NATIONAL INSTITUTE FOR APPLIED STATISTICS RESEARCH AUSTRALIA


# National Institute for Applied Statistics Research Australia 

University of Wollongong, Australia

## Working Paper

$$
01-23
$$

Computational Implementation of Supernodal Variants for Cholesky Factorization and Calculation of the SIS

Luke Mazur and Robin Thompson

# Computational implementation of supernodal variants for Cholesky factorization and calculation of the SIS 

Luke Mazur, Robin Thompson

September 4, 2023

## 1. The AI algorithm

This chapter briefly discusses the various quantities required in the calculation of the AI algorithm (Gilmour et al., 1995).

### 1.1 Introducing the general linear mixed model

If the $n_{y}$-vector of observations is denoted by $\boldsymbol{y}$, the linear mixed model can be written as

$$
\begin{equation*}
\boldsymbol{y}=\boldsymbol{X} \boldsymbol{\tau}+\boldsymbol{Z} \boldsymbol{u}+e \tag{1.1}
\end{equation*}
$$

where $\boldsymbol{\tau}$ is the $n_{\boldsymbol{\tau}}$-vector of fixed effects, $\boldsymbol{X}$ is the $n_{\boldsymbol{y}} \times n_{\boldsymbol{\tau}}$ design matrix that associates observations with the appropriate combination of fixed effects, $\boldsymbol{u}$ is the $n_{\boldsymbol{u}}$-vector of random effects, $\boldsymbol{Z}$ is the $n_{\boldsymbol{y}} \times n_{\boldsymbol{u}}$ design matrix which associates observations with the appropriate combination of random effects, and $\boldsymbol{e}$ is the $n_{\boldsymbol{y}}$-vector of residuals. The column rank of $\boldsymbol{X}$ is denoted $d_{\boldsymbol{X}}$ and $d_{\boldsymbol{X}} \leq n_{\boldsymbol{\tau}}$. The model, in Equation (1.1), is called a linear mixed model or linear mixed-effects model. It is assumed

$$
\left[\begin{array}{l}
\boldsymbol{u}  \tag{1.2}\\
\boldsymbol{e}
\end{array}\right] \sim N\left(\left[\begin{array}{l}
\mathbf{0} \\
\mathbf{0}
\end{array}\right],\left[\begin{array}{cc}
\boldsymbol{G} & \mathbf{0} \\
\mathbf{0} & \boldsymbol{R}
\end{array}\right]\right) .
$$

The variance models given by the matrices $\boldsymbol{G}$ and $\boldsymbol{R}$ are called $G$-structures and $R$-structures respectively. In this report the only the uni-variate case is considered, however extensions to the multi-variate case are straightforward by replacing $\boldsymbol{y}, \boldsymbol{\tau}, \boldsymbol{u}$ and $\boldsymbol{e}$ by matrix equivalents, where each column corresponds to a response variable.

Typically $\boldsymbol{G}$ and $\boldsymbol{R}$ are functions of parameters that need to be estimated

$$
\begin{equation*}
G=\boldsymbol{G}\left(\sigma_{G}\right) \tag{1.3}
\end{equation*}
$$

and

$$
\begin{equation*}
\boldsymbol{R}=\boldsymbol{R}\left(\sigma_{R}\right) \tag{1.4}
\end{equation*}
$$

The vectors $\boldsymbol{\sigma}_{\boldsymbol{G}}$ and $\boldsymbol{\sigma}_{\boldsymbol{R}}$ are parameter vectors associated with the random effects $(\boldsymbol{u})$ and residuals $(\boldsymbol{e})$ respectively.

Then a matrix $\boldsymbol{V}$ is defined such that

$$
V=R+Z G Z^{\top}
$$

so that

$$
\begin{equation*}
\boldsymbol{y} \sim N(\boldsymbol{X} \boldsymbol{\tau}, \boldsymbol{V}) \tag{1.5}
\end{equation*}
$$

### 1.1.1 Variance structures for the residuals

In many datasets the vector of residuals represents the residuals from a single experiment and it is assumed that $\boldsymbol{R}$ is simply a scaled identity matrix, that is $\boldsymbol{R}=\sigma \boldsymbol{I}_{n_{\boldsymbol{y}}}$, where $\sigma$ is the first and only element of the parameter vector $\boldsymbol{\sigma}_{\boldsymbol{R}}$ implying the the residuals were assumed independent and identically distributed.

In the datasets of interest in this report, that is MET datasets (Smith et al., 2001), the vector $\boldsymbol{e}$ is a series of vectors indexed by a factor or factors. The subvectors relate to sections of the data, each with their own associated variance matrix. Thus in general it is written $\boldsymbol{e}=\left(\boldsymbol{e}_{1}^{\top}, \boldsymbol{e}_{2}^{\top}, \ldots, \boldsymbol{e}_{s}^{\top}\right)^{\top}$, so that $\boldsymbol{e}_{j}$ represents the vector of residuals of the $j^{\text {th }}$ section of the data. The variance matrix for each section may differ, but it is assumed that the residuals from different sections are independent (if they are not then the dependent components can be coalesced into a single component and hence the independence structure can be maintained).

In matrix terms this gives

$$
\boldsymbol{R}=\oplus_{j=1}^{s} \boldsymbol{R}_{j}
$$

where $\oplus$ is the direct sum operator. In MET datasets usually the indexing factor for a section is either Trial or Environment.

### 1.1.2 Variance structures for the random effects

The vector of random effects is often composed of $\phi$ sub-vectors $\boldsymbol{u}=\left(\boldsymbol{u}_{1}^{\top}, \boldsymbol{u}_{2}^{\top}, \ldots, \boldsymbol{u}_{\phi}^{\top}\right)^{\top}$ where the sub-vectors $\boldsymbol{u}_{i}$ are of length $n_{\boldsymbol{G}_{i}}$. These sub-vectors are assumed independent normally distributed with variance matrices $\boldsymbol{G}_{i}$. Thus, as for $\boldsymbol{R}$, this leads to

$$
\boldsymbol{G}=\oplus_{i=1}^{\phi} \boldsymbol{G}_{i}
$$

There is a corresponding partition in $\boldsymbol{Z}$, namely $\boldsymbol{Z}=\left[\boldsymbol{Z}_{1}, \boldsymbol{Z}_{2}, \ldots, \boldsymbol{Z}_{\phi}\right]$.

### 1.2 Prediction in the linear mixed model

The general aim is the prediction of the linear combination $\boldsymbol{\zeta}_{\mathbf{1}}^{\top} \boldsymbol{\tau}+\boldsymbol{\zeta}_{\mathbf{2}}^{\top} \boldsymbol{u}$ of fixed and random effects where $\boldsymbol{\zeta}_{1}$ and $\boldsymbol{\zeta}_{2}$ are known $n_{\tau^{-}}$and $n_{u^{\prime}}$-vectors respectively, and further $\boldsymbol{\zeta}_{1}^{\top} \boldsymbol{\tau}$ is estimable, i.e. $\boldsymbol{\zeta}_{1}^{\top}$ lies in the row space of $\boldsymbol{X}$ (Searle \& Khuri,

2017, p. 266). If $\boldsymbol{V}$ is known, implying $\boldsymbol{\sigma}_{G}$ and $\boldsymbol{\sigma}_{\boldsymbol{R}}$ are known, the predictor which has the minimum mean square error (MSE) among the class of linear unbiased predictors is given by $\boldsymbol{\zeta}_{1}^{\top} \hat{\boldsymbol{\tau}}+\boldsymbol{\zeta}_{\mathbf{2}}^{\top} \tilde{\boldsymbol{u}}$ where

$$
\begin{align*}
& \hat{\boldsymbol{\tau}}=\left(\boldsymbol{X}^{\top} \boldsymbol{V}^{-1} \boldsymbol{X}\right)^{-} \boldsymbol{X}^{\top} \boldsymbol{V}^{-1} \boldsymbol{y}  \tag{1.6}\\
& \tilde{\boldsymbol{u}}=\boldsymbol{G} \boldsymbol{Z}^{\top} \boldsymbol{P} \boldsymbol{y} \tag{1.7}
\end{align*}
$$

$\boldsymbol{P}=\boldsymbol{V}^{-1}-\boldsymbol{V}^{-1} \boldsymbol{X}\left(\boldsymbol{X}^{\top} \boldsymbol{V}^{-1} \boldsymbol{X}\right)^{-} \boldsymbol{X}^{\top} \boldsymbol{V}^{-1}$ and $\left(\boldsymbol{X}^{\top} \boldsymbol{V}^{-1} \boldsymbol{X}\right)^{-}$is a generalised inverse of $\boldsymbol{X}^{\top} \boldsymbol{V}^{-1} \boldsymbol{X}$ (Henderson, 1950).

These can be obtained by solving a system of equations known as the mixed model equations (MMEs), as proposed by Henderson (1950, 1973). They can be written in matrix-vector notation by:

$$
\left[\begin{array}{cc}
\boldsymbol{X}^{\top} \boldsymbol{R}^{-1} \boldsymbol{X} & \boldsymbol{X}^{\top} \boldsymbol{R}^{-1} \boldsymbol{Z}  \tag{1.8}\\
\boldsymbol{Z}^{\top} \boldsymbol{R}^{-1} \boldsymbol{X} & \boldsymbol{Z}^{\top} \boldsymbol{R}^{-1} \boldsymbol{Z}+\boldsymbol{G}^{-1}
\end{array}\right]\left[\begin{array}{c}
\hat{\boldsymbol{\tau}} \\
\tilde{\boldsymbol{u}}
\end{array}\right]=\left[\begin{array}{c}
\boldsymbol{X}^{\top} \boldsymbol{R}^{-1} \boldsymbol{y} \\
\boldsymbol{Z}^{\top} \boldsymbol{R}^{-1} \boldsymbol{y}
\end{array}\right]
$$

Letting $\boldsymbol{W}=[\boldsymbol{X} \boldsymbol{Z}], \tilde{\boldsymbol{\beta}}^{\top}=\left[\hat{\boldsymbol{\tau}}^{\top} \tilde{\boldsymbol{u}}^{\top}\right], \boldsymbol{b}=\boldsymbol{W}^{\top} \boldsymbol{R}^{-1} \boldsymbol{y}$,

$$
\begin{aligned}
\boldsymbol{C} & =\boldsymbol{W}^{\top} \boldsymbol{R}^{-1} \boldsymbol{W}+\boldsymbol{G}^{*} \\
\boldsymbol{G}^{*} & =\left[\begin{array}{cc}
\mathbf{0} & \mathbf{0} \\
\mathbf{0} & \boldsymbol{G}^{-1}
\end{array}\right],
\end{aligned}
$$

leads to a more abbreviated representation of the MMEs

$$
\begin{equation*}
C \tilde{\boldsymbol{\beta}}=b \tag{1.9}
\end{equation*}
$$

It can be shown that a specific solution to the MMEs is in fact equivalent to the estimates $\hat{\boldsymbol{\tau}}$ and $\tilde{\boldsymbol{u}}$ in Equations (1.6) and (1.7), known as empirical BLUEs (E-BLUEs) and empirical BLUPs (E-BLUPs) respectively.

### 1.3 The REML log-likelihood function

The REML log-likelihood (Patterson \& Thompson, 1971) for Equation (1.1) can be written as

$$
\begin{equation*}
\ell_{\boldsymbol{R}}=-\frac{1}{2}\left\{2 \log |\boldsymbol{F}|-\log \left|\boldsymbol{X}_{\mathbf{0}}^{\top} \boldsymbol{X}_{\mathbf{0}}\right|+\log |\boldsymbol{V}|+\log \left|\boldsymbol{X}_{\mathbf{0}}^{\top} \boldsymbol{V}^{-1} \boldsymbol{X}_{\mathbf{0}}\right|+\boldsymbol{y}^{\top} \boldsymbol{P}_{\mathbf{0}} \boldsymbol{y}\right\}, \tag{1.10}
\end{equation*}
$$

where the choice of $\boldsymbol{F}=\boldsymbol{I}_{\left(n_{\boldsymbol{y}}-\boldsymbol{d}_{\boldsymbol{X}}\right)}$ and $\boldsymbol{X}_{\mathbf{0}}$ refers to the linearly independent columns of $\boldsymbol{X}$, see Mazur (2022) for more information on that.

### 1.4 Score equations and AI updating

Letting $\dot{\boldsymbol{R}}_{\boldsymbol{i}}=\partial \boldsymbol{R} / \partial \sigma_{\boldsymbol{R}_{\boldsymbol{i}}}$ and $\boldsymbol{q}_{\boldsymbol{R}_{\boldsymbol{i}}}=\dot{\boldsymbol{R}}_{\boldsymbol{i}} \boldsymbol{P}_{\mathbf{o}} \boldsymbol{y}$, the REML score for a variance parameter $\sigma_{\boldsymbol{R}_{\boldsymbol{i}}}$ associated with the residuals $\boldsymbol{e}$, following Gilmour et al. (1995), is given by

$$
\begin{align*}
\boldsymbol{U}_{\boldsymbol{R}}\left(\sigma_{\boldsymbol{R}_{\boldsymbol{i}}}\right) & =-\frac{1}{2}\left\{\operatorname{tr}\left(\boldsymbol{R}^{-1} \dot{\boldsymbol{R}}_{i}\right)-\operatorname{tr}\left(\boldsymbol{C}_{\mathbf{0}}^{-1} \boldsymbol{W}_{\mathbf{0}}^{\top} \boldsymbol{R}^{-1} \dot{\boldsymbol{R}}_{\boldsymbol{i}} \boldsymbol{R}^{-1} \boldsymbol{W}_{\mathbf{0}}\right)-\right. \\
& \left.\boldsymbol{y}^{\top} \boldsymbol{P}_{\mathbf{0}} \boldsymbol{q}_{\boldsymbol{R}_{\boldsymbol{i}}}\right\} \tag{1.11}
\end{align*}
$$

where $\boldsymbol{C}_{\mathbf{0}}$ refers to $\boldsymbol{C}$ with all instances of $\boldsymbol{X}$ replaced by $\boldsymbol{X}_{\mathbf{0}}$. Letting $\dot{\boldsymbol{G}}_{\boldsymbol{i}}=$ $\partial \boldsymbol{G} / \partial \sigma_{\boldsymbol{G}_{\boldsymbol{i}}}, \boldsymbol{C}^{\boldsymbol{Z}}$ be the partition of $\boldsymbol{C}_{\mathbf{0}}^{-1}$ which corresponds to $\boldsymbol{u}$ and $\boldsymbol{q}_{\boldsymbol{G}_{\boldsymbol{i}}}=$ $\boldsymbol{Z} \dot{\boldsymbol{G}}_{\boldsymbol{i}} \boldsymbol{Z}^{\top} \boldsymbol{P}_{\mathbf{0}} \boldsymbol{y}$, the REML score for a variance parameter $\sigma_{\boldsymbol{G}_{\boldsymbol{i}}}$ is given by

$$
\begin{align*}
& \boldsymbol{U}_{\boldsymbol{R}}\left(\sigma_{\boldsymbol{G}_{i}}\right)=-\frac{1}{2}\left\{\operatorname{tr}\left(\boldsymbol{G}^{-1} \dot{\boldsymbol{G}}_{i}\right)-\operatorname{tr}\left(\boldsymbol{G}^{-1} \dot{\boldsymbol{G}}_{i} \boldsymbol{G}^{-1} \boldsymbol{C}^{\boldsymbol{Z} \boldsymbol{Z}}\right)-\right. \\
&\left.\boldsymbol{y}^{\top} \boldsymbol{P}_{\mathbf{0}} \boldsymbol{q}_{\boldsymbol{G}_{i}}\right\} . \tag{1.12}
\end{align*}
$$

