, H_{4 k} \\
F_{4 k+1}, G_{4 k+1}, H_{4 k+1}
\end{array}\right\} \quad\left\{\begin{array}{l}
F_{4 k+2}, G_{4 k+2}, H_{4 k+2} \\
F_{4 k+3}, G_{4 k+3}, H_{4 k+3}
\end{array}\right\}
$$

Figure 2: Subdivision process during an iteration of the basis construction. Each set of conditions is subdivided into two components using the paired-interleaving strategy. Corresponding components for each set of conditions are processed together.

This method will prevent collinearity in conditions corresponding to identity-matrix blocks. We can still evaluate the basis efficiently at the split nodes, since we have effectively subdivided the data by four while processing two sets at once. While paired interleaving requires twice as many FFTs to evaluate the basis for a subproblem, these FFTs are half of the size they would be for standard interleaving. As a result, the number of operations is essentially unchanged.

Paired interleaving can greatly reduce the number of difficult points generated, as it removes most collinearity problems. Since the difficult points are processed with the fast-only basis constructor, there is a drastic decrease in the number of operations necessary to construct the basis with this strategy. Additionally, since difficult points generate numerical instability, this approach tends to be less errorprone.

While paired interleaving helps, a large number of difficult points may still arise if the interpolation conditions for the three block rows are processed in isolation. Regardless of the subdivision strategy, constructing a basis using only the conditions $\left\{H_{k}\right\}$ amounts to attempting to solve the equation

$$
\sigma_{2}-L x=0
$$

without any knowledge of $\sigma_{2}$ or $x$. In the absence of additional information, there are not enough restrictions on possible solutions to yield a meaningful result. Accordingly, many difficult points will still be encountered.

Therefore, rather than process the conditions separately, we directly construct the basis from all sets of the interpolation conditions. To retain our ability to evaluate the bases efficiently, our subdivisions use a paired-interleaving split on each set of conditions, as illustrated in Fig. 2. While lacking the benefits of reduced basis size and pre-computability, this approach is usually much more efficient for the Tikhonov problems since it generates many fewer difficult points.

### 6.2 Data sizes

We have so far assumed that the algorithm will subdivide the conditions until arriving at a single interpolation point. As noted in [8], this is a poor strategy in practice. At some level, further subdivision results in a higher overhead cost than serial construction. Hence, our implementation calls the fast-only solver when the number of interpolation conditions for a given subproblem satisfies $K \leq N_{\text {lim }}$ for some specified $N_{\text {lim }}$. The optimal value of $N_{\text {lim }}$ is machine-dependent; we have empirically found that it tends to lie in the range of $256-512$ for our machines.

We have also implicitly assumed that the total number of conditions we wish to process is a power of two. If this is indeed the case, the subdivided data points $\left\{\sigma_{2 k-1}\right\}$ are roots of -1 , and we can compute $B_{L}\left(\sigma_{2 k-1}\right)$ using length- $N / 2$ FFTs. When $N$ is not a power of two, however, we require longer FFTs.

For instance, consider the case when $N$ is prime; there is no cancellation in the complex exponentials of $\left\{\sigma_{2 k-1}\right\}$. As a result, we require a full length- $N$ FFT to evaluate $B_{L}(z)$ at these nodes regardless of how many points there are in the subdivision. This increased overhead can cause the computation time to vary drastically across problem sizes.

To counteract this effect, we return to the circulant extensions of (3). For convenience, we developed our basis construction algorithm using 1-circulant extensions. However, we may instead choose to perform $k$-circulant extensions, where we choose $k \geq 0$ to allow us to compute shorter FFTs when evaluating $B_{L}(z)$. We have found empirically that setting the arbitrary entries in the extensions to zero when $k \geq 1$ can lead to very ill-conditioned problems. To avoid this, we choose them to be similar in value to the known matrix coefficients. While we have no theoretical results to justify such a tactic, it appears to be more stable for the problems we have tested.

The question of how to choose an optimal $k$ remains. One immediate option, following the lead of [8], is to choose $k$ such that $N$ is a power of two; we refer to this as the "power-of-two" method. The problem with such an approach is that it will cause the algorithm's complexity function to become highly quantized. In particular, the number of operations required for $2^{p}$ points will be significantly lower than it will be for $2^{p}+1$.

Instead, we propose an alternative which may also be applied to the algorithms of [8] and [10]. Since we would like each subdivision to produce cancellation in the complex exponentials of the roots of unity, our requirements are:

1. $N=2^{p} M$, where $M$ is an integer;
2. $3 M \leq N_{\text {lim }}$; and
3. $N \geq \tilde{N}$, where $\tilde{N}$ is the number of conditions we would have in each stage with 0 -circulant extensions.

The first condition guarantees cancellation in the complex exponential for the first $p$ subdivisions. The second condition guarantees that after $p$ subdivisions, the number of remaining points is sufficiently small that the fast-only basis constructor is called. The final condition assures that the extensions will indeed produce partial-circulant blocks. These conditions are devised for an interleaving data split; since we are using a paired-interleaving split in practice, the right side of the inequality in the second condition is replaced by $N_{\text {lim }} / 2$ to ensure that we can evaluate the nodes $\left\{\sigma_{4 k+i}\right\}$ efficiently.

Alg. 2 provides a simple method for choosing the extension size $k$ to satisfy these criteria. The algorithm is formulated for interleaving data splitting, and can be easily modified for paired-interleaving splitting by changing the while loop condition. It is also more efficient to force $M$ to be even; this ensures that none of the smallest FFT sizes are prime numbers, reducing the overhead cost of the FFT calls.

Fig. 3 plots the number of interpolation conditions processed for the power-of-two method versus our proposed method. For all data points in the curves, sufficient cancellation occurs to halve the required FFT length each time the nodes are subdivided. Since the number of conditions is lower, the complexity for our proposed method should increase more smoothly with the problem size. However, the actual gain in efficiency may depend on the FFT implementation.

```
ditions \(\tilde{N}\) and fast-only level \(N_{\text {lim }}\)
    procedure \(k=\) OptExtend \(\left(\tilde{N}, N_{\text {lim }}\right)\)
        \(M \leftarrow \tilde{N}\)
        \(M_{\text {tot }} \leftarrow 3 M\)
        \(p \leftarrow 0\)
        while \(M_{t o t}>N_{\text {lim }}\) do
            \(p \leftarrow p+1\)
            \(M \leftarrow\lceil M / 2\rceil\)
            \(M_{\text {tot }} \leftarrow 3 M\)
        end while
        \(k \leftarrow 2^{p} M-\tilde{N}\)
    end procedure
```

Algorithm 2 Routine for calculating the optimal circulant extension size for minimum number of con-

## 7 Numerical simulations

We now describe several experiments that demonstrate our algorithm's utility. First, we present the results of example problems with randomly generated matrices to illustrate that the algorithm's computational complexity increases as $\mathcal{O}\left(n \log ^{2} n\right)$. Next, we compare the tangential interpolator to an iterative Tikhonov solver based on the Conjugate Gradient (CG) method by expressing its execution time in terms of an equivalent number of CG iterations. Finally, we compute the regularized recovery of a low-frequency signal from its non-uniform Fourier samples.

Our algorithm is coded in C++ with MEX-function interfaces to MATLAB. All experiments were run on a 3.16 GHz Intel Core 2 Duo machine with 3.0 GB of RAM under the Ubuntu 12.04 LTS operating system and MATLAB R2012A. ${ }^{6}$ For each experiment, we used paired-interleaving data splitting and the extension strategy of Section 6.2.

### 7.1 Computational complexity

To confirm the asymptotic cost of our algorithm, we solved a large number of Tikhonov problems across a range of matrix sizes. Since an exact solution to the Tikhonov-regularized system $T x=b$ will not return a vector identical to $x$, we instead used our algorithm to solve systems of the form

$$
\left(G_{T}+G_{L}\right) x=y
$$

for known input vectors $x$. These systems are effectively equivalent to Tikhonov-regularization problems, as we could replace $y$ with $T^{H} T x$. However, by comparing our solutions

$$
\hat{x}=\left(G_{T}+G_{L}\right)^{-1} y
$$

with the known input vectors, we can analyze our algorithm's accuracy in applying the matrix $\left(G_{T}+G_{L}\right)^{-1}$ (see Section 7.2).

In each experiment, the system matrix $T$ and Tikhonov matrix $L$ were $n \times n$ Toeplitz matrices whose coefficients were drawn from a (complex) standard normal distribution, for a total of $4 n-2$ free parameters

[^5]

Figure 3: Number of conditions processed versus the minimal possible number of conditions for the power-of-two method and our proposed method of circulant extension, with $N_{\text {lim }}=256$. The "quantization" in the number of conditions processed with our proposed method is significantly milder than that of the power-of-two method.
per experiment. The value of $n$ was varied in even logarithmic steps between $2^{9}$ and $2^{15}$, with the resulting average execution times shown in Fig. 4(a). In the figure, the execution time is plotted as a function of the number of free parameters rather than the matrix side length.