The elements of the average information matrix are given by:

$$
\begin{equation*}
\boldsymbol{I}_{\boldsymbol{A}}\left(\kappa_{i}, \kappa_{j}\right)=\frac{1}{2}\left(\dot{\boldsymbol{V}}_{\boldsymbol{i}} \boldsymbol{P}_{\mathbf{0}} \boldsymbol{y}\right)^{\top} \boldsymbol{P}_{\mathbf{0}}\left(\dot{\boldsymbol{V}}_{\boldsymbol{j}} \boldsymbol{P}_{\mathrm{o}} \boldsymbol{y}\right) . \tag{1.13}
\end{equation*}
$$

## 2. Cholesky factorization and calculation of the sparse inverse subset by recursion

### 2.1 Defining the CF and the SIS

Typically $\boldsymbol{C}$ is a sparse matrix, meaning it has many zero elements, and the goal is to exploit this as much as possible to minimize runtime by avoiding redundant operations involving multiplication by zeroes. This report concentrates on the two most computationally intensive steps of the algorithm, that is the Cholesky factorization (CF) used in solving the MMEs and the calculation of the sparse inverse subset (SIS) which is the set of elements of the inverse of $\boldsymbol{C}$ necessary for the scores, i.e. the elements of $\boldsymbol{C}_{\mathbf{0}}^{-1}$ corresponding to non-zeroes in $\boldsymbol{C}_{\mathbf{0}}$. A CF typically refers to the calculation of a lower triangular matrix $\boldsymbol{L}(\boldsymbol{L}$ is the convention in the literature as L stands for Lower) such that $\boldsymbol{L} \boldsymbol{L}^{\top}=\boldsymbol{C}$. In the algorithms present in this report, following Gilmour et al. (1995) and Meyer (1989), an operator Factorize is defined such that $\boldsymbol{L}=\operatorname{Factorize}(\boldsymbol{A})$ denotes finding a lower triangular matrix $\boldsymbol{L}$ such that $\boldsymbol{L}^{\top} \boldsymbol{L}=\boldsymbol{A}$ for some SPD matrix $\boldsymbol{A}$. This could be considered an "inversely permuted" CF, and is simply referred to as a CF where the meaning is clear from context. This can be viewed as a form of extension of the $\sqrt{ }$ operator to matrices, and if Factorize is applied to a $1 \times 1$ matrix then it behaves as the $\sqrt{ }$ operator. For example, if $\boldsymbol{A}$ and $\boldsymbol{L}$ are $1 \times 1$ matrices then $\boldsymbol{L}=\operatorname{Factorize}(\boldsymbol{A})$ is equivalent to $l_{11}=\sqrt{a_{11}}$ in this case.
CF of $\boldsymbol{C}$ can be used to transform a system of equations:

$$
\begin{equation*}
C \tilde{\boldsymbol{\beta}}=b \tag{2.1}
\end{equation*}
$$

to a triangular system:

$$
\begin{equation*}
\boldsymbol{L} \tilde{\boldsymbol{\beta}}=\left(\boldsymbol{L}^{\top}\right)^{-1} \boldsymbol{b}=\boldsymbol{h} \tag{2.2}
\end{equation*}
$$

The calculation of $\boldsymbol{h}$ is known as the back solution and this could occur after the CF is complete, but it is shown how it can be performed as part of the CF process.

## 2．2 Recursion using multi－nodes

The recursion occurs by considering a multi－node at each level of the recursion． The lower triangle of a symmetric matrix $\boldsymbol{A}$ can be partitioned such that：

$$
\boldsymbol{A}=\left[\begin{array}{ccc}
\boldsymbol{A}_{11} & * & * \\
\vdots & \ddots & \vdots \\
\boldsymbol{A}_{m 1} & \ldots & \boldsymbol{A}_{m m}
\end{array}\right]
$$

Then it can be partitioned into $m$ multi－nodes such that the $i^{\text {th }}$ multi－node consists of all blocks of $\boldsymbol{A}$ in the $i^{\text {th }}$ block row，i．e．$\left[\boldsymbol{A}_{i 1} \boldsymbol{A}_{i 2} \ldots \boldsymbol{A}_{i i}\right]$ ．Each multi－ node consists of two parts：a diagonal part（of which only the lower triangle is stored）and an off－diagonal part．The diagonal part of the $i^{\text {th }}$ multi－node is simply the diagonal block $\boldsymbol{A}_{i i}$ ，while the off－diagonal part refers to all of the off－diagonal blocks，i．e． $\boldsymbol{A}_{i, 1: i-1}$ ．The first multi－node only has a diagonal part i．e．the off－diagonal part of the first multi－node has no columns．A visual representation of the first two multi－nodes is shown in Figure 2．1．

Figure 2．1：A visual representation of the top two multi－nodes in terms of the elements that need to be stored．First multi－node（red）with diagonal part （yellow）and off－diagonal part（blue）of second multi－node．


In particular，the rows of $\boldsymbol{C}$ and $\boldsymbol{L}$ can be grouped into multi－nodes with $m$ denoting the number of such multi－nodes．Then $\left\{\mathfrak{r}_{i}\right\}$ denotes the set of rows in the $i^{\text {th }}$ multi－node，numbering $\left|⿷_{i}\right|$ ．Further $[i]$ is denoted as a multi－node comprised of the first $i$ multi－nodes and $\left\{\mathfrak{r}_{[i]}\right\}=\cup\left\{\mathfrak{r}_{j}\right\}, j \leq i$ ．Then the number of rows associated with the first $i$ multi－nodes is denoted by $\left|⿷_{[i]}\right|$（so $\left\{\mathfrak{r}_{[m]}\right\}$ refers to all the rows and $\left|⿷_{[m]}\right|=n$ is the total number of rows of $\left.\boldsymbol{C}\right)$ ．As a result，$\left\{\mathbb{r}_{i}\right\}=\left\{\left|⿷_{[i-1]}\right|+1:\left|\mathfrak{r}_{[i]}\right|\right\}$ Similarly $\left\{\mathbb{C}_{i}\right\}$ denotes the columns of the off－ diagonal part of the $i^{\text {th }}$ multi－node，numbering $\left|\mathbb{C}_{i}\right|$ ．The columns of the diagonal part correspond to $\left\{\mathfrak{r}_{i}\right\}$ ．
There are two main advantages of multi－nodes．The first is that the off－diagonal part can be treated as a dense rectangular matrix．The second is that the diagonal part of a multi－node of $\boldsymbol{C}$ can be treated as a dense symmetric matrix while the diagonal part of a multi－node of $\boldsymbol{L}$ is treated as a dense triangular matrix．
The recursive process is initialized such that：

$$
\begin{align*}
\boldsymbol{C}^{(m)} & =\boldsymbol{C}  \tag{2.3}\\
\boldsymbol{b}^{(m)} & =\boldsymbol{b} \tag{2.4}
\end{align*}
$$

To denote the step of the recursion, $k$ is used. Finally, in order to save memory in computer implementations, a separate object for $\boldsymbol{L}$ is not created. Instead, $\boldsymbol{C}$ is overwritten by its Cholesky factor, and then its SIS. This occurs using the assignment operator $:=$ as the recursion moves ahead.
Additionally, there are situations where $\boldsymbol{C}$ of size $n \times n$ is stored in contiguous memory but operations are only required on $C_{:\{\mathbb{A}\}\{\mathbb{B}\}}$ : of size $|\mathbb{A}| \times|\mathbb{B}|$ which is not stored in contiguous memory, nor can it be treated as contiguous through the use of a fixed offset. Such a memory structure that is unsuitable for BLAS operations must be copied into a memory structure which is suitable, and this is denoted by the subscript of the partition being pre- and post-ceded by ":" symbols. Then for the sake of argument say it is desired to apply a function Function to $\boldsymbol{C}_{:\{\mathbb{A}\}\{\mathrm{B}\} \text { : }}$ but because it is not stored in contiguous memory they are instead applied to $\boldsymbol{A}=\boldsymbol{C}_{:\{\mathrm{A}\}\{\mathrm{B}\}:}$. This requires gathering $\boldsymbol{C}_{:\{\mathrm{A}\}\{\mathrm{B}\}}$ : into the shallow copy $\boldsymbol{A}$, performing the function (or some other operation) and then scattering $\boldsymbol{A}$ back into the relevant locations of $\boldsymbol{C}_{:\{\mathbb{A}\}\{\mathbb{B}\}:}$ :
The function Gather gives a name to the targeted copying of rows of $C$ contained in $\{\mathbb{A}\}$ and columns of $\boldsymbol{C}$ contained in $\{\mathbb{B}\}$ into $\boldsymbol{A}$. Letting $\left\{\mathbb{A}_{i}\right\}$ be a set containing only the $i^{\text {th }}$ element of $\{\mathbb{A}\}$ such that $\{\mathbb{A}\}=\cup\left\{\mathbb{A}_{i}\right\}$ and similarly letting $\left\{\mathbb{B}_{j}\right\}$ be a set containing only the $j^{\text {th }}$ element of $\{\mathbb{B}\}$. Then:

$$
\boldsymbol{A}_{i j}=\boldsymbol{C}_{\left\{\mathbb{A}_{i}\right\}\left\{\mathbb{B}_{j}\right\}}
$$

Similarly the function Scatter gives a name to the copying of $\boldsymbol{A}$ back into $\boldsymbol{C}$ such that:

$$
\boldsymbol{C}_{\left\{\mathbb{A}_{i}\right\}\left\{\mathbb{B}_{j}\right\}}:=\boldsymbol{A}_{i j} .
$$

The symbol $:=$ : is used to refer to this procedure of gathering, assigning to and scattering back (if a term on the right hand side is not stored in a convenient memory structure this does not pose a problem, it is simply gathered and used as necessary without a change in notation) and so equivalently:

$$
C_{:\{\mathbb{A}\}\{\mathbb{B}\}:}:=: \text { Function }\left(C_{:\{\mathbb{A}\}\{\mathbb{B}\}}\right) \text {. }
$$

### 2.3 Sparsity and supernodes

To incorporate sparsity, sets $\left\{\mathbb{Z}_{\boldsymbol{i}}\right\}$ are defined such that $a \in\left\{\mathbb{Z}_{i}\right\}$ if $a<i$ and $c_{i a}=$ 0 after the CF is complete, i.e. the zeroes in the $i^{\text {th }}$ row of the Cholesky factor, with $i$ taking values from 1 to $n$. Similarly sets $\left\{\mathbb{C}_{i}\right\}$ are defined such that $a \in\left\{\mathbb{C}_{i}\right\}$ if $a<i$ and $c_{i a} \neq 0$, i.e. the non-zeroes in the $i^{\text {th }}$ row of the Cholesky factor. Supernodes are multi-nodes with the constraint that all of the rows in the off-diagonal part of a supernode have the same non-zero structure. These supernodes are discovered in a matrix using a supernodal symbolic factorization which merges rows into sets $\left\{\mathfrak{r}_{k}\right\}$, as well as the sets of non-zero columns for the supernodes denoted by $\left\{\mathbb{C}_{k}\right\}$ and the sets of zero columns for the supernodes denoted by $\left\{\mathbb{Z}_{\boldsymbol{k}}\right\}$. These are calculated using an elimination tree, see Liu
(1990). Then another set of merging occurs through the supernodal symbolic factorization such that if $\left\{\mathbb{z}_{\boldsymbol{d}}\right\}$ is sufficiently similar to $\left\{\mathbb{z}_{(d+1)}\right\}$ then these two supernodes are merged together, even if this means explicitly storing some zero entries, as it was found to reduce runtime in the majority of cases (Ashcraft \& Grimes, 1989).

Prior to this symbolic factorization, typically a permutation is used to minimize the number of structural non-zero entries in $\boldsymbol{L}$, by reducing the number of fillin that occurs as zeroes are added to/subtracted from during the CF process: turning them into non-zeroes. After this permutation, a post-ordering occurs which permutes the children of each node in the elimination tree prior to its parents. This ordering can be computed using a recursive depth first search, where in each step it does a depth first search through all of the children of the current node, until it encounters leaf nodes (nodes with no children). Then it moves its way back up the stack of recursion, numbering the nodes in the order they were visited, i.e. with children being numbered before the parents. The procedure starts at the root, and the root is given its number once all of its children have been visited, i.e. at the end. This post-ordering has no impact on the number of non-zeroes in $\boldsymbol{L}$ but it means that children are typically placed consecutively, and these are the most likely to be able to be merged into supernodes as they are likely to have similar non-zero structures. In the following it is assumed that both of these orderings have been applied to $\boldsymbol{C}$ prior to the CF.

The code written by the author uses the supernodal amalgamation rules of CHOLMOD (Chen et al., 2008) (an open source alternative discussed in Section 5.4 ), as it uses CHOLMOD for the entire symbolic factorization phase. The rules are that two supernodes are merged if any of the following conditions:

- The number of rows in the merged supernode is less than some parameter (by default 4)
- No new structural non-zeroes are formed by merging them
- The number of rows in the merged supernode is less than some parameter (by default 16) and the ratio of stored explicit zeroes to total structural non-zeroes in the merged supernode is less than some parameter (by default 0.8)
- The number of rows in the merged supernode is less than some parameter (by default 48) and the ratio of stored explicit zeroes to total structural non-zeroes in the merged supernode is less than some parameter (by default 0.1)
- The ratio of stored explicit zeroes to total structural non-zeroes in the merged supernode is less than some parameter (by default 0.05)
are met.

Then Equation (2.1) can be partitioned into:

$$
\left[\begin{array}{cc}
\boldsymbol{C}_{\{\mathbb{\mathbb { T }}}^{(k)\}\left\{\mathbb{T}^{(k)}\right\}} & \boldsymbol{C}_{\left\{\mathbb{T}^{(k)}\right\}\left\{\mathbb{B}^{(k)}\right\}}^{(k)}  \tag{2.5}\\
\boldsymbol{C}_{\left\{\mathbb{B}^{(k)}\right\}\left\{\mathbb{T}^{(k)}\right\}}^{(k)} & \boldsymbol{C}_{\left\{\mathbb{B}^{(k)}\right\}\left\{\mathbb{B}^{(k)}\right\}}^{(k)}
\end{array}\right]\left[\begin{array}{c}
\tilde{\boldsymbol{\beta}}_{\left\{\mathbb{T}^{(k)}\right\}} \\
\tilde{\boldsymbol{\beta}}_{\left\{\mathbb{B}^{(k)}\right\}}
\end{array}\right]=\left[\begin{array}{c}
\boldsymbol{b}_{\left\{\mathbb{T}^{(k)}\right\}}^{(k)} \\
\boldsymbol{b}_{\left\{\mathbb{B}^{(k)}\right\}}^{(k)}
\end{array}\right],
$$

with at the $k^{\text {th }}$ step of the algorithm $\left\{\mathbb{T}^{(k)}\right\}=\left\{\mathfrak{a}_{[k-1]}\right\},\left\{\mathbb{B}^{(k)}\right\}=\left\{\mathfrak{r}_{k}\right\}$. Then defining the following sets $\left\{\mathbb{Z}^{(k)}\right\}=\left\{\mathbb{Z}_{\left\{\mathfrak{r}_{k}\right\}}\right\}$, $\left\{\mathbb{S}^{(k)}\right\}=\left\{\mathbb{C}_{k}\right\}=\left\{\mathbb{T}^{(k)}\right\}-\left\{\mathbb{Z}^{(k)}\right\}$, the relevant elements at the $k^{\text {th }}$ step can be accessed using $\left\{\mathscr{S}^{(k)}\right\}$ and $\left\{\mathbb{B}^{(k)}\right\}$. In the case where the matrix is entirely dense then there is only one supernode with $\left\{n_{1}\right\}=\{1: n\}$, and this is equivalent to the multi-nodal case with one multi-node.
A direct consequence of supernodal methods is that $\boldsymbol{C}_{\left\{\mathbb{T}^{(k)}\right\}\left\{\mathbb{T}^{(k)}\right\}}$ is no longer stored, as many of the elements within it would be zero. Rather, $\boldsymbol{C}_{\left[\mathbb{F}^{(k-1)}\right]}$ is stored such that at the $k^{\text {th }}$ step of the algorithm $\left[\mathbb{F}^{(k-1)}\right]=\cup\{a, b\}$ for $a \in$ $\left\{\mathbb{r}_{i}\right\}, b \in\left\{\mathbb{C}_{i}\right\}$ for $i<k$, i.e. it is the union of all the elements present in the first $k-1$ supernodes. Similarly letting $\left[\mathbb{F}^{(k)}\right]=\cup\{a, b\}$ for $a \in\left\{\mathfrak{r}_{i}\right\}, b \in$ $\left\{\mathbb{C}_{i}\right\}$ for $i<=k$, i.e. it is the union of all the elements present in the first $k$ supernodes.