To verify the algorithm's complexity, we compared the observed execution times to the theoretical behavior. Based on the algorithm's operation count, if the execution time is primarily a function of the number of operations it should be characterized by a function of the form

$$
E(n)=c_{1} n \log ^{2} n+c_{2} n \log n+\varepsilon(n)+\operatorname{ov}(n)
$$

where $c_{1}$ and $c_{2}$ are constants, $\varepsilon(n)$ consists of lower-order computations, and ov $(n)$ reflects differences in overhead cost for various problem sizes. For most problems, the $\mathcal{O}\left(n \log ^{2} n\right)$ calculations dominate the computational cost, but the contribution of $\mathcal{O}(n \log n)$ calculations is non-negligible for the problem sizes we consider. We can therefore make the approximation

$$
E(n) \approx c_{1} n \log ^{2} n+c_{2} n \log n
$$

Using this model, we computed a least-squares fit of the constants $c_{1}$ and $c_{2}$ and superimposed the estimated computational cost $E(n)$ on the observed calculation times in Fig. 4(a).

To compare how the execution time of the tangential interpolator is reduced for $\ell_{2}$-norm penalization and the Toeplitz-Gramian problem, we repeated the experiments for these cases. For $\ell_{2}$-norm penalization, we fixed the regularization parameter $\beta$ for each matrix size. Similarly, we fixed the main diagonal $a_{0}$ of $G_{T}$ for each of the Toeplitz-Gramian experiments. Neither of these choices affect the execution time; we have specified these values only to better control the matrix conditioning, allowing for more meaningful comparisons of accuracy (as will be explained in the next section). As a result of these choices, the number of free parameters for these problems are $2 n-1$ and $3 n-2$, respectively. In Fig. 4(b) and (c),


Figure 4: Execution time vs. total number of free parameters for (a) square Toeplitz matrices $T$ and $L$ with coefficients drawn from a complex standard normal distribution; (b) square Toeplitz matrices $T$ with coefficients drawn from a complex standard normal distribution and $\ell_{2}$-norm penalization; and (c) Toeplitz Gramian matrices $G_{T}$ and Tikhonov matrices $L$ with coefficients drawn from a complex standard normal distribution. For each data point, the execution time was averaged over 1000 trials. Also plotted in the figures are the least-squares estimates of the underlying complexity functions, which closely match the observed data.
the average execution times for these problems are plotted as functions of the number of free parameters, and are again compared to least-squares fits of the underlying cost function.

The execution-time curves in Fig. 4 suggest that our algorithm achieves the predicted asymptotic complexity. Moreover, these experiments give insight into the algorithm's overhead cost for the various problems. In each plot, the $i^{\text {th }}$ data point corresponds to systems of the same size. The $\ell_{2}$-normpenalization problem and the Toeplitz-Gramian problem each require less overhead, as their computation times are significantly smaller than the general problem for systems of the same size. This result is not surprising; the two special cases involve fewer interpolation conditions and smaller basis sizes than the general problem.

However, we can also consider the cost as a function of free parameters. In this sense, there is minimal difference in the interpolator's performance for the general problem and for $\ell_{2}$-norm penalization, but it is more efficient for the Toeplitz-Gramian problem. While the general problem requires more operations than $\ell_{2}$-norm penalization, it also contains roughly twice as many parameters. When considering cost as a function of free parameters, these two effects nearly "cancel out," and the execution time - as a function of the number of free parameters - is similar for the two problems. By contrast, when the Gramian $G_{T}$ is Toeplitz, the amount of calculation is the same as for a $\ell_{2}$-norm-penalization problem of the same size, while the number of free parameters increases to $3 n-2$. Therefore, the execution time as a function of the number of parameters drops by a factor of $2 / 3$ for the Toeplitz-Gramian problem.

### 7.2 Accuracy of results

In addition to confirming our algorithm's complexity, we used the experiments of Section 7.1 to verify that it is able to apply the inverse of the matrix $\left(G_{T}+G_{L}\right)$ with acceptable accuracy. For each experiment, we recorded the maximum error between the source vector $x$ and the recovered vector $\hat{x}$. We then took the maximum error across all trials for each type of problem and matrix size.

Table 1: Maximum error between the source vector and the solution returned from the inversion program for the experiments of Section 7.1. For each problem type and matrix size, the maximum errors in the recovered vectors were taken across all 1000 trials. "Matrix size" refers to the side length of the matrices for the experiments.

| Matrix <br> Size | General <br> Problem | $\ell_{2}$-Norm <br> Penalization | Toeplitz <br> Gramian |
| :---: | :---: | :---: | :---: |
| 512 | $9.86 \mathrm{E}-12$ | $1.55 \mathrm{E}-11$ | $3.73 \mathrm{E}-12$ |
| 1024 | $2.68 \mathrm{E}-11$ | $4.38 \mathrm{E}-11$ | $1.38 \mathrm{E}-11$ |
| 2048 | $7.56 \mathrm{E}-11$ | $1.34 \mathrm{E}-10$ | $3.62 \mathrm{E}-11$ |
| 4096 | $1.77 \mathrm{E}-10$ | $3.79 \mathrm{E}-10$ | $1.19 \mathrm{E}-10$ |
| 8192 | $4.46 \mathrm{E}-10$ | $1.13 \mathrm{E}-9$ | $3.18 \mathrm{E}-10$ |
| 16384 | $1.19 \mathrm{E}-9$ | $3.39 \mathrm{E}-9$ | $1.05 \mathrm{E}-9$ |
| 32768 | $2.88 \mathrm{E}-9$ | $1.07 \mathrm{E}-8$ | $3.04 \mathrm{E}-9$ |

While the tangential interpolator's execution time is seemingly independent of the matrix conditioning, its accuracy is not. As is to be expected, there are larger errors that propagate throughout the basis construction as the matrix $\left(G_{T}+G_{L}\right)$ becomes more poorly conditioned, resulting in less exact solutions. To address this, we fixed certain parameters in our $\ell_{2}$-norm-penalization and Toeplitz-Gramian experiments to ensure that the matrices generated for each of the three problems had similar condition numbers when all else was equal. The resulting data better reflects the interpolator's accuracy across the three problems when conditions are effectively the same.

We have empirically found that when the matrix entries are drawn from a standard normal distribution, the sum of the two Gramians $\left(G_{T}+G_{L}\right)$ tends to be fairly well conditioned for the general problem, even as the matrix size increases. However, for the $\ell_{2}$-norm-penalization problem, the conditioning of the single Gramian $G_{T}$ worsens with increasing size. To compensate, we set the regularization parameter $|\beta|^{2}=\sqrt{N}$. Similarly, the conditioning of the matrix $\left(G_{T}+L^{H} L\right)$ in the Toeplitz-Gramian problem appears to be largely a function of the diagonal dominance of $G_{T}$. We therefore set the main diagonal entry of $G_{T}$ to be $a_{0}=10 \sqrt{N}$. The values of $\beta$ and $a_{0}$ were both found through trial-and-error.

The maximum-error results are listed in Table 1, and suggest that our algorithm applies the inverse with acceptable accuracy for each problem. These results can be further improved with iterative refinement as described in [8] (though this is beyond the scope of this work). However, the numbers in Table 1 do not include any such adjustments, and report the errors after a single pass of the inversion process. All stabilization measures are intrinsic to the basis-construction algorithm and are factored into the computation times.

### 7.3 Equivalence in Conjugate Gradient Iterations

The results of Sections 7.1 and 7.2 give an absolute measure of the performance of the tangential interpolator. To gain a sense of perspective, we now translate this performance into a comparison with CG. Since CG is one of the most celebrated iterative methods for solving least-squares problems, it serves as a reasonable benchmark.

Unfortunately, a direct comparison with CG is difficult, as it is an iterative method while the tangential interpolator is not. Moreover, the convergence speed of CG is dependent on the conditioning of the matrix

Table 2: Number of CG iterations corresponding to the tangential-interpolation Tikhonov solver for the three different Tikhonov-regularization problems. The table values were calculated by averaging the equivalent number of iterations in each scenario over 1000 trials. "Matrix size" refers to the matrix side length for the experiments.