## 3. Function notation

### 3.1 Cholesky factorization of the coefficient matrix of the MMEs

First $\boldsymbol{C}_{\left\{\mathbb{B}^{(k)}\right\}\left\{\mathbb{B}^{(k)}\right\}}$ is factorized then $\boldsymbol{C}_{\left\{\mathbb{B}^{(k)}\right\}\left\{\mathbb{S}^{(k)}\right\}}$ and $\boldsymbol{b}_{\left\{\mathbb{B}^{(k)}\right\}}$ are back-solved for and finally $\boldsymbol{C}_{:\left\{\mathscr{S}^{(k)}\right\}\left\{\mathbb{S}^{(k)}\right\}:}$ and $\boldsymbol{b}_{:\left\{\mathcal{S}^{(k)}\right\}:}$ are updated using the back-solutions, where the Schur complement is a specific case of a back-update.

$$
\begin{aligned}
& C_{\left\{\mathbb{B}^{(k)}\right\}\left\{\mathbb{B}^{(k)}\right\}}:=\text { Factorize }\left(\boldsymbol{C}_{\left\{\mathbb{B}^{(k)}\right\}\left\{\mathbb{B}^{(k)}\right\}}\right), \\
& C_{\{\mathbb{B}(k)\}\left\{S^{(k)}\right\}}:=\left(\left(\boldsymbol{C}_{\{\mathbb{B}(k)\}\{\mathbb{B}(k)\}}\right)^{-1}\right)^{\top} C_{\{\mathbb{B}(k)\}\left\{S^{(k)}\right\}}, \\
& \boldsymbol{b}_{\left\{\mathbb{B}^{(k)}\right\}}:=\left(\left(\boldsymbol{C}_{\left\{\mathbb{B}^{(k)}\right\}\left\{\mathbb{B}^{(k)}\right\}}\right)^{-1}\right)^{\top} \boldsymbol{b}_{\left\{\mathbb{B}^{(k)}\right\}}, \\
& C_{:\left\{S^{(k)}\right\}\left\{S^{(k)}\right\}}:=: C_{:\left\{S^{(k)}\right\}\left\{S^{(k)}\right\}}:- \\
& \left(C_{\left\{\mathbb{B}^{(k)}\right\}\left\{S^{(k)}\right\}}\right)^{\top} C_{\left\{\mathbb{B}^{(k)}\right\}\left\{S^{(k)}\right\}}, \\
& \boldsymbol{b}_{:\left\{S^{(k)}\right\}}:=: \boldsymbol{b}_{:\left\{S^{(k)}\right\}}=-\left(\boldsymbol{C}_{\{\mathbb{B}(k)\}\left\{S^{(k)}\right\}}\right)^{\top} \boldsymbol{b}_{\left\{\mathbb{B}^{(k)}\right\}},
\end{aligned}
$$

$$
\begin{align*}
& C_{\left\{\mathbb{B}^{(k)}\right\}\left\{\mathbb{B}^{(k)}\right\}}:=\text { Factorize }\left(\boldsymbol{C}_{\left\{\mathbb{B}^{(k)}\right\}\left\{\mathbb{B}^{(k)}\right\}}\right),  \tag{3.1}\\
& C_{\left\{\mathbb{B}^{(k)}\right\}\left\{\mathrm{S}^{(k)}\right\}}:=\operatorname{Back}-\operatorname{Solve}\left(C_{\left\{\mathbb{B}^{(k)}\right\}\left\{\mathbb{B}^{(k)}\right\}}, C_{\left\{\mathbb{B}^{(k)}\right\}\left\{\mathrm{S}^{(k)}\right\}}\right),  \tag{3.2}\\
& \boldsymbol{b}_{\{\mathbb{B}(k)\}}:=\operatorname{Back}-\operatorname{Solve}\left(\boldsymbol{C}_{\{\mathbb{B}(k)\}\left\{\mathbb{B}^{(k)}\right\}}, \boldsymbol{b}_{\{\mathbb{B}(k)\}}\right), \\
& C_{:\left\{S^{(k)}\right\}\left\{S^{(k)}\right\}:}:=\text { SchurComplement }\left(C_{:\left\{S^{(k)}\right\}\left\{S^{(k)}\right\}:,} C_{\left\{\mathbb{B}^{(k)}\right\}\left\{S^{(k)}\right\}}\right) \text {, }  \tag{3.3}\\
& \boldsymbol{b}_{:\left\{S^{(k)}\right\}}:=: \text { Back-UpdateSCP }\left(\boldsymbol{b}_{:\left\{\mathcal{S}^{(k)}\right\}}, \boldsymbol{C}_{\{\mathbb{B}(k)\}\left\{S^{(k)}\right\}}, \boldsymbol{b}_{\{\mathbb{B}(k)\}}\right) \text {. } \tag{3.4}
\end{align*}
$$

### 3.2 Calculation of the sparse inverse subset of the coefficient matrix of the MMEs

The first Assignment forward-solves for $\boldsymbol{A}$, while the second multiplies $\boldsymbol{C}_{\left\{\mathbb{B}^{(k)}\right\}\left\{\mathbb{S}^{(k)}\right\}}$ by a symmetric matrix. After this $\boldsymbol{C}_{\left\{\mathbb{B}^{(k)}\right\}\left\{\mathbb{B}^{(k)}\right\}}$ is inverted then multiplied by its transpose and finally $\boldsymbol{C}_{\left\{\mathbb{B}^{(k)}\right\}\left\{\mathbb{B}^{(k)}\right\}}$ is forward-updated.

$$
\begin{array}{rlrl}
\boldsymbol{A} & :=\left(\boldsymbol{C}_{\left\{\mathbb{B}^{(k)}\right\}\left\{\mathbb{B}^{(k)}\right\}}\right)^{-1} \boldsymbol{C}_{\left\{\mathbb{B}^{(k)}\right\}\left\{\mathbb{S}^{(k)}\right\}}, \\
\boldsymbol{C}_{\left\{\mathbb{B}^{(k)}\right\}\left\{\mathbb{S}^{(k)}\right\}} & :=-\boldsymbol{A} \boldsymbol{C}_{:\left\{\mathbb{S}^{(k)}\right\}\left\{\mathbb{S}^{(k)}\right\}:} \\
\boldsymbol{C}_{\left\{\mathbb{B}^{(k)}\right\}\left\{\mathbb{B}^{(k)}\right\}} & :=\left(\boldsymbol{C}_{\left\{\mathbb{B}^{(k)}\right\}\left\{\mathbb{B}^{(k)}\right\}}\right)^{-1}, \\
\boldsymbol{C}_{\left\{\mathbb{B}^{(k)}\right\}\left\{\mathbb{B}^{(k)}\right\}} & :=\boldsymbol{C}_{\left\{\mathbb{B}^{(k)}\right\}\left\{\mathbb{B}^{(k)}\right\}}\left(\boldsymbol{C}_{\left\{\mathbb{B}^{(k)}\right\}\left\{\mathbb{B}^{(k)}\right\}}\right)^{\top}, \\
\boldsymbol{C}_{\left\{\mathbb{B}^{(k)}\right\}\left\{\mathbb{B}^{(k)}\right\}}: & =\boldsymbol{C}_{\left\{\mathbb{B}^{(k)}\right\}\left\{\mathbb{B}^{(k)}\right\}}-\boldsymbol{C}_{\left\{\mathbb{B}^{(k)}\right\}\left\{\mathbb{S}^{(k)}\right\}} \boldsymbol{A}^{\top}, & \boldsymbol{C}_{\left\{\mathbb{B}^{(k)}\right\}\left\{\mathbb{S}^{(k)}\right\}}:=\operatorname{SymmetricMultiply}\left(-\boldsymbol{A}, \boldsymbol{C}_{:\left\{\mathbb{S}^{(k)}\right\}\left\{\mathbb{S}^{(k)}\right\}:}\right),  \tag{3.8}\\
\boldsymbol{C}_{\left\{\mathbb{B}^{(k)}\right\}\left\{\mathbb{B}^{(k)}\right\}}:=\text { InvertCrossMultiply}\left(\boldsymbol{I}, \boldsymbol{C}_{\left\{\mathbb{B}^{(k)}\right\}\left\{\mathbb{B}^{(k)}\right\}}\right), \\
\left.\boldsymbol{C}_{\left\{\mathbb{B}^{(k)}\right\}\left\{\mathbb{B}^{(k)}\right\}}:=\text { Forward-UpdateInv( } \boldsymbol{C}_{\left\{\mathbb{B}^{(k)}\right\}\left\{\mathbb{B}^{(k)}\right\}}, \boldsymbol{C}_{\left\{\mathbb{B}^{(k)}\right\}\left\{\mathbb{S}^{(k)}\right\}}, \boldsymbol{A}\right) .
\end{array}
$$

There is an alternate formulation worth considering (after re-ordering Assignments (3.5-3.8) and initializing $\boldsymbol{S}:=\boldsymbol{C}_{\left\{\mathbb{B}^{(k)}\right\}\left\{\mathbb{B}^{(k)}\right\}}, \boldsymbol{A}:=\boldsymbol{C}_{\left\{\mathbb{B}^{(k)}\right\}\left\{\mathbb{S}^{(k)}\right\}}$ and $\left.\boldsymbol{C}_{\left\{\mathbb{B}^{(k)}\right\}\left\{\mathbb{B}^{(k)}\right\}}:=\boldsymbol{I}\right)$ which first multiplies $\boldsymbol{C}_{\left\{\mathbb{B}^{(k)}\right\}\left\{\mathbb{S}^{(k)}\right\}}$ by a symmetric matrix then $\boldsymbol{C}_{\left\{\mathbb{B}^{(k)}\right\}\left\{\mathbb{B}^{(k)}\right\}}$ is forward-updated. After this $\boldsymbol{C}_{\left\{\mathbb{B}^{(k)}\right\}\left\{\mathbb{S}^{(k)}\right\}}$ is forward solved for and finally $\boldsymbol{C}_{\left\{\mathbb{B}^{(k)}\right\}\left\{\mathbb{B}^{(k)}\right\}}$ is pre- and post-multiplied by the inverse of $\boldsymbol{S}$ :


## 4. Singularity-handling supernodal AI algorithm

### 4.1 Literature review

While CF is typically thought to only work for SPD matrices, it can work for SPSD (Symmetric Positive Semi-Definite) matrices in two main ways. The first is known as a modified CF (Bunch, 1971; Schnabel \& Eskow, 1990; Eskow \& Schnabel, 1991) where small numbers are added to the diagonal elements to effectively make the matrix SPD. A block approach to this is discussed in Dayde \& Camichel (1995). If this matrix was the coefficient matrix for a linear system then such a modification modifies the solution of the system and so will not be discussed further herein.

The second way is to simply avoid the rows that have diagonal entries below a certain threshold in some way. There are two methods to avoid singularities when solving a linear system where the coefficient matrix is SPSD, without affecting the solution. In the first method they are set to extremely large numbers with off-diagonals set to zero and corresponding elements of the solution set to zero, essentially "fixing" them so that they do not cause numerical instability. In fact numerically they will have no impact on the remainder of the operations. This method is denoted the "zeroing-out" method. The second method is to mark them and avoid them entirely, such that they take no part in any future computation, denoted the "avoiding" method. The first method is discussed in detail in Gupta et al. (1997); Gupta (2000).
Of the two methods the avoiding should be more efficient as redundant operations are avoided, and this is the choice method in the dense linear algebra space, and in some sparse single-nodal codes such as ASReml. This method has also been applied for sparse unsymmetric nondefinite matrices and sparse symmetric indefinite matrices, but these problems are different as the tiny pivots cannot immediately be flagged as singularities, as it is possible that a negative pivot further in the matrix will modify the tiny pivot back to a suitable size. As a result, dynamic permutation must occur of the matrix in order to delay this pivot, and this typically causes fill-in to occur as it represents a departure from
the original ordering scheme. As far as the author is aware there has not been an implementation that avoids these singularities in a supernodal algorithm for the factorization of an SPSD matrix, and this report presents such an algorithm that while it features dynamic permutation, it does not lead to an increase in fill-in, as rows within a supernode can be interchanged without introducing any fill-in.

Since the software ASReml uses a CF that can handle singularities for every iteration, a supernodal CF for an SPSD is demonstrated herein through an adjustment to the Factorize operator, in Subsection 4.1.1. However, typically in the AI algorithm the singularities are only present in $\boldsymbol{C}_{\boldsymbol{X} \boldsymbol{X}}$ due to aliasing in the fixed effects and so such a CF could simply be applied to this partition of $\boldsymbol{C}$, which is relatively small for most models used for plant variety trials. The singularities could be detected, then a full rank parameterization could be selected and used for all remaining iterations since the singularities are typically the same for each iteration as they are independent of the variance parameters. The time for this CF would be a relatively minor portion of the time required for fitting most models, then the AI algorithm could proceed with an SPD $C$ matrix. Nevertheless, other applications such as chemical physics or linear programming feature matrices that can have singularities throughout the coefficient matrix, and only require one CF to be performed in their respective algorithms. In these applications having a faster method to factorize an SPSD matrix could potentially be quite useful.

For a more comprehensive review of the Cholesky factorization of SPSD matrices see Mazur (2022).