| Matrix <br> Size | General <br> Problem | $\ell_{2}$-Norm <br> Penalization | Toeplitz <br> Gramian |
| :---: | :---: | :---: | :---: |
| 512 | 70.0 | 54.2 | 41.9 |
| 1024 | 74.6 | 55.7 | 41.2 |
| 2048 | 84.3 | 60.0 | 43.8 |
| 4096 | 118.6 | 82.0 | 56.1 |
| 8192 | 126.9 | 84.7 | 58.3 |
| 16384 | 169.7 | 104.4 | 77.4 |
| 32768 | 240.2 | 123.8 | 86.6 |

(among other factors), resulting in variable solution times across different problems of the same size. By contrast, our algorithm is nearly static in computation time for a given matrix size, as the complexity is primarily dependent on the number of interpolation conditions. However, we can compare the relative efficiency of the two algorithms by determining how many iterations of CG can be performed in the same amount of time that our superfast solver requires.

We repeated the experiments of Section 7.1 for matrices with side-lengths of $n=2^{k}, k=9, \ldots, 15$. For each trial, the time required to compute the explicit inverse of $\left(G_{T}+G_{L}\right)$ using tangential interpolation was recorded. We then passed the matrix parameters to an implementation of CG to obtain the solution to the Tikhonov problem, stopping the program when the total elapsed time surpassed the direct-inversion time. When CG terminated, we recorded the number of completed iterations (discarding the results of any partial iterations), and averaged it across all trials. The resulting equivalent iteration counts are given in Table 2.

To arrive at a fair comparison, we implemented routines that allowed the CG solver to apply the matrix $\left(G_{T}+G_{L}\right)$ with minimal complexity. In each case, since either the matrices $T$ and $L$ or their Gramians are Toeplitz, the individual matrices can be applied with FFTs of length $3 n-2$. However, the minimum number of FFTs is different for each problem:

- General problem: nine - five to compute $T x$ and $L x$ and four to compute $T^{H}(T x)$ and $L^{H}(L x)$;
- $\ell_{2}$-norm penalization: five to compute $T^{H}(T x)$;
- Toeplitz-Gramian problem: seven - five to compute $G_{T} x$ and $L x$ and two to compute $L^{H}(L x)$.

In these tallies we have used the fact that the FFT of the generating vector for the matrix $T^{H}$ can be obtained in $\mathcal{O}(n)$ operations from the FFT of the generating vector for the matrix $T$ (and similarly for $\left.L^{H}\right)$.

As indicated in Table 2, the equivalent number of CG iterations increases with the matrix sidelength. This is to be expected; the CG iterations use $\mathcal{O}(n \log n)$ operations while our algorithm contains procedures requiring $\mathcal{O}\left(n \log ^{2} n\right)$ operations. Therefore, it is unsurprising that as $n$ increases, more CG iterations can be run in the same amount of time it takes for the tangential interpolator to build a solution.

Table 3: Maximum errors between the source vector and the returned solutions of both the CG method and the tangential-interpolation algorithm for the three Tikhonov problem types. For each problem type and matrix size, the maximum errors in the recovered vectors were computed across all trials. In all trials, the CG method was terminated after surpassing the time required for the inversion program to return a solution to the same problem. "Matrix size" refers to the side length of the matrices for the experiments.

| Matrix <br> Size | General Problem |  | $\ell_{2}$-Norm Penalization |  | Toeplitz Gramian |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | CG | Direct | CG | Direct | CG | Direct |
| 512 | $1.64 \mathrm{E}-7$ | $8.66 \mathrm{E}-12$ | $1.02 \mathrm{E}-3$ | $1.45 \mathrm{E}-11$ | $1.21 \mathrm{E}-4$ | $3.84 \mathrm{E}-12$ |
| 1024 | $1.61 \mathrm{E}-7$ | $2.52 \mathrm{E}-11$ | $5.26 \mathrm{E}-3$ | $4.45 \mathrm{E}-11$ | $4.18 \mathrm{E}-3$ | $1.30 \mathrm{E}-11$ |
| 2048 | $6.03 \mathrm{E}-8$ | $7.49 \mathrm{E}-11$ | $1.21 \mathrm{E}-2$ | $1.28 \mathrm{E}-10$ | $1.22 \mathrm{E}-3$ | $4.02 \mathrm{E}-11$ |
| 4096 | $1.14 \mathrm{E}-9$ | $1.70 \mathrm{E}-10$ | $1.22 \mathrm{E}-2$ | $4.00 \mathrm{E}-10$ | $1.48 \mathrm{E}-4$ | $1.29 \mathrm{E}-10$ |
| 8192 | $2.02 \mathrm{E}-10$ | $4.32 \mathrm{E}-10$ | $1.62 \mathrm{E}-2$ | $1.14 \mathrm{E}-9$ | $3.51 \mathrm{E}-3$ | $3.53 \mathrm{E}-10$ |
| 16384 | $2.28 \mathrm{E}-10$ | $1.08 \mathrm{E}-9$ | $5.85 \mathrm{E}-2$ | $3.51 \mathrm{E}-9$ | $2.20 \mathrm{E}-4$ | $1.01 \mathrm{E}-9$ |
| 32768 | $1.09 \mathrm{E}-14$ | $2.75 \mathrm{E}-9$ | $3.82 \mathrm{E}-2$ | $1.10 \mathrm{E}-8$ | $1.01 \mathrm{E}-2$ | $3.23 \mathrm{E}-9$ |

In addition, we may compare the accuracy of the two algorithms for the same execution time. For each experiment, we recorded the maximum error in the solutions returned by the two algorithms. We then took the maximum of these errors across all trials, with the results shown in Table 3.

Table 3 reflects the potential performance gains that can be realized with our algorithm. The tangential interpolator is only markedly outperformed by the CG solver when the matrices of the general problem become very large. For $\ell_{2}$-norm penalization and the Toeplitz-Gramian problem, CG requires much more time to achieve a comparable level of accuracy than the tangential interpolator even when the condition numbers are kept reasonable.

### 7.4 Non-Uniform Fourier Experiments

To demonstrate a practical use of our algorithm, we applied it to the task of reconstructing a signal from non-uniform samples of its spectrum. This is a common problem in signal processing, and one that is particularly relevant (albeit in a multidimensional variant) for magnetic resonance imaging (MRI). For one-dimensional signals, the spectrum of a discrete signal $x=\left[x_{s}\right]$ may be sampled at an arbitrary frequency $f_{0} \in[-1 / 2,1 / 2)$ by evaluating the sum

$$
X\left(f_{0}\right)=\sum_{s} x_{s} \mathbf{e}^{-\mathbf{j} 2 \pi f_{0} s}
$$

Collectively, a set of $K$ spectral samples at frequencies $\left\{f_{k}\right\}$ may be obtained by evaluating the matrixvector product $X=A x$, where $A$ is a Fourier-like matrix with entries

$$
A_{k s}=\mathbf{e}^{-\mathbf{j} 2 \pi f_{k} s}
$$

When $K=N$ and the $\left\{f_{k}\right\}$ are uniformly spaced in $[-1 / 2,1 / 2)$, the matrix $A$ is a Fourier matrix. Accordingly, the signal $x$ may be recovered from its spectral samples by

$$
A^{H} X=A^{H} A x=x
$$

However, when the frequencies are not uniformly spaced, the matrix $A$ is often severely ill-conditioned, and a regularized solution to the system $X=A x$ is required. The Tikhonov regularization for the problem is typically formulated as

$$
\hat{x}=\left(A^{H} W A+L^{H} L\right)^{-1} A^{H} W X
$$

where $W$ is a diagonal weighting matrix that compensates for the sampling density in the Fourier domain and $L$ is the regularizer. For MRI reconstruction problems, a Voronoi-cell weighting is typically used to compute the sampling density [23], yielding the entries of $W$.

Straightforward calculations show that the "weighted Gramian" $A^{H} W A$ is structured, with entries

$$
\left(G_{A}\right)_{r s}=\left(A^{H} W A\right)_{r s}=\sum_{k=1}^{K} w_{k k} \mathbf{e}^{\mathbf{j} 2 \pi f_{k}(r-s)} .
$$

Since the entries of $G_{A}$ depend only on the index difference $(r-s), G_{A}$ is Toeplitz. If the Tikhonov matrix $L$ is also Toeplitz, this amounts to the Toeplitz-Gramian problem, and our algorithm may be used to determine the closed-form solution to the Tikhonov-regularization problem.

To demonstrate, we generated length-4096 input signals consisting of linear combinations of sinusoids of three randomly-chosen frequencies in the digital frequency range $[0,0.02]$. For each input signal, we acquired 4096 spectral samples at random frequencies chosen from a triangular distribution over $[-1 / 2,1 / 2)$. As the number of Fourier samples is equal to the length of the input signals, the corresponding Gramian matrices $G_{A}$ were severely rank-deficient.