### 4.1.1 Singularity-Handling Factorization Operator

The operator Factorize ${ }^{+}$is defined to refer to a Factorize operation that can handle a singular input matrix. While for example Factorize $\left(\boldsymbol{C}_{\{\mathbb{B}(k)\}\{\mathbb{B}}{ }^{(k)\}}\right)$ only modifies $\boldsymbol{C}_{\left\{\mathbb{B}^{(k)}\right\}\left\{\mathbb{B}^{(k)}\right\}}$,
Factorize ${ }^{+}\left(\boldsymbol{C}_{\left\{\mathbb{B}^{(k)}\right\}\left\{\mathbb{B}^{(k)}\right\}}\right)$ may modify either $\boldsymbol{C}_{\left\{\mathbb{B}^{(k)}\right\}\left\{\mathbb{B}^{(k)}\right\}}$ to make it nonsingular (as in the zeroing-out method), or modify $\left\{\mathbb{B}^{(k)}\right\}$ to create a $\left\{\mathbb{B}^{(\mathbb{k})+}\right\}$ to ignore the linearly dependent rows of $\boldsymbol{C}_{\left\{\mathbb{B}^{(k)}\right\}\left\{\mathbb{B}^{(k)}\right\}}$ (as in the avoiding method). This is achieved by permuting all of the linearly dependent rows in the supernode to the top of the supernode such that $\left\{\mathbb{B}^{(\mathbb{k})+}\right\}$ remains a contiguous set, and keeping track of their number denoted as $\zeta_{\{\mathbb{B}(k)\}}$. As a result $d_{\{\mathbb{B}(k)\}}=\left|\mathfrak{r}_{k}\right|-\zeta_{\{\mathbb{B}(k)\}}$ denotes the rank of $\boldsymbol{C}_{\left\{\mathbb{B}^{(k)}\right\}\left\{\mathbb{B}^{(k)}\right\}}$. In either case Factorize ${ }^{+}$denotes finding an $\boldsymbol{L}_{\left\{\mathbb{B}^{(\mathbb{K})+\}\left\{\mathbb{B}^{(\mathbb{K})+}\right\}}\right.}=$ Factorize $^{+}\left(\boldsymbol{C}_{\left\{\mathbb{B}^{(k)}\right\}\left\{\mathbb{B}^{(k)}\right\}}\right)$ such that $\left(\boldsymbol{L}_{\left\{\mathbb{B}^{(\mathbb{k})+\}\left\{\mathbb{B}^{(\mathbb{K})+}\right\}}\right.}\right)^{\top} \boldsymbol{L}_{\left\{\mathbb{B}^{(\mathbb{K})+\}\{\mathbb{B}}(\mathbb{K})+\right\}}=$ $C_{\left\{\mathbb{B}^{(\mathbb{K})+}{ }_{\}\left\{\mathbb{B}^{(k)}+\right\}}\right.}$, sans potential modification of the linearly dependent rows, and $\left\{\mathbb{B}^{(\mathbb{k})+}\right\}=\left\{\mathbb{B}^{(k)}\right\}$ for the zeroing-out method (and thus $\zeta_{\left\{\mathbb{B}^{(k)}\right\}}=0$ and $\left.d_{\left\{\mathbb{B}^{(k)}\right\}}=\left|\mathbb{B}^{(\mathbb{k})}\right|\right)$. In the SPD case Factorize and Factorize ${ }^{+}$achieve the same result.

Further bookkeeping is required for the avoiding method, as all of the permutations are applied only within each supernode locally during the CF, and once the CF has been completed for all supernodes, they can be "synchronised" such that they all share the new matrix ordering, and the columns corresponding to linearly dependent rows can be avoided at the cost of some additional bookkeeping which leads to a potential reduction in the size of $\left\{\mathbb{S}^{(k)}\right\}$.

## 5. Computational Implementation

### 5.1 Introduction

This chapter shows how the functions (defined in Chapter 3) are then implemented in an efficient fashion, primarily through the use of highly optimized third party library routines. These routines are first introduced in Section 5.2, and then it is shown how with the use of some minor pre-processing steps (including pre- and post-multiplication by permutation matrices see Appendix A), the routines can be used to perform the computations necessary for the functions (Section 5.3). Existing open source implementations of parts of the AI algorithm based on some of these routines are also introduced in Section 5.4. Concluding remarks are presented in Section 5.5.

### 5.2 Basic Linear Algebra Subprograms (BLAS), Linear Algebra PACKage (LAPACK) and BLAS-like subroutines

The following is a brief alphabetically ordered list of Basic Linear Algebra Subprograms (BLAS), LAPACK and BLAS-like subroutines. Their purpose is to implement some specific mathematical operation on a dense matrix which is stored as a collection of some form of decimal number. The first letter of the routine determines the type of decimal number the subroutine operates on (except for TWOSIDEDTRSM because it is written in $\mathrm{C}++$ and is templated, so it automatically works for any of the various data types), in this work all subroutines work with double precision numbers, so all subroutines are preceded with a "D" which is prepended when the subroutine is formally referred to. In the following this first letter is omitted. The names are typically somewhat informative regarding what the subroutine does but there is no fixed rule to determine it. The following has the name and level of the subroutine (level 3
is typically order $n^{3}$ FLOPs: matrix-matrix operations while level 2 is typically order $n^{2}$ matrix-vector operations and level 1 is typically order $n$ vector scalar operations), followed by the purpose of the subroutine, followed by the algebra, followed by roughly how many FLOPs it requires (many of the calculations are of the form of the number of FLOPs per element multiplied by the number of elements, for the derivation of the non-trivial FLOP results see Mazur (2022)), and finally where it is used. The below are implemented in the Intel MKL and are documented in Intel (2019), with the exception of TWOSIDEDTRSM is documented in Poulson et al. (2013).

LAUUM-3

POTRF-3
"Computes a matrix-matrix product with general matrices. $\boldsymbol{C}:=\alpha \times \operatorname{op}(\boldsymbol{A}) \times \operatorname{op}(\boldsymbol{B})+\beta \times \boldsymbol{C}$, where $\operatorname{op}(\boldsymbol{X})$ is one of $\operatorname{op}(\boldsymbol{X})=\boldsymbol{X}$, or $\operatorname{op}(\boldsymbol{X})=\boldsymbol{X}^{\top}, \alpha$ and $\beta$ are scalars, $\boldsymbol{A}, \boldsymbol{B}$ and $\boldsymbol{C}$ are matrices: $\operatorname{op}(\boldsymbol{A})$ is an $m$-by- $k$ matrix, $\operatorname{op}(\boldsymbol{B})$ is a $k$-by- $n$ matrix, $\boldsymbol{C}$ is an $m$-by- $n$ matrix", roughly $2 m k n$ FLOPs. Used in back-updating and in updating the inverse.
"Computes a matrix-matrix product with general matrices but updates only the upper or lower triangular part of the result matrix. The operation is defined as $\boldsymbol{C}:=\alpha \times \mathrm{op}(\boldsymbol{A}) \times \mathrm{op}(\boldsymbol{B})+\beta \times \boldsymbol{C}$, where $\operatorname{op}(\boldsymbol{X})$ is one of $\operatorname{op}(\boldsymbol{X})=\boldsymbol{X}$, or $\operatorname{op}(\boldsymbol{X})=\boldsymbol{X}^{\top}, \alpha$ and $\beta$ are scalars, $\boldsymbol{A}, \boldsymbol{B}$ and $\boldsymbol{C}$ are matrices: op $(\boldsymbol{A})$ is an $n$-by- $k$ matrix, $\operatorname{op}(\boldsymbol{B})$ is a $k$-by- $n$ matrix, $\boldsymbol{C}$ is an $n$-by- $n$ upper or lower triangular matrix", roughly $n^{2} k$ FLOPs. Used in forward-updating the inverse.
"Computes a matrix-vector product using a general matrix." "The operation is defined as $\boldsymbol{y}:=\alpha \times \boldsymbol{A} \times$ $\boldsymbol{x}+\beta \times \boldsymbol{y}$ or $\boldsymbol{y}:=\alpha \times \boldsymbol{A}^{\top} \times \boldsymbol{x}+\beta \times \boldsymbol{y}$ where $\alpha$ and $\beta$ are scalars, $\boldsymbol{x}$ and $\boldsymbol{y}$ are vectors and $\boldsymbol{A}$ is an $m$-by- $n$ matrix.", roughly $2 m n$ FLOPs. Used in forward-updating.
"Performs a rank-1 update of a general matrix. $\boldsymbol{A}:=$ $\alpha \times \boldsymbol{x} \times \boldsymbol{y}^{\top}+\boldsymbol{A}$ where $\alpha$ is a scalar, $\boldsymbol{x}$ is an $m$ element vector, $\boldsymbol{y}$ is an $n$-element vector, $\boldsymbol{A}$ is an $m$-by- $n$ general matrix", roughly $2 m n$ FLOPs. Used in factorization.
"Computes the product $[\boldsymbol{A}:=] \boldsymbol{U} \times \boldsymbol{U}^{\top}$ or $[\boldsymbol{A}:=]$ $\boldsymbol{L}^{\top} \times \boldsymbol{L}^{\prime}$, roughly $n^{3} / 3$ FLOPs where $n$ is the order of the matrix. Used in invert and cross multiply.
"Computes the Cholesky factorization of a symmetric (Hermitian) positive-definite matrix." Finds the
$\boldsymbol{L}$ (order $n$ ) such that $\boldsymbol{A}=\boldsymbol{L} \boldsymbol{L}^{\top}$ then $\boldsymbol{A}:=\boldsymbol{L}$, roughly $n^{3} / 3$ FLOPs. Used in factorization.

PSTRF-3

SCAL-1

SWAP-1

SYMM-3

SYRK-3

TRSM-3

TRTRI-3
"Computes the Cholesky factorization with complete pivoting of a real symmetric (complex Hermitian) positive semi-definite matrix." Finds the $\boldsymbol{L}$ (order $n$ ) such that $\boldsymbol{P}^{\top} \boldsymbol{A} \boldsymbol{P}=\boldsymbol{L} \boldsymbol{L}^{\top}$ "where $\boldsymbol{P}$ is a permutation matrix stored as vector piv" then $\boldsymbol{A}:=\boldsymbol{L}$, roughly $n^{3} / 3$ FLOPs. Used in factorization.
"Computes the product of a vector by a scalar. $\boldsymbol{x}:=$ $a \times \boldsymbol{x}$ where $a$ is a scalar, $\boldsymbol{x}$ is an $n$-element vector", roughly $n$ FLOPs. Used in factorization.
"Swaps a vector with another vector" $\boldsymbol{x}, \boldsymbol{y}:=\boldsymbol{y}, \boldsymbol{x}$, roughly $2 n$ FLOPs where $n$ is the size of the vector. Used in factorization.
"Computes a scalar-matrix-matrix product with one symmetric matrix and add the result to a scalarmatrix product . The operation is defined as $\boldsymbol{C}:=$ $\alpha \times \boldsymbol{A} \times \boldsymbol{B}+\beta \times \boldsymbol{C}$ or $\boldsymbol{C}:=\alpha \times \boldsymbol{B} \times \boldsymbol{A}+\beta \times \boldsymbol{C}$, where where $\alpha$ and $\beta$ are scalars, $\boldsymbol{A}$ is a symmetric matrix, $\boldsymbol{B}$ and $\boldsymbol{C}$ are $m$-by- $n$ matrices", roughly $m^{2} n$ FLOPs in the first case. Used in symmetric multiplication.
"Performs a symmetric rank-k update." "The operation is defined as $\boldsymbol{C}:=\alpha \times \boldsymbol{A} \times \boldsymbol{A}^{\top}+\beta \times \boldsymbol{C}$, or $\boldsymbol{C}:=\alpha \times \boldsymbol{A}^{\top} \times \boldsymbol{A}+\beta \times \boldsymbol{C}$, where $\alpha$ and $\beta$ are scalars, $\boldsymbol{C}$ is an $n$-by- $n$ symmetric matrix, $\boldsymbol{A}$ is an $n$-by- $k$ matrix in the first case and a $k$-by- $n$ matrix in the second case", roughly $n^{2} k$ FLOPs. Used in back-updating.
"Solves a triangular matrix equation" (and thus essentially multiplies a matrix by the inverse of a triangular matrix in a cheaper way) "op $(\boldsymbol{A}) \times \boldsymbol{X}=$ $\alpha \times \boldsymbol{B}$, or $\boldsymbol{X} \times \operatorname{op}(\boldsymbol{A})=\alpha \times \boldsymbol{B}$, where $\alpha$ is a scalar, $\boldsymbol{X}$ and $\boldsymbol{B}$ are $m$-by- $n$ matrices, $\boldsymbol{A}$ is a unit, or nonunit, upper or lower triangular matrix, $\operatorname{op}(\boldsymbol{A})$ is one of $\operatorname{op}(\boldsymbol{A})=\boldsymbol{A}$, or $\operatorname{op}(\boldsymbol{A})=\boldsymbol{A}^{\top}$ " "The matrix $\boldsymbol{B}$ is overwritten by the solution matrix $\boldsymbol{X}$ ", i.e. $\boldsymbol{X}:=\boldsymbol{B}$, roughly $m^{2} n$ FLOPs. Used in forward-solution and back-solution.
"Computes the inverse of a triangular matrix". $\boldsymbol{A}:=$ $\boldsymbol{A}^{-1}$ where $\boldsymbol{A}$ is a triangular matrix, roughly $n^{3} / 3$ FLOPs where $n$ is the order of the matrix. Used in invert and cross multiply.

TWOSIDEDTRSM-3 "Performs a two-sided triangular solves with multiple right-hand sides which preserves the symmetry of the input matrix, either $\boldsymbol{A}:=\boldsymbol{L}^{-1} \boldsymbol{A} \boldsymbol{L}^{-\top}$ or $\boldsymbol{A}:=\boldsymbol{U}^{-\top} \boldsymbol{A} \boldsymbol{U}^{-1} "$, roughly $n^{3}$ FLOPs where $n$ is the order of the matrix. Used in invert and cross multiply.

### 5.3 BLAS implementations

In this section function names from Section 3 can be shown on the left with their BLAS/LAPACK/BLAS-like implementations on the right.

### 5.3.1 Implementation of Factorize

### 5.3.1.1 LAPACK variant

This variant factorizes $\boldsymbol{C}_{\left\{\mathbb{B}^{(k)}\right\}\left\{\mathbb{B}^{(k)}\right\}}$ using a LAPACK optimized routine built for this purpose which achieves the factorization presumably via some recursively blocked process, as used in the open source version. It is impossible to know exactly how the Intel MKL implementation operates because it is not open source. The subroutine DPOTRF computes the factorization $\boldsymbol{L} \boldsymbol{L}^{\top}$ rather than $\boldsymbol{L}^{\top} \boldsymbol{L}$ as the CF described previously requires. The matrix $\boldsymbol{T} \equiv \boldsymbol{C}_{\left\{\mathbb{B}^{(k)}\right\}\left\{\mathbb{B}^{(k)}\right\}}$ is denoted as the diagonal part of the $k^{\text {th }}$ supernode, where in this case there is no copying as $\boldsymbol{T}$ is not actually formed, but it is purely for notational purposes and any changes that occur to $\boldsymbol{T}$ occur to $\boldsymbol{C}_{\left\{\mathbb{B}^{(k)}\right\}\left\{\mathbb{B}^{(k)}\right\}}$-in other words a deep copy. Reverse permutations were used to create a $\boldsymbol{F}$ such that $\boldsymbol{F}=\boldsymbol{P}_{\boldsymbol{R e v}} \boldsymbol{T} \boldsymbol{P}_{\boldsymbol{R e v}}$. Then $\left\{\mathbb{A}^{(k)}\right\}=\left\{\left|\mathbb{B}^{(k)}\right|-i+1\right\},\left\{\mathbb{C}^{(k)}\right\}=\left\{\left|\mathbb{B}^{(k)}\right|-j+1\right\},\left\{\mathbb{D}^{(k)}\right\}=\{j\}$ and $\left\{\mathbb{E}^{(k)}\right\}=\{i\}$ with $j$ iterating from $\left|\mathbb{B}^{(k)}\right|$ to 1 and $i$ iterating from 1 to $j$ :

$$
\begin{equation*}
\boldsymbol{F}_{\left\{\mathbb{A}^{(k)}\right\}\left\{\mathbb{C}^{(k)}\right\}}:=\boldsymbol{T}_{\left\{\mathbb{D}^{(k)}\right\}\left\{\mathbb{E}^{(k)}\right\}} \mid \quad \boldsymbol{F}_{\{\mathbb{A}(k)\}\left\{\mathbb{C}^{(k)}\right\}}=\boldsymbol{T}_{\left\{\mathbb{D}^{(k)}\right\}\left\{\mathbb{E}^{(k)}\right\}} . \tag{5.1}
\end{equation*}
$$

Then $\boldsymbol{F}$ is factorized:

$$
\begin{equation*}
\boldsymbol{F}:=\operatorname{Factorize}(\boldsymbol{F}) \quad \mid \quad \operatorname{DPOTRF}\left(" L ",\left|\mathbb{B}^{(\mathbb{k})}\right|, \boldsymbol{F},\left|\mathbb{B}^{(\mathbb{k})}\right|, \text { INF0 }\right), \tag{5.2}
\end{equation*}
$$

where INFO holds an error value that provides information in case something went wrong. Finally $\boldsymbol{F}$ needs to be reversely permuted back into $\boldsymbol{T}$ :

$$
\begin{equation*}
\boldsymbol{T}_{\left\{\mathbb{A}^{(k)}\right\}\{\mathbb{C}(k)\}}:=\boldsymbol{F}_{\left\{\mathbb{D}^{(k)}\right\}\left\{\mathbb{E}^{(k)}\right\}} \mid \quad \boldsymbol{T}_{\left\{\mathbb{A}^{(k)}\right\}\left\{\mathbb{C}^{(k)}\right\}}=\boldsymbol{F}_{\left\{\mathbb{D}^{(k)}\right\}\left\{\mathbb{E}^{(k)}\right\}} \tag{5.3}
\end{equation*}
$$

### 5.3.2 Implementation of Factorize ${ }^{+}$

### 5.3.2.1 LAPACK variant with singularity handling

This variant is an alternative LAPACK variant that uses the avoiding method to handle singularities, which works in the same way as the above except it uses DPSTRF instead of DPOTRF.
$\boldsymbol{F}_{\left\{\mathbb{A}^{(k)}\right\}\left\{\mathbb{C}^{(k)}\right\}}:=\boldsymbol{T}_{\left\{\mathbb{D}^{(k)}\right\}\left\{\mathbb{E}^{(k)}\right\}} \mid \quad \boldsymbol{F}_{\{\mathbb{A}(k)\}\left\{\mathbb{C}^{(k)}\right\}}=\boldsymbol{T}_{\left\{\mathbb{D}^{(k)}\right\}\left\{\mathbb{E}^{(k)}\right\}}$.
The matrix $\boldsymbol{T}$ could have been permuted in place without the use of $\boldsymbol{F}$ but such permutation requires some temporary space anyway, so $\boldsymbol{F}$ simplifies operations in this regard. Then $\boldsymbol{F}$ is factorized:

$$
\begin{equation*}
\boldsymbol{F}:=\text { Factorize }^{+}(\boldsymbol{F}) \quad \operatorname{DPSTRF}\left(" L ",\left|\mathbb{B}^{(k)}\right|, \boldsymbol{F},\left|\mathbb{B}^{(k)}\right|, \operatorname{PERM}, d_{\left\{\mathbb{B}^{(k)}\right\}}, 0.0001, \text { WORK, INF0 }\right) . \tag{5.5}
\end{equation*}
$$

The subroutine DPSTRF factorizes the input matrix column by column starting from the left and after each column it permutes the column with the largest pivot value to be the next column factorized, as this is what the open source version does. It is hard to know exactly what goes on in the Intel MKL implementation as it is closed source, but this is the best guess of the author. This permutation is stored in PERM. Once none of the remaining pivots are greater than the tolerance ( 0.0001 was chosen for this) then the rest of the columns are classed as singular and then the rank of $\boldsymbol{C}_{\left\{\mathbb{B}^{(k)}\right\}\left\{\mathbb{B}^{(k)}\right\}}$ is reported in $d_{\left\{\mathbb{B}^{(k)}\right\}}$. The vector WORK is just some extra memory (presumably to facilitate the permutations) and INFO holds an error value that provides information in case something went wrong. Finally $\boldsymbol{F}$ needs to be reversely permuted back into $\boldsymbol{T}$ :

$$
\begin{equation*}
\boldsymbol{T}_{\{\mathbb{A}(k)\}\{\mathbb{C}(k)\}}:=\boldsymbol{F}_{\left\{\mathbb{D}^{(k)}\right\}\left\{\mathbb{E}^{(k)}\right\}} \mid \quad \boldsymbol{T}_{\{\mathbb{A}(k)\}\left\{\mathbb{C}^{(k)}\right\}}=\boldsymbol{F}_{\left\{\mathbb{D}^{(k)}\right\}\left\{\mathbb{E}^{(k)}\right\}} \tag{5.6}
\end{equation*}
$$

then $\zeta_{\left\{\mathbb{B}^{(k)}\right\}}=\left|\mathbb{B}^{(\mathbb{k})}\right|-d_{\left\{\mathbb{B}^{(k)}\right\}}$ and finally $\left\{\mathbb{B}^{(\mathbb{k})+}\right\}:=\left\{\mathbb{B}^{(k)}\right\}_{1+\zeta_{\left\{\mathbb{B}^{(k)}\right\}}:\left|\mathbb{B}^{(\mathbb{k})}\right|}$

### 5.3.2.2 Matrix variant with the zeroing-out method for singularity-handling

This variant factorizes $C_{\left\{\mathbb{B}^{(k)}\right\}\left\{\mathbb{B}^{(k)}\right\}}$ row by row with $j$ iterating from $\left|\mathbb{B}^{(\mathbb{k})}\right|$ to 1 and $\boldsymbol{T}$ is as above. Then defining $\left\{\mathbb{T}^{*}\right\}=\{1:(j-1)\},\left\{\mathbb{B}^{*}\right\}=\{j\}$,

$$
\boldsymbol{T}=\left[\begin{array}{cc}
\boldsymbol{T}_{:\left\{\mathbb{T}^{*}\right\}\left\{\mathbb{T}^{*}\right\}:} & \left(\boldsymbol{T}_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{T}^{*}\right\}}\right)^{\top}  \tag{5.7}\\
\boldsymbol{T}_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{T}^{*}\right\}} & t_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{B}^{*}\right\}}
\end{array}\right] .
$$

If $\left(t_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{B}^{*}\right\}}<0.0001\right)$ Then
$t_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{B}^{*}\right\}}:=100000000$
$\boldsymbol{T}_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{T}^{*}\right\}}:=0$
Else

$$
\begin{aligned}
& t_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{B}^{*}\right\}}:=t_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{B}^{*}\right\}}^{1 / 2} \\
& \boldsymbol{T}_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{T}^{*}\right\}}:=t_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{B}^{*}\right\}}^{-1} \boldsymbol{T}_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{T}^{*}\right\}} \\
& \boldsymbol{T}_{:\left\{\mathbb{T}^{*}\right\}\left\{\mathbb{T}^{*}\right\}:}:=\boldsymbol{T}_{:\left\{\mathbb{T}^{*}\right\}\left\{\mathbb{T}^{*}\right\}:}- \\
& \left(\boldsymbol{T}_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{T}^{*}\right\}}\right)^{\top} \boldsymbol{T}_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{T}^{*}\right\}}
\end{aligned}
$$

End If

$$
\begin{gather*}
\text { If }\left(t_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{B}^{*}\right\}}<0.0001\right) \text { Then } \\
\quad t_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{B}^{*}\right\}}:=100000000  \tag{5.8}\\
\boldsymbol{T}_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{T}^{*}\right\}}:=0 \tag{5.9}
\end{gather*}
$$

Else
$t_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{B}^{*}\right\}}:=\operatorname{SQRT}\left(t_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{B}^{*}\right\}}\right)$
$\operatorname{DSCAL}\left(j-1,\left(t_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{B}^{*}\right\}}\right)^{-1}, \boldsymbol{T}_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{T}^{*}\right\}}, 1\right)$
$\operatorname{DSYRK}\left(" L ", " T ", j-1,1,-1.0, \boldsymbol{T}_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{T}^{*}\right\}}\right.$,

$$
\begin{equation*}
\left.1,1.0, \boldsymbol{T}_{\{1\},\{1\}},\left|\mathbb{B}^{(k)}\right|\right) \tag{5.12}
\end{equation*}
$$

End If
and finally $\left\{\mathbb{B}^{(\mathbb{k})+}\right\}:=\left\{\mathbb{B}^{(k)}\right\}$.

### 5.3.2.3 Matrix variant with the avoiding method for singularity-handling

This variant factorizes $C_{\left\{\mathbb{B}^{(k)}\right\}\left\{\mathbb{B}^{(k)}\right\}}$. row by row with $j$ iterating from $\left|\mathbb{B}^{(\mathbb{k})}\right|$ to 1 , and finishing prematurely if $j=\zeta_{\{\mathbb{B}(k)\}}$ as this means all the remaining rows are linearly dependent. The matrix $\boldsymbol{T}$ is as above and denoting $\boldsymbol{S} \equiv \boldsymbol{C}_{\{\mathbb{B}(k)\}\left(\left\{\mathcal{S}^{(k)}\right\} \cup\left\{\mathbb{B}^{(k)}\right\}\right)} \equiv\left[{ }^{\left.C_{\{\mathbb{B}}(k)\right\}\left\{\mathbb{S}^{(k)}\right\}}{ }^{\left.C_{\{\mathbb{B}}(k)\right\}\left\{\mathbb{B}^{(k)}\right\}}{ }\right] \equiv[\boldsymbol{R T}]$ (i.e. $\boldsymbol{S}$ refers to the entire supernode while $\boldsymbol{R}$ refers to the off-diagonal part) and $\boldsymbol{f} \equiv \boldsymbol{b}_{\left\{\mathbb{B}^{(k)}\right\}}$. Both $\boldsymbol{S}$ and $\boldsymbol{f}$ are dense. Then defining $\left\{\mathbb{T}^{*}\right\}=\left\{\left(1+\zeta_{\{\mathbb{B}(k)\}}\right):(j-1)\right\}$, $\left\{\mathbb{B}^{*}\right\}=\{j\}$. Finally denoting a temporary $a$ and $\boldsymbol{A}$ that are meaningless and serve only to temporarily store values due to swapping (the equivalent of WORK in DPSTRF)

$$
\begin{aligned}
& \text { If }\left(t_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{B}^{*}\right\}}<0.0001\right) \text { Then } \\
& \zeta_{\left\{\mathbb{B}^{(k)}\right\}}:=\zeta_{\left\{\mathbb{B}^{(k)}\right\}}+1 \\
& \boldsymbol{A}:=\boldsymbol{T}_{1:\left|\mathbb{B}^{(k)}\right|, \zeta_{\left\{\mathbb{B}^{(k)}\right\}}} \\
& \boldsymbol{T}_{\left.1:\left|\mathbb{B}^{(k)}\right|, \zeta_{\{\mathbb{B}}(k)\right\}}:=\boldsymbol{T}_{1:|\mathbb{B}(\mathbb{k})|, j} \\
& \boldsymbol{T}_{1:\left|\mathbb{B}^{(k)}\right|, j}:=\boldsymbol{A} \\
& \boldsymbol{A}:=\boldsymbol{S}_{j, 1:\left(\left|\mathbb{B}^{(\mathbb{k})}\right|+\mid \mathbb{S}^{(k) \mid)}\right.} \\
& \boldsymbol{S}_{j, 1:\left(\left|\mathbb{B}^{(\mathbb{K})}\right|+\mid \mathbb{S}^{(k) \mid)}\right.}:=\boldsymbol{S}_{\zeta_{\left\{\mathbb{B}^{(k)}\right\}}, 1:\left(\left|\mathbb{B}^{(k)}\right|+\mid \mathbb{S}^{(k) \mid)}\right.} \\
& \boldsymbol{S}_{\zeta_{\left\{\mathbb{B}^{(k)}\right\}}, 1:\left(\left|\mathbb{B}^{(k)}\right|+\mid \mathbb{S}^{(k) \mid)}\right.}:=\boldsymbol{A} \\
& a:=\boldsymbol{f}_{\left.\zeta_{\{\mathbb{B}}(k)\right\}} \\
& \boldsymbol{f}_{\left.\zeta_{\{\mathbb{B}}(k)\right\}}:=\boldsymbol{f}_{j} \\
& \boldsymbol{f}_{j}:=a
\end{aligned}
$$

Else

$$
\begin{aligned}
& t_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{B}^{*}\right\}}:=t_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{B}^{*}\right\}}^{1 / 2} \\
& \boldsymbol{T}_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{T}^{*}\right\}}:=t_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{B}^{*}\right\}}^{-1} \boldsymbol{T}_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{T}^{*}\right\}} \\
& \boldsymbol{T}_{:\left\{\mathbb{U}^{*}\right\}\left\{\mathbb{T}^{*}\right\}:}:=\boldsymbol{T}_{:\left\{\mathbb{U}^{*}\right\}\left\{\mathbb{T}^{*}\right\}:}- \\
& \left(\boldsymbol{T}_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{T}^{*}\right\}}\right)^{\top} \boldsymbol{T}_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{T}^{*}\right\}}
\end{aligned}
$$

End If

$$
\text { If }\left(t_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{B}^{*}\right\}}<0.0001\right) \text { Then }
$$

$$
\begin{equation*}
\zeta_{\left\{\mathbb{B}^{(k)}\right\}}:=\zeta_{\left\{\mathbb{B}^{(k)}\right\}}+1 \tag{5.13}
\end{equation*}
$$

$$
\operatorname{DSWAP}\left(\left|\mathbb{B}^{(\mathbb{k})}\right|, \boldsymbol{T}_{\left.1:\left|\mathbb{B}^{(\mathbb{k})}\right|, \zeta_{\{\mathbb{B}}(k)\right\}}, 1,\right.
$$

$$
\begin{equation*}
\left.\boldsymbol{T}_{1:\left|\mathbb{B}^{(k)}\right|, j}, 1\right) \tag{5.14}
\end{equation*}
$$

$$
\begin{align*}
& \operatorname{DSWAP}\left(\left|\mathbb{S}^{(k)}\right|, \boldsymbol{S}_{j, 1:\left(\left|\mathbb{B}^{(k)}\right|+\left|\mathbb{S}^{(k) \mid}\right|\right.}, 1,\right. \\
& \left.\boldsymbol{S}_{\zeta_{\left\{\mathbb{B}^{(k)}\right\}}, 1:\left(\left|\mathbb{B}^{(k)}\right|+\left|\mathbb{S}^{(k) \mid}\right|\right.}, 1\right)  \tag{5.15}\\
& \operatorname{SWAPDOUBLE}\left(\boldsymbol{f}_{\zeta_{\left\{\mathbb{B}^{(k)}\right\}}}, \boldsymbol{f}_{j}\right) \tag{5.16}
\end{align*}
$$

Else

$$
\begin{align*}
& t_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{B}^{*}\right\}}:=\operatorname{SQRT}\left(t_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{B}^{*}\right\}}\right)  \tag{5.17}\\
& \operatorname{DSCAL}\left(j-1,\left(t_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{B}^{*}\right\}}\right)^{-1}, \boldsymbol{T}_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{T}^{*}\right\}}, 1\right)  \tag{5.18}\\
& \operatorname{DSYRK}\left(" L ", " T^{\prime}, j-1-\zeta_{\left\{\mathbb{B}^{(k)}\right\}}, 1,-1.0, \boldsymbol{T}_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{T}^{*}\right\}},\right. \\
& \left.1,1.0, \boldsymbol{T}_{\left\{1+\zeta_{\left\{\mathbb{B}^{(k)}\right\}}\right\},\left\{1+\zeta_{\left\{\mathbb{B}^{(k)}\right\}}\right\}},\left|\mathbb{B}^{(\mathrm{k})}\right|\right) \tag{5.19}
\end{align*}
$$

End If

Where SWAPDOUBLE is a subroutine written by the author to swap two double precision numbers rather than using any BLAS for that task. Finally $\left\{\mathbb{B}^{(\mathbb{k})+}\right\}:=$ $\left\{\mathbb{B}^{(k)}\right\}_{1+\zeta_{\left\{\mathbb{B}^{(k)}\right\}}:\left|\mathbb{B}^{(\mathbb{k})}\right|}$.