Since the input signals contained only low-frequency sinusoids, a scaled second-order difference matrix

$$
L_{r-s}= \begin{cases}-1 \mathrm{E}-4 & |r-s|=1 \\ 2 \mathrm{E}-4 & r-s=0 \\ 0 & \text { else }\end{cases}
$$

serves as a mild but effective regularizer, penalizing solutions with high-frequency content.
A typical reconstruction is shown in Fig. 5. In this example, the Gramian $G_{A}$ had a numerical rank of 3050 despite being $4096 \times 4096$, and the condition number of the matrix $\left(G_{A}+L^{H} L\right)$ was approximately 9.7 E 6 . Despite this somewhat poor conditioning, we were able to acquire a reasonable reconstruction through direct inversion in less than 0.7 seconds.

Our reconstruction may be compared to one achieved with CG. Again time-limiting the CG reconstruction based on the runtime of the tangential interpolator, we computed the iterative solution to the system. With both solutions available, we obtained the residual vector $x-\hat{x}$ for each method and plotted the results on the same scale in Fig. 6. While both methods achieve reasonable reconstructions, our tangential-interpolation algorithm outperforms CG in reconstruction quality for an equal amount of computation time.

## 8 Summary and future extensions

In this work, we have proposed an algorithm for solving three different Toeplitz-structured Tikhonovregularization problems with complexity $\mathcal{O}\left(N \log ^{2} N\right)$, where $N$ is the total number of free parameters. This algorithm solves a tangential-interpolation problem in place of a linear-algebraic problem, much like the superfast pseudoinversion algorithm of [10], and is based on the "extension-and-transformation" approach of [17]. Further, it is stabilized and non-iterative.

Nonuniform Fourier Reconstruction (4096 Samples)


Figure 5: A segment of the reconstruction of a low-frequency input signal using 4096 non-uniformly spaced Fourier samples via superfast Tikhonov regularization, where the Tikhonov matrix was a scaled second-order difference matrix. Despite the poor conditioning of the matrix, a reasonable reconstruction was produced in under a second.


Figure 6: Residuals of the reconstructed signals from our tangential-interpolation method and the CG method. Both methods yield a reasonable reconstruction in the same amount of time, with the tangential-interpolation method having a smaller residual.

We have demonstrated through a series of experiments that our implementation of the proposed algorithm produces accurate results and does so with a computational cost that closely matches the
theoretical asymptotic bound. While a direct comparison between the two algorithms is difficult, we were able to frame our results in terms of the number of iterations that a CG solver may perform in the same amount of time. In comparing the accuracy of the two solution methods for the same amount of time, our algorithm appears to be preferable across a wide range of scenarios. In addition, we showed that our algorithm may be used in practical settings by recovering a signal from its non-uniform spectral samples.

It is possible to extend our approach to multilevel Toeplitz matrices, which exhibit Toeplitz structure in multiple scales. The extension-and-transformation approach can be applied directly to such problems, but unfortunately requires a much higher asymptotic cost than desired. However, we have recently presented an efficient inversion algorithm for two-level Toeplitz matrices that exhibit triangularity in one or more of their levels [24]. In a future work, we will combine the theoretical results that form the basis of the inversion method with the regularization algorithm of this paper to produce a least-squares solver for this class of matrices.