### 5.3.3 Implementation of Back-Solve

There is only one variant of this which uses the dense BLAS triangular solve to do the solve. Then newly defining $\boldsymbol{T}^{+} \equiv \boldsymbol{C}_{\left\{\mathbb{B}^{(\mathbb{k})+}\right\}\left\{\mathbb{B}^{(\mathbb{k})+}\right\}}, \boldsymbol{R}^{+} \equiv \boldsymbol{C}_{\left\{\mathbb{B}^{(k)+}\right\}\left\{\mathbb{S}^{(k)}\right\}}$ to be the off-diagonal part of the supernode :

$$
\begin{array}{ll}
\left.\boldsymbol{C}_{\left\{\mathbb{B}^{(\mathbb{K})+}\right\}\left\{\mathbb{S}^{(k)}\right\}}:=\left(\boldsymbol{C}_{\left\{\mathbb{B}^{(\mathbb{k})+}\right\}\left\{\mathbb{B}^{(\mathbb{k})+}\right\}}\right)^{-1}\right)^{\top} \boldsymbol{C}_{\left\{\mathbb{B}^{(\mathbb{k})+}\right\}\left\{\mathbb{S}^{(k)}\right\}} & \operatorname{DTRSM}\left(" L ", " L ", " T ", " N ",\left|\mathbb{B}^{(\mathbb{k})+}\right|,\left|\mathbb{S}^{(k)}\right|, 1.0,\right. \\
& \left.\boldsymbol{T}_{\{1\},\{1\}}^{+},\left|\mathbb{B}^{(\mathbb{k})}\right|, \boldsymbol{R}_{\{1\},\{1\}}^{+},\left|\mathbb{B}^{(\mathbb{k})}\right|\right) \tag{5.20}
\end{array}
$$

The implementation of Forward-Solve occurs in essentially the same way except that the triangular matrix is not transposed.

### 5.3.4 Implementation of Factorize and Back-Solve together with the zeroing-out method for singularity-handling

This vector-vector variant achieves the Factorize and Back-Solve steps together row by row with $j$ iterating from $\left|\mathbb{B}^{(\mathbb{k})}\right|$ to 1 . The matrices $\boldsymbol{S}$ and $\boldsymbol{f}$ are defined as above. Then defining $\lambda=\left(\left|\mathbb{S}^{(k)}\right|+j\right), \iota=\left(\left|\mathbb{S}^{(k)}\right|+1\right),\left\{\mathbb{T}^{*}\right\}=\{1:(j-1)\}$, $\left\{\mathbb{B}^{*}\right\}=\{j\},\left\{\mathbb{L}^{*}\right\}=\{1:(\lambda-1)\},\left\{\mathbb{R}^{*}\right\}=\{\lambda\},\left\{\mathbb{D}^{*}\right\}=\{\iota:(\lambda-1)\}$.
Then the following partitioning can occur:

$$
\begin{align*}
& \boldsymbol{S}=\left[\begin{array}{cc}
\boldsymbol{S}_{:\left\{\mathbb{U}^{*}\right\}\left\{\mathbb{L}^{*}\right\}:} & \left(s_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{L}^{*}\right\}}\right)^{\top} \\
s_{\left.\left\{\mathbb{B}^{*}\right\} \mathbb{L}^{*}\right\}} & s_{\left.\left\{\mathbb{B}^{*}\right\} \mathbb{R}^{*}\right\}}
\end{array}\right],  \tag{5.21}\\
& \boldsymbol{f}=\left[\begin{array}{c}
\boldsymbol{f}_{\left\{\mathbb{T}^{*}\right\}} \\
f_{\left\{\mathbb{B}^{*}\right\}}
\end{array}\right] . \tag{5.22}
\end{align*}
$$

$$
\begin{aligned}
& \text { If }\left(s_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{R}^{*}\right\}}<0.0001\right) \text { Then } \\
& s_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{R}^{*}\right\}}:=100000000 \\
& f:=0 \\
& s_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{L}^{*}\right\}}:=0
\end{aligned}
$$

Else

$$
\begin{aligned}
& \quad s_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{R}^{*}\right\}}:=\left(s_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{R}^{*}\right\}}\right)^{1 / 2} \\
& \quad f_{\left\{\mathbb{B}^{*}\right\}}:=\left(s_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{R}^{*}\right\}}\right)^{-1} f_{\left\{\mathbb{B}^{*}\right\}} \\
& \quad s_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{L}^{*}\right\}}:=\left(s_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{R}^{*}\right\}}\right)^{-1} s_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{Q}^{*}\right\}} \\
& \quad \boldsymbol{f}_{\left\{\mathbb{T}^{*}\right\}}:=\boldsymbol{f}_{\left\{\mathbb{T}^{*}\right\}}- \\
& \quad\left(s_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{D}^{*}\right\}}\right)^{\top} f_{\left\{\mathbb{B}^{*}\right\}} \\
& \quad \boldsymbol{S}_{:\left\{\mathbb{T}^{*}\right\}\left\{\mathbb{L}^{*}\right\}:}:=\boldsymbol{S}_{:\left\{\mathbb{T}^{*}\right\}\left\{\mathbb{Q}^{*}\right\}:}- \\
& \quad\left(\boldsymbol{s}_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{D}^{*}\right\}}\right)^{\top} \boldsymbol{s}_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{L}^{*}\right\}} \\
& \text { End If }
\end{aligned}
$$

If $\left(s_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{R}^{*}\right\}}<0.0001\right)$ Then
$s_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{R}^{*}\right\}}:=100000000$
$f:=0$
$s_{\left\{\mathbb{B}^{*}\right\}\left\{\mathrm{L}^{*}\right\}}:=0$
Else

$$
\begin{align*}
& s_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{R}^{*}\right\}}:=\operatorname{SQRT}\left(s_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{R}^{*}\right\}}\right)  \tag{5.26}\\
& \operatorname{DSCAL}\left(1,\left(s_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{R}^{*}\right\}}\right)^{-1}, f_{\left\{\mathbb{B}^{*}\right\}}, 1\right)  \tag{5.27}\\
& \operatorname{DSCAL}\left(\lambda-1,\left(s_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{R}^{*}\right\}}\right)^{-1}, s_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{L}^{*}\right\}}\right)  \tag{5.28}\\
& \operatorname{DGER}\left(j-1,1,-1.0,\left(s_{\left.\left\{\mathbb{B}^{*}\right\}\{ \}\right\}}\right)^{\top}, 1,\right. \\
& \left.f_{\left\{\mathbb{B}^{*}\right\}}, 1,1, \boldsymbol{f}_{\left\{\mathbb{T}^{*}\right\}},\left.\right|_{\mathbb{B}^{(\mathbb{k})} \mid} \mid\right)  \tag{5.29}\\
& \operatorname{DGER}\left(j-1, \lambda-1,-1.0,\left(\boldsymbol{s}_{\left\{\mathbb{B}^{*}\right\}\{\iota\}}\right)^{\top}, 1,\right. \\
& \left.s_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{L}^{*}\right\}}, 1, \boldsymbol{S}_{\{1\}\{1\}},\left|\mathbb{B}^{(\mathbb{k})}\right|\right) \tag{5.30}
\end{align*}
$$

End If

Finally $\left\{\mathbb{B}^{(k)+}\right\}:=\left\{\mathbb{B}^{(k)}\right\}$.

### 5.3.5 Implementation of Factorize and Back-Solve together with the avoiding method for singularity-handling

This vector-vector variant achieves the Factorize and Back-Solve steps together row by row with $j$ iterating from $\left|\mathbb{B}^{(\mathbb{k})}\right|$ to 1 , and finishing prematurely if $j=\zeta_{\left\{\mathbb{B}^{(k)}\right\}}$ as this means all the remaining rows are linearly dependent. $\boldsymbol{S}$ and
$f$ are defined as above. Then defining $\lambda=\left(\left|\mathbb{S}^{(k)}\right|+j\right), \iota=\left(\left|\mathcal{S}^{(k)}\right|+1+\zeta_{\{\mathbb{B}(k)\}}\right)$, $\left\{\mathbb{T}^{*}\right\}=\left\{\left(1+\zeta_{\left\{\mathbb{B}^{(k)}\right\}}\right):(j-1)\right\},\left\{\mathbb{B}^{*}\right\}=\{j\},\left\{\mathbb{L}^{*}\right\}=\{1:(\lambda-1)\},\left\{\mathbb{R}^{*}\right\}=\{\lambda\}$, $\left\{\mathbb{D}^{*}\right\}=\{\iota:(\lambda-1)\}$, while also defining $a$ and $\boldsymbol{A}$ as meaningless temporaries.

If $\left(s_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{R}^{*}\right\}}<0.0001\right)$ Then
$\zeta_{\left\{\mathbb{B}^{(k)}\right\}}:=\zeta_{\{\mathbb{B}(k)\}}+1$
$s_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{R}^{*}\right\}}:=100000000$
$f:=0$
$s_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{Q}^{*}\right\}}:=0$
Else

$$
\begin{aligned}
& s_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{R}^{*}\right\}}:=\left(s_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{R}^{*}\right\}}\right)^{1 / 2} \\
& f_{\left\{\mathbb{B}^{*}\right\}}:=\left(s_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{R}^{*}\right\}}\right)^{-1} f_{\left\{\mathbb{B}^{*}\right\}} \\
& s_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{Q}^{*}\right\}}:=\left(s_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{R}^{*}\right\}}\right)^{-1} s_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{Q}^{*}\right\}} \\
& \boldsymbol{f}_{\left\{\mathbb{U}^{*}\right\}}:=\boldsymbol{f}_{\left\{\mathbb{\mathbb { T }}^{*}\right\}}- \\
& \left(s_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{D}^{*}\right\}}\right)^{\top} f_{\left\{\mathbb{B}^{*}\right\}} \\
& \boldsymbol{S}_{:\left\{\mathbb{T}^{*}\right\}\left\{\mathbb{L}^{*}\right\}}:=\boldsymbol{S}_{\left.:\left\{\mathbb{U}^{*}\right\}\left\{\mathbb{L}^{*}\right\}\right\}}- \\
& \left(s_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{D}^{*}\right\}}\right)^{\top} s_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{Q}^{*}\right\}}
\end{aligned}
$$

End If

If $\left(s_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{R}^{*}\right\}}<0.0001\right)$ Then
$\zeta_{\left\{\mathbb{B}^{(k)}\right\}}:=\zeta_{\left\{\mathbb{B}^{(k)}\right\}}+1$
$s_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{R}^{*}\right\}}:=100000000$
$f:=0$
$s_{\left\{\mathbb{B}^{*}\right\}\left\{\mathrm{L}^{*}\right\}}:=0$
Else

$$
\begin{align*}
& s_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{R}^{*}\right\}}:=\operatorname{SQRT}\left(s_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{R}^{*}\right\}}\right)  \tag{5.35}\\
& \operatorname{DSCAL}\left(1,\left(s_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{R}^{*}\right\}}\right)^{-1}, f_{\left\{\mathbb{B}^{*}\right\}}, 1\right)  \tag{5.36}\\
& \operatorname{DSCAL}\left(\lambda-1,\left(s_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{R}^{*}\right\}}\right)^{-1}, s_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{Q}^{*}\right\}}\right) \tag{5.37}
\end{align*}
$$

End If

Finally $\left\{\mathbb{B}^{(k)+}\right\}:=\left\{\mathbb{B}^{(k)}\right\}_{1+\zeta_{\left\{\mathbb{B}^{(k)}\right\}}\left|\mathbb{B}^{(k)}\right|}$

### 5.3.6 Implementation of SchurComplement

The following establishes a right-looking scheme and as a result rather than updating supernodes $\boldsymbol{C}_{\left\{\mathbb{T}^{(k)}\right\}\left\{\mathrm{T}^{(k)}\right\}}$ all at once the updating process is split across the various supernodes. Then denoting $\left\{\mathbb{T}^{*}\right\}=\left(\left\{\mathbb{C}_{k}\right\} \cap\left\{\mathfrak{q}_{i}\right\}\right),\left\{\mathbb{L}^{*}\right\}=\left(\left\{\mathbb{C}_{k}\right\} \cap\right.$ $\left.\left\{\mathbb{C}_{i}\right\}\right)$ with $i$ iterating in any order such that $\left\{\mathbb{T}^{*}\right\} \neq \emptyset$ and letting $\boldsymbol{R}^{+} \equiv$ $\boldsymbol{C}_{\{\mathbb{B}(\mathbb{k})+\}\left\{\mathbb{T}^{*}\right\}}$ and $\boldsymbol{Q}^{+} \equiv \boldsymbol{C}_{\left\{\mathbb{B}^{(k)+\}} \boldsymbol{Q}^{*}\right\}}$, and a single copy $\boldsymbol{A}^{+}:=\left(\boldsymbol{R}^{+}\right)^{\top}$ made to avoid having to transpose $\boldsymbol{R}^{+}$multiple times.

$$
\begin{align*}
& C_{:\left\{\mathbb{T}^{*}\right\}\left\{\mathbb{T}^{*}\right\}:}:=: C_{:\left\{\mathbb{T}^{*}\right\}\left\{\mathbb{T}^{*}\right\}:}= \\
& \left(\boldsymbol{C}_{\left\{\mathbb{B}^{(\mathrm{K})+}+\left\{\mathbb{T}^{*}\right\}\right.}\right)^{\top} \boldsymbol{C}_{\left\{\mathbb{B}^{(\mathrm{K})}+\right\}\left\{\mathbb{T}^{*}\right\}}  \tag{5.40}\\
& C_{:\left\{\mathbb{U}^{*}\right\}\left\{\text { ® }^{*}\right\}:}:=: C_{:\left\{\mathbb{T}^{*}\right\}\left\{\text { ® }^{*}\right\}:}- \\
& \left(\boldsymbol{C}_{\left\{\mathbb{B}^{(k)}+\right\}\left\{\mathbb{T}^{*}\right\}}\right)^{\top} * \boldsymbol{C}_{\left\{\mathbb{B}^{(k)}+\right\}\left\{\mathbb{Q}^{*}\right\}} \tag{5.41}
\end{align*}
$$

Additionally the overhead in performing the gathering and scattering is reused where possible in the performing of Back-UpdateSCP. In the case of updates to $r$ the overhead of gathering and scattering into copies can be avoided since it is entirely dense.

### 5.3.7 Implementation of Back-UpdateSCP

Similarly updating the right hand side is split across the various supernodes, letting $\left\{\mathbb{T}^{*}\right\}=\left(\left\{\mathbb{C}_{k}\right\} \cap\left\{\mathbb{c}_{i}\right\}\right)$, with $i$ iterating in any order such that $\left\{\mathbb{T}^{*}\right\} \neq \emptyset$. The matrix $\boldsymbol{A}^{+}$and vector $\boldsymbol{f}^{+}$are defined as above.

$$
\begin{align*}
& \boldsymbol{b}_{\left\{\mathbb{T}^{*}\right\}}:=: \boldsymbol{b}_{\left\{\mathbb{T}^{*}\right\}}- \\
& \left(\boldsymbol{C}_{\left\{\mathbb{B}^{(\mathbb{k})+}+\left\{\left\{\mathbb{L}^{*}\right\}\right.\right.}\right)^{\top} \boldsymbol{b}_{\left\{\mathbb{B}^{(\mathbb{R})+}\right\}} \tag{5.42}
\end{align*}
$$

$$
\begin{aligned}
& \operatorname{DGEMM}\left(" N ", " N ",\left|\mathbb{T}^{*}\right|, 1,\left|\mathbb{B}^{(\mathbb{k})+}\right|,-1.0, \boldsymbol{A}_{\{1\}\{1\}}^{+},\right. \\
& \left.\left|\mathbb{B}^{(\mathbb{k})}\right|+\left|\mathbb{S}^{(k)}\right|, \boldsymbol{f}_{\left\{1: \mid \mathbb{B}^{(\mathbb{k})+\mid\}}\right.}^{+},\left|\mathbb{B}^{(\mathbb{k})+}\right|, 1.0, \boldsymbol{b}_{\left\{\mathbb{T}^{*}\right\}},\left|\mathbb{T}^{*}\right|\right) .
\end{aligned}
$$

### 5.3.8 Implementation of SymmetricMultiply

This demonstration of SymmetricMultiply uses $\boldsymbol{A}=\left(\boldsymbol{C}_{\left\{\mathbb{B}^{(\mathbb{k})+}\right\}\left\{\mathbb{B}^{(k)}+{ }^{(k)}\right.}\right)^{-1} \boldsymbol{C}_{\left\{\mathbb{B}^{(k)}+\right\}\left\{\mathbb{S}^{(k)}\right\}}$ assuming that it has already been calculated in a previous Forward-Solve step. Then letting $\boldsymbol{T}^{+} \equiv \boldsymbol{C}_{\left\{\mathbb{B}^{(\mathbb{k})+}{ }^{(\{\mathbb{B}}{ }^{(\mathbb{k})+}{ }^{+} \text {, which needs to be gathered as it is not stored }\right.}$ in contiguous memory and also $\boldsymbol{R}^{+}$is as before.

$$
\boldsymbol{C}_{\left\{\mathbb{B}^{(k)+}\right\}\left\{\mathbb{S}^{(k)}\right\}}:=-\boldsymbol{A} \boldsymbol{C}_{:\left\{\mathbb{S}^{(k)}\right\}\left\{\mathbb{S}^{(k)}\right\}:} \quad \operatorname{DSYMM}\left(" R ", " L ",\left|\mathbb{B}^{(\mathbb{k})+}\right|,\left|\mathbb{S}^{(k)}\right|,-1.0, \boldsymbol{T}_{\{1\}\{1\}},\left|\mathbb{S}^{(k)}\right|,\right.
$$

$$
\begin{equation*}
\left.\boldsymbol{A},\left|\mathbb{B}^{(k)+}\right|, 0, \boldsymbol{R}_{\{1\}\{1\}}^{+},\left|\mathbb{B}^{(\mathbb{k})}\right|\right) \tag{5.43}
\end{equation*}
$$

### 5.3.9 Implementation of InvertCrossMultiply

There are multiple implementations of this step, the DTRTRI and DLAUUM method which exploits symmetry and Intel MKL BLAS, a method based on twosidedtrsm from the Elemental (Poulson et al., 2013) library which exploits symmetry in a more general way and the standard two calls of TRSM (the way CHOLMOD-Extra does it) which does not exploit symmetry.

### 5.3.9.1 Implementation using DTRTRI and DLAUUM

First $\boldsymbol{C}_{\left\{\mathbb{B}^{(k)}\right\}\left\{\mathbb{B}^{(k)}\right\}}$ is inverted using DTRTRI then the final result is calculated using DLAUUM, however because DLAUUM computes $\boldsymbol{L}^{\top} \boldsymbol{L}$ rather than $\boldsymbol{L} \boldsymbol{L}^{\top}$ then reverse permutations must be used. With $\boldsymbol{T}^{+} \equiv \boldsymbol{C}_{\left\{\mathbb{B}^{(\mathbb{k})+}{ }^{(1) \mathbb{B}}{ }^{(\mathbb{k})+\}}\right.}$ and $\boldsymbol{F}^{+}=$ $\boldsymbol{P}_{\boldsymbol{R e v}} \boldsymbol{T}^{+} \boldsymbol{P}_{\text {Rev }}$. Then denoting $\left\{\mathbb{A}^{(k)}\right\}=\left\{\left|\mathbb{B}^{(\mathrm{k})+}\right|-i+1\right\},\left\{\mathbb{C}^{(k)}\right\}=\left\{\left|\mathbb{B}^{(\mathrm{k})+}\right|-j+1\right\}$,
$\left\{\mathbb{D}^{(k)}\right\}=\{j\}$ and $\left\{\mathbb{E}^{(k)}\right\}=\{i\}$ with $j$ iterating from $\left|\mathbb{B}^{(k)+}\right|$ to 1 and $i$ iterating from 1 to $j$ :

$$
\begin{equation*}
\boldsymbol{F}_{\left\{\mathbb{A}^{(k)}\right\}\left\{\mathbb{C}^{(k)}\right\}}^{+}:=\boldsymbol{T}_{\left\{\mathbb{D}^{(k)}\right\}\left\{\mathbb{E}^{(k)}\right\}}^{+} \quad \boldsymbol{F}_{\left\{\mathbb{A}^{(k)}\right\}\left\{\mathbb{C}^{(k)}\right\}}^{+}=\boldsymbol{T}_{\left\{\mathbb{D}^{(k)}\right\}\left\{\mathbb{E}^{(k)}\right\}}^{+} . \tag{5.44}
\end{equation*}
$$

Then $\boldsymbol{F}^{+}$is inverted and cross multiplied in place:

$$
\begin{array}{l|l}
\boldsymbol{F}^{+}:=\left(\boldsymbol{F}^{+}\right)^{-1} & \operatorname{DTRTRI}\left(" L ", " N ",\left|\mathbb{B}^{(k)+}\right|, \boldsymbol{F}^{+},\left|\mathbb{B}^{(k)+}\right|, \text { INF0 }\right), \\
\boldsymbol{F}^{+}:=\boldsymbol{F}^{+}\left(\boldsymbol{F}^{+}\right)^{\top} & \operatorname{DLAUUM}\left(" L ",\left|\mathbb{B}^{(k)+}\right|, \boldsymbol{F}^{+},\left|\mathbb{B}^{(k)+}\right|, \text { INF0 }\right) . \tag{5.46}
\end{array}
$$

where INFO holds an error value that provides information in case something went wrong. Finally $\boldsymbol{F}^{+}$needs to be reversely permuted back into $\boldsymbol{T}^{+}$:

$$
\begin{equation*}
\boldsymbol{T}_{\left\{\mathbb{A}^{(k)}\right\}\left\{\mathbb{C}^{(k)}\right\}}^{+}:=\boldsymbol{F}_{\left\{\mathbb{D}^{(k)}\right\}\left\{\mathbb{E}^{(k)}\right\}}^{+} \quad \boldsymbol{T}_{\left\{\mathbb{A}^{(k)}\right\}\left\{\mathbb{C}^{(k)}\right\}}^{+}=\boldsymbol{F}_{\left\{\mathbb{D}^{(k)}\right\}\left\{\mathbb{E}^{(k)}\right\}}^{+} . \tag{5.47}
\end{equation*}
$$

### 5.3.9.2 Implementation based on twosidedtrsm

This implementation operates differently, and it requires a temporary $\boldsymbol{S}^{+}:=$ $\boldsymbol{C}_{\left\{\mathbb{B}^{(\mathbb{k})+\}\left\{\mathbb{B}^{(k)+}\right\}}\right.}=\boldsymbol{L}_{\left\{\mathbb{B}^{(k)}+\right\}\left\{\mathbb{B}^{(\mathbb{k})+}{ }^{(k)}\right.}$, set prior to the beginning of the SIS process for the $k^{\text {th }}$ supernode, because twosidedtrsm cannot multiply a matrix by its own inverse, instead it cross-multiplies a matrix by the inverse of another matrix. Then $\boldsymbol{T}^{+} \equiv \boldsymbol{C}_{\left\{\mathbb{B}^{(k)+\}}\right\}\left\{\mathbb{B}^{(k)+}\right\}}$ is pre and post-multiplied by the inverse of $\boldsymbol{S}^{+}$ via the solution of triangular systems. An interface was written by the author called forttwosidedtrsm that was used since Elemental is not written in FORTRAN (the implementation of this is in Appendix C).

$$
\boldsymbol{C}_{\left\{\mathbb{B}^{(k)+}\right\}\left\{\mathbb{B}^{(\mathbb{k})+}\right\}}:=\left(\boldsymbol{S}^{+}\right)^{-1} \boldsymbol{C}_{\left\{\mathbb{B}^{(\mathbb{k})+}\right\}\left\{\left\{\mathbb{B}^{(\mathbb{k})+}{ }^{2}\right.\right.}\left(\left(\boldsymbol{S}^{+}\right)^{-1}\right)^{\top} \left\lvert\, \begin{align*}
& \text { forttwosidedtrsm }\left(\left|\mathbb{B}^{(\mathbb{k})+}\right|,\left|\mathbb{B}^{(\mathbb{k})+}\right|,\left|\mathbb{B}^{(\mathbb{k})}\right|,\right. \\
& \left.\boldsymbol{S}_{:,\left\{1: \mid \mathbb{B}^{(\mathbb{k})+\mid\}}\right.}^{+},\left|\mathbb{B}^{(\mathbb{k})}\right|, \boldsymbol{T}_{:,\left\{1: \mid \mathbb{B}^{(\mathbb{k})+\mid\}}\right.}^{+}, \zeta_{\left\{\mathbb{B}^{(k)}\right\}}\right) . \tag{5.48}
\end{align*}\right.
$$

### 5.3.9.3 Implementation using two TRSM calls

This implementation is very similar to the implementation using twosidedtrsm except it does this using two TRSM calls, which fails to exploit symmetry.

$$
\begin{align*}
& C_{\left\{\mathbb{B}^{(\mathbb{k})+}{ }^{(1)} \mathbb{B}^{(\mathbb{k})+\}}\right.}:=\left(\boldsymbol{S}^{+}\right)^{-1} \boldsymbol{C}_{\left\{\mathbb{B}^{(\mathbb{k})+}\right\}\left\{\mathbb{B}^{(\mathbb{k})+}\right\}} \quad \operatorname{DTRSM}\left(" L ", " L ", " N ", " N ",\left|\mathbb{B}^{(\mathbb{k})+}\right|,\left|\mathbb{B}^{(\mathbb{k})+}\right|, 1.0,\right. \\
& \left.\boldsymbol{S}_{\{1\}\{1\}}^{+},\left|\mathbb{B}^{(\mathbb{k})}\right|, \boldsymbol{T}_{\{1\}\{1\}}^{+},\left|\mathbb{B}^{(\mathbb{k})}\right|\right) \text {, }  \tag{5.49}\\
& \boldsymbol{C}_{\left\{\mathbb{B}^{(\mathbb{K})+}+\left\{\mathbb{B}^{(\mathbb{R})+}\right\}\right.}:=\boldsymbol{C}_{\left\{\mathbb{B}^{(\mathbb{k})+}\right\}\left\{\mathbb{B}^{(\mathbb{R})+}\right\}}\left(\left(\boldsymbol{S}^{+}\right)^{-1}\right)^{\top} \\
& \operatorname{DTRSM}\left(" R ", " L ", " T ", " N ",\left|\mathbb{B}^{(k)+}\right|,\left|\mathbb{B}^{(k)+}\right|, 1.0\right. \text {, } \\
& \left.\boldsymbol{S}_{\{1\}\{1\}}^{+},\left|\mathbb{B}^{(\mathbb{k})}\right|, \boldsymbol{T}_{\{1\}\{1\}}^{+},\left|\mathbb{B}^{(\mathbb{k})}\right|\right) \text {. } \tag{5.50}
\end{align*}
$$

In theory the two TRSM calls could be executed in reverse order without consequence, but for reasons unknown to the author calling them in this order was faster by a non-trivial amount in preliminary tests.

### 5.3.10 Implementation of Forward-UpdateInv

### 5.3.10.1 Implementation to exploit symmetry DGEMMT

This version is very similar to Back-UpdateSCP, except that since $\boldsymbol{C}_{\left\{\mathbb{B}^{(k)+}\right\}\left\{\mathbb{B}^{(\mathbb{k})+}\right\}}$ is symmetric, only the lower triangle of it needs to be updated. There is an Intel MKL subroutine DGEMMT that exploits this symmetry in the result matrix however, so this is used instead. Afterwards a subroutine called COMPLETESQUARE written by the author is used to copy the lower triangle of the diagonal part into the top triangle, which may be necessary for the invert and cross multiply step. This is significantly cheaper than computing both triangles, to the point of being completely dominated by the calculation of the lower triangle. Symmetry is exploited using DGEMMT and COMPLETESQUARE. Using $\boldsymbol{R}^{+}$and $\boldsymbol{T}^{+}$as before,


$$
\begin{array}{l|l}
\boldsymbol{C}_{\left\{\mathbb{B}^{(k)+}+\right\}\left\{\mathbb{B}^{(k)+}\right\}}:=\boldsymbol{C}_{\left\{\mathbb{B}^{(k)+}+\left\{\mathbb{B}^{(k)+\}}\right.\right.}- & \operatorname{DGEMMT}\left(" L ", " N ", " T ",\left|\mathbb{B}^{(\mathbb{k})+}\right|,\left|\mathbb{S}^{(k)}\right|,-1.0,\right. \\
\boldsymbol{C}_{\left\{\mathbb{B}^{(k)}+\right\}\left\{\mathbb{S}^{(k)}\right\}} \boldsymbol{A}^{\top} & \left.\boldsymbol{R}_{\{1\}\{1\}}^{+},\left|\mathbb{B}^{(\mathbb{k})}\right|, \boldsymbol{A}_{\{1\}\{1\}},\left|\mathbb{B}^{(\mathbb{k})}\right|, 1.0, \boldsymbol{T}_{\{1\}\{1\}}^{+},\left|\mathbb{B}^{(\mathbb{k})}\right|\right), \\
& \text { COMPLETESQUARE }\left(\boldsymbol{T}_{\left\{1:\left|\mathbb{B}^{(k)+}\right|\right\}\left\{1: \mid \mathbb{B}^{(\mathbb{k})+\mid\}}\right.}^{+}\right) .
\end{array}
$$

### 5.3.10.2 Implementation using DGEMM

This is the same as above, except that symmetry of $\boldsymbol{C}_{\left\{\mathbb{B}^{(k)}+\right\}\left\{\mathbb{B}^{(\mathbb{k})+}\right\}}$ is not exploited:

$$
\begin{array}{l|l}
\boldsymbol{C}_{\left\{\mathbb{B}^{(\mathbb{k})+}{ }_{\}\left\{\mathbb{B}^{(k)+}+\right.}\right.}:=\boldsymbol{C}_{\left\{\mathbb{B}^{(\mathbb{k})+}\right\}\left\{\mathbb{B}^{(\mathbb{k})+}\right\}}- & \operatorname{DGEMM}\left(" N ", " T ",\left|\mathbb{B}^{(\mathbb{k})+}\right|,\left|\mathbb{B}^{(\mathbb{k})+}\right|,\left|\mathbb{S}^{(k)}\right|,-1.0,\right. \\
\boldsymbol{C}_{\left\{\mathbb{B}^{(\mathbb{k})+\}\left\{S^{(k)}\right\}}\right.} \boldsymbol{A}^{\top} & \left.\boldsymbol{R}_{\{1\}\{1\}}^{+},\left|\mathbb{B}^{(\mathbb{k})}\right|, \boldsymbol{A}_{\{1\}\{1\}},\left|\mathbb{B}^{(\mathbb{k})}\right|, 1.0, \boldsymbol{T}_{\{1\}\{1\}}^{+},\left|\mathbb{B}^{(\mathbb{k})}\right|\right) . \tag{5.53}
\end{array}
$$

### 5.3.11 Leading dimension

Many of the BLAS and LAPACK routines have arguments associated with the dimensions of the matrices, as well as leading dimensions. When the matrix
arguments are entire matrices, the latter are effectively redundant, however when the matrices to be operated on are partitions of larger dense matrices then these can be used to avoid unnecessary copying. Matrices are stored as vectors in a computer, with additional information regarding the number of rows and columns in order to treat this vector as a matrix. FORTRAN stores matrices by columns, which means that all of the elements in the same column are stored contiguously. As a result of this, partitioning a matrix by column is very straightforward for the processor, the relevant columns can simply be taken and put together in the partition since it knows how many elements there are per column. In particular, taking the first however many columns of a matrix requires no copying, one simply takes the elements associated with these columns from the vector of memory. However, partitioning by row is not as straight forward, if one wanted to take a number of rows of a matrix these would have to be found and then copied into a new contiguous memory vector, regardless where these rows are in the original matrix. The leading dimension argument is a way that the user can supply a matrix with say $b$ rows but only take a partition of the first $a \leq b$ rows. In this case the processor takes the first $a$ elements in the first column, then skips the next $b-a$ elements and moves onto the next column and so on.