## References

[1] N. Levinson, "The Wiener RMS error criterion in filter design and prediction," J. Math. Phys., vol. 25, pp. 261-278, 1947.
[2] J. Durbin, "The fitting of time series models.," Rev. Inst. Int. Stat., vol. 28, pp. 233-243, 1960.
[3] W.F. Trench, "An algorithm for the inversion of finite Toeplitz matrices," J. SIAM, vol. 12, pp. 512-522, 1964.
[4] S. Zohar, "Toeplitz matrix inversion: the algorithm of W. F. Trench," J. ACM, vol. 16, no. 4, pp. 592-601, Oct. 1969.
[5] E.H. Bareiss, "Numerical solution of linear equations with Toeplitz and vector Toeplitz matrices," Num. Math., vol. 13, pp. 404-424, 1969.
[6] G. Heinig and K. Rost, "Fast algorithms for Toeplitz and Hankel matrices," Linear Algebra Appl., vol. 435, pp. 1-59, 2011.
[7] G.S. Ammar and W.B. Gragg, "Superfast solution of real positive definite Toeplitz systems," SIAM J. Matrix Anal. Appl., vol. 9, pp. 61-76, 1988.
[8] M. Van Barel, G. Heinig, and P. Kravanja, "A stabilized superfast solver for nonsymmetric Toeplitz systems," SIAM J. Matrix Anal. Appl., vol. 23, pp. 494-510, 2001.
[9] S. Chandrasekaran, M. Gu, X. Sun, J. Xia, and J. Zhu, "A superfast algorithm for Toeplitz systems of linear equations," SIAM J. Matrix Anal. Appl., vol. 29, no. 4, pp. 1247-1266, 2007.
[10] M. Van Barel, G. Heinig, and P. Kravanja, "A superfast method for solving Toeplitz linear least squares problems," Linear Algebra Appl., vol. 366, pp. 441-457, 2003.
[11] G. Heinig, "Fast algorithms for Toeplitz least squares problems," Operator Theory: Advances and Applications, vol. 149, pp. 167-197, 2004.
[12] D.R. Sweet, "Fast Toeplitz orthogonalization," Num. Math., vol. 43, pp. 1-21, 1984.
[13] D.A. Bini, G. Codevico, and M. Van Barel, "Solving Toeplitz least squares problems by means of Newton's iteration," Numer. Algorithms, vol. 33, pp. 93-103, 2003.
[14] V.Y. Pan, M. Van Barel, Xinmao Wang, and G. Codevico, "Iterative inversion of structured matrices," Theor. Comput. Sci., vol. 315, pp. 581-592, 2004.
[15] T. Kailath, S. Kung, and M. Morf, "Displacement ranks of matrices and linear equations," J. Math. Anal. Appl., vol. 68, no. 2, pp. 395-407, 1979.
[16] S.K. Kimitei, Algorithms for Toeplitz Matrices with Applications to Image Deblurring, Mathematics theses, Georgia State University, April 2008.
[17] G. Heinig, "Solving Toeplitz systems after extension and transformation," Calcolo, vol. 33, pp. 115-129, 1996.
[18] I.P. Fechina, "Tangential Nevanlinna-Pick problem with multiple points (in Russian)," Doklady Akad. Nauk Arm. SSR, vol. 61, pp. 214-218, 1975.
[19] M. Van Barel and A. Bultheel, "A general module theoretic framework for vector M-Padé and matrix rational interpolation," Num. Alg., vol. 3, pp. 451-461, 1992.
[20] M. Van Barel and A. Bultheel, The "look-ahead" philosophy applied to matrix rational interpolation problems, vol. Systems and Networks: Mathematical Theory and Applications, Vol. II: Invited and Contributed Papers of Mathematical Research 79, pp. 891-894, Akademie-Verlag, Berlin, 1994.
[21] M. Van Barel and P. Kravanja, "A stabilized superfast solver for indefinite Hankel systems," Linear Algebra Appl., vol. 284, pp. 335-355, 1998.
[22] T.H. Cormen, C.E. Leiserson, R.L. Rivest, and C. Stein, Introduction to Algorithms, MIT Press and McGraw-Hill, second edition, 2001.
[23] V. Rasche, R. Proksa, R. Sinkus, P. Börnert, and H. Eggers, "Resampling of data between arbitrary grids using convolution interpolation," IEEE T. Med. Imaging, vol. 18, pp. 385-392, 1999.
[24] C. Turnes, D. Balcan, and J. Romberg, "Image deconvolution via superfast inversion of a class of two-level Toeplitz matrices," Proc. IEEE ICIP, pp. 3073-3076, 2012.


[^0]:    ${ }^{*}$ C. Turnes and J. Romberg are with the School of Electrical and Computer Engineering at the Georgia Institute of Technology. Doru Balcan is with the School of Interactive Computing at the Georgia Institute of Technology. This work was partially supported by a Packard Fellowship.

[^1]:    ${ }^{1}$ The code package is available at http://users.ece.gatech.edu/~cturnes3.
    ${ }^{2}$ The term "ridge regression" is used in the field of statistics.

[^2]:    ${ }^{3}$ For the Toeplitz-Gramian problem, we need only decouple the matrix $G_{L}$.

[^3]:    ${ }^{4}$ It is "artificial" because it adds no new information to the system.

[^4]:    ${ }^{5}$ In its original form (in [19]), the $\tau$-degree was referred to as the $\vec{s}$-degree. This terminology was later changed in [8] and [10].

[^5]:    ${ }^{6}$ The code for these experiments can be obtained from http://users.ece.gatech.edu/~cturnes3/software.html.