### 5.4 Existing open source alternatives

### 5.4.1 CHOLMOD and CHOLMOD-Extra

It is easiest to call the functions from C (ISO, 2011) as CHOLMOD (Chen et al., 2008) and CHOLMOD-Extra (Luttinen, 2018) are written in C.

### 5.4.1.1 Symbolic factorization and reverse permutation

CHOLMOD computes a $\boldsymbol{L} \boldsymbol{L}^{\top}$ rather than a $\boldsymbol{L}^{\top} \boldsymbol{L}$ so reverse permutations need to be applied, so letting CHOLMODPERM $=n-1, n-2, \ldots 0$ (as CHOLMOD is zero-indexed):
$L:=$ cholmod_l_analyze_p2(TRUE, $\boldsymbol{C}$, CHOLMODPERM, NULL, $0, \&$ CHOLMODCOMMON);
Where TRUE refers to a CF as opposed to an LDLT factorization, NULL is an option to not subset the matrix and CHOLMODCOMMON (passed by reference hence the \&) stores auxiliary information at different steps in the process. This creates the symbolic factorization and stores it in $\boldsymbol{L}$.

### 5.4.1.2 Cholesky factorization

$$
\begin{equation*}
\text { cholmod_l_factorize }(\boldsymbol{C}, \boldsymbol{L}, \& \text { CHOLMODCOMMON }) ; \tag{5.55}
\end{equation*}
$$

CHOLMOD uses a forward left-looking algorithm for the factorization.

### 5.4.1.3 Forward-solving and back-solving

CHOLMOD does not do any solving in factorization, so both of these steps must be applied, but this is all wrapped up in the following function. Curiously $\boldsymbol{L}$ does not overwrite $\boldsymbol{C}$ but $\boldsymbol{C}$ can simply be deleted at this point.

$$
\begin{equation*}
\tilde{\boldsymbol{\beta}}:=\text { cholmod_l_solve(CHOLMOD_A, } \boldsymbol{L}, \boldsymbol{b}, \& \text { CHOLMODCOMMON }) ; \tag{5.56}
\end{equation*}
$$

CHOLMOD_A simply tells CHOLMOD to use $\boldsymbol{C}$ as the coefficient matrix rather than $\boldsymbol{C} \boldsymbol{C}^{\top}$ which is another option in the case of rectangular $\boldsymbol{C}$. Similarly $\tilde{\boldsymbol{\beta}}$ does not overwrite $\boldsymbol{b}$ but $\boldsymbol{b}$ can be deleted at this point.

### 5.4.1.4 Calculation of the sparse inverse subset

$$
\begin{equation*}
\boldsymbol{A}:=\text { cholmod_l_spinv }(\boldsymbol{L}, \& C H O L M O D C O M M O N) \tag{5.57}
\end{equation*}
$$

where in this case $\boldsymbol{A}$ is used to store the SIS of $\boldsymbol{C}$. The software CHOLMOD-Extra is used to compute the SIS as CHOLMOD in its base form does not have this functionality. CHOLMOD-Extra uses a back left-looking algorithm to compute the SIS.

### 5.4.1.5 Simplicial mode

Within CHOLMOD (and by extension CHOLMOD-Extra) there is the option to run in either supernodal mode, or simplicial mode (which is single-nodal mode). This is denoted as CHOLMOD-sn and CHOLMOD-Extra-sn respectively.

### 5.4.1.6 Key BLAS libraries

The factorization occurs using LAPACK variant (Subsubsection 5.3.1.1), Back-Solve (Subsubsection 5.3.3) and SchurComplement (Subsubsection 5.3.6), while the calculation of the SIS occurs using Forward-Solve (Subsubsection 5.3.3), SymmetricMultiply (Subsubsection 5.3.8), two TRSM calls (Subsubsection 5.3.9.3) and DGEMM (Subsubsection 5.3.10.2).

### 5.4.2 CHOMPACK

It is easiest to call the various functions from Python (Van Rossum \& Others, 2007) as CHOMPACK (Vandenberghe \& Andersen, 2015) is written in Python.

### 5.4.2.1 Symbolic factorization and reverse permutation

CHOMPACK computes a $\boldsymbol{L} \boldsymbol{L}^{\top}$ rather than a $\boldsymbol{L}^{\top} \boldsymbol{L}$ so reverse permutations need to be applied, so letting CHOLMODPERM $=n-1, n-2, \ldots 0$ (as CHOMPACK is zero-indexed):

$$
\begin{equation*}
\text { symb }:=\text { symbolic }(\boldsymbol{C}, \mathrm{p}=\text { CHOMPACKPERM, merge_function }=\text { fmerge }) \tag{5.58}
\end{equation*}
$$

Where merge_function $=$ fmerge is an option to add duplicate elements if they appear, which doesn't matter since they are not present. This creates the symbolic factorization and stores it in symb.

### 5.4.2.2 Cholesky factorization

$$
\begin{equation*}
\text { cholesky }(\boldsymbol{C}) \tag{5.59}
\end{equation*}
$$

CHOMPACK uses a forward multifrontal algorithm for the factorization.

### 5.4.2.3 Forward-solving and back-solving

CHOMPACK does not do any solving in factorization, so both of these steps must be applied, but this is all wrapped up in the following function.

$$
\begin{equation*}
\operatorname{trsm}(\boldsymbol{C}, \boldsymbol{b}) \tag{5.60}
\end{equation*}
$$

This trsm is not to be confused with the level 3 BLAS DTRSM.

### 5.4.2.4 Calculation of the sparse inverse subset

$$
\begin{equation*}
\text { projected_inverse }(\boldsymbol{C}) \tag{5.61}
\end{equation*}
$$

CHOMPACK uses a back multifrontal algorithm for the calculation of the SIS.

### 5.4.2.5 Key BLAS Libraries

CHOMPACK uses an LDLT algorithm for the factorization, which means its factorization has a forward solve and a backward solve, but its SIS calculation doesn't have a backward solve component. Its factorization uses the LAPACK variant (Subsubsection 5.3.1.1), Back-Solve (Subsubsection 5.3.3), Forward-Solve (Subsubsection 5.3.3) and a scheme similar to SchurComplement (Subsubsection 5.3.6) except that it does update $\boldsymbol{C}_{\{\mathbb{T}(k)\}\{\mathbb{T}(k)\}}$ all at once (as mentioned) using the level 3 BLAS DSYRK as it is multifrontal. For the SIS it uses, SymmetricMultiply (Subsubsection 5.3.8), DTRTRI and DSYRK (taking the place of DLAUUM in Subsubsection 5.3.9.1) and DGEMM (Subsubsection 5.3.10.2).

### 5.4.3 MUMPS

It is easiest to call the various functions from FORTRAN (Reid, 2008) as MUMPS (Amestoy et al., 2000) is written in mainly FORTRAN.

### 5.4.3.1 Symbolic factorization and reverse permutation

MUMPS computes a $\boldsymbol{L} \boldsymbol{L}^{\top}$ rather than a $\boldsymbol{L}^{\top} \boldsymbol{L}$ so reverse permutations need to be applied, so letting MUMPSPERM $=n, n-1, \ldots 1$ (as MUMPS is one-indexed):

$$
\begin{align*}
& \text { mumps_par\%JOB = } 1  \tag{5.62}\\
& \text { DMUMPS(mumps_par) } \tag{5.63}
\end{align*}
$$

Where mumps_par is a common object that stores everything related to the matrix in question, i.e. $\boldsymbol{C}$ and $\boldsymbol{b}$ are parts of it and DMUMPS is the subroutine that is called for each step, with mumps_par\%JOB being used to select which operation is to be performed, in this case it is the symbolic factorization due to the value of 1 . A value of 2 refers to CF and a value of 3 refers to one of the types of solves.

### 5.4.3.2 Cholesky factorization

$$
\begin{align*}
& \text { mumps_par\%JOB }=2  \tag{5.64}\\
& \text { DMUMPS(mumps_par) } \tag{5.65}
\end{align*}
$$

MUMPS uses a forward multifrontal algorithm for the factorization.

### 5.4.3.3 Forward-solving and back-solving

MUMPS has the capability to do forward solving in factorization but this feature is used for a specific application, and is more likely to be less efficient in the general case, so for that reason it is turned off. Therefore both of these steps must be applied, but this is all wrapped up in the following function.

$$
\begin{align*}
& \text { mumps_par\%J0B = } 3  \tag{5.66}\\
& \text { DMUMPS(mumps_par) } \tag{5.67}
\end{align*}
$$

### 5.4.3.4 Calculation of the sparse inverse subset

$$
\begin{align*}
& \text { mumps_par\%ICNTL(30) }=1  \tag{5.68}\\
& \text { mumps_par\%JOB }=3  \tag{5.69}\\
& \text { DMUMPS(mumps_par) } \tag{5.70}
\end{align*}
$$

MUMPS does not have a true SIS interface, instead it has a selected inversion interface where if it is provided with various indices of elements it will find those elements of $\boldsymbol{C}^{-1}$. The relevant elements are placed into mumps_par and the option mumps_par\%ICNTL (30) must be set to 1 to enable this option. MUMPS finds these elements via its forward-solve and back-solve routines, and so mumps_par\%JOB $=3$. MUMPS does this by solving for a number of columns of the identity matrix, and exploiting sparsity such that it only performs the necessary calculations to find the selected elements of the inverse. This number of columns is controlled by a blocking parameter mumps_par\%ICNTL(27) and for the remainder of this thesis the value of 4096 is chosen. This value was chosen using some initial testing, and as the MUMPS user guide offers little guidance in this matter this was what the author decided was sensible.

### 5.4.3.5 Key BLAS Libraries

MUMPS uses the LAPACK variant (Subsubsection 5.3.1.1), Back-Solve (Subsubsection 5.3.3), Forward-Solve (Subsubsection 5.3.3) and a scheme similar to SchurComplement (Subsubsection 5.3.6) except that it does update $\boldsymbol{C}_{\left\{\mathbb{T}^{(k)}\right\}\left\{\mathbb{T}^{(k)}\right\}}$ all at once (as mentioned) using the level 3 BLAS DSYRK as it is multifrontal. For the calculation of the SIS it uses only Back-Solve (Subsubsection 5.3.3), Forward-Solve (Subsubsection 5.3.3) and DGEMM (Subsubsection 5.3 .10 .2 ).

### 5.5 Conclusions

This chapter introduced the Basic Linear Algebra Subprograms, as well as LAPACK and BLAS-Like routines, which are highly optimized dense linear algebra libraries that can be used to perform various standard matrix operations very efficiently in Section 5.2. The way that these routines could be used to perform the various functions required for the AI algorithm was demonstrated in Section 5.3. Existing open source codes that could perform the most computationally intensive steps of the AI algorithm were also introduced, in Section 5.4.

## Bibliography

Amestoy, Patrick R, Duff, Iain S, L'Excellent, Jean-Yves \& Koster, JACKo (2000). MUMPS: a general purpose distributed memory sparse solver. In International Workshop on Applied Parallel Computing. Springer, Boston, MA, 121-130.

Ashcraft, Cleve \& Grimes, Roger (1989). The influence of relaxed supernode partitions on the multifrontal method. ACM Transactions on Mathematical Software (TOMS) 15(4), 291-309.

Bunch, James R (1971). Analysis of the diagonal pivoting method. SIAM Journal on Numerical Analysis 8(4), 656-680.

Chen, Yanqing, Davis, Timothy A, Hager, William W \& Rajamanickam, Sivasankaran (2008). Algorithm 887: CHOLMOD, supernodal sparse Cholesky factorization and update/downdate. ACM Transactions on Mathematical Software (TOMS) 35(3), 22.

Dayde, Michel J \& Camichel, Rue (1995). A Block Version of the EskowSchnabel Modified Cholesky Factorization.
URL http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1. $25.9977\{\&\}$ rep=rep1 $\{\&\}$ type=pdf

Eskow, Elizabeth \& Schnabel, Robert B (1991). Algorithm 695: Software for a new modified Cholesky factorization. ACM Transactions on Mathematical Software (TOMS) 17(3), 306-312.

Gilmour, A R (1998). ASREML, Technical Details. Technical report, NSW DPI.

Gilmour, A R, Gogel, B J, Cullis, B R, Welham, S J \& Thompson, R (2015). ASReml User Guide Release 4.1 Functional Specification. URL www.vsni.co.uk

Gilmour, Arthur R, Thompson, Robin \& Cullis, Brian R (1995). Average information REML: An efficient algorithm for variance parameter estimation in linear mixed models. Biometrics 51(4), 1440-1450.

Gupta, Anshul (2000). WSMP: Watson sparse matrix package (Part-I: direct solution of symmetric sparse systems). IBM TJ Watson Research Center, Yorktown Heights, NY, Tech. Rep. RC 21886.

Gupta, Anshul, Joshi, Mahesh \& Kumar, Vipin (1997). WSSMP: Watson symmetric sparse matrix package. IBM TJ Watson Research Center.

Henderson, C R (1950). Estimation of genetic parameters (abstract). Annals of Mathematical Statistics 21, 309-310.

Henderson, C R (1973). Sire Evaluation and Genetic Trends. Journal of Animal Science 1973(Symposium), 10-41.
URL https://doi.org/10.1093/ansci/1973.Symposium. 10
Intel (2019). Developer Reference for Intel Math Kernel Library - Fortran. URL https://software.intel.com/en-us/ mkl-developer-reference-fortran

ISO (2011). IEC 9899: 2011 Information technology—Programming lan-guages-C. International Organization for Standardization, Geneva, Switzerland 27, 59.

Liu, Joseph W H (1990). The role of elimination trees in sparse factorization. SIAM Journal on Matrix Analysis and Applications 11(1), 134-172.

Luttinen, JaAkko (2018). cholmod-extra Documentation. URL https://buildmedia.readthedocs.org/media/pdf/ cholmod-extra/stable/cholmod-extra.pdf

Mazur, Luke (2022). Computational Methods for the Fitting of Factor Analytic Linear Mixed Models with Applications to Plant Variety Trials. Ph.D. thesis, University of Wollongong.

Meyer, Karin (1989). Restricted maximum likelihood to estimate variance components for animal models with several random effects using a derivativefree algorithm. Genetics Selection Evolution 21(3), 317.

Patterson, H D \& Thompson, R (1971). Recovery of interblock information when block sizes are unequal. Biometrika 31, 545-554.

Poulson, Jack, Geijn, Robert A V A N D E \& Bennighof, Jeffrey (2012). ( Parallel ) Algorithms for Two-sided Triangular Solves and Matrix Multiplication.
URL https://www.cs.utexas.edu/users/flame/pubs/ ElementalGenEig.pdf

Poulson, Jack, Marker, Bryan, de Geijn, Robert A, Hammond, Jeff R \& Romero, Nichols A (2013). Elemental: A new framework for distributed memory dense matrix computations. ACM Transactions on Mathematical Software (TOMS) 39(2), 1-24.

Reid, John (2008). The new features of Fortran 2008. In ACM SIGPLAN Fortran Forum. ACM New York, NY, USA, 27(2), 8-21.

Schnabel, Robert B \& Eskow, Elizabeth (1990). A new modified Cholesky factorization. SIAM Journal on Scientific and Statistical Computing 11(6), 1136-1158.

Searle, Shayle Roy \& Khuri, Andre I (2017). Matrix Algebra Useful for Statistics. John Wiley \& Sons.

Smith, Alison B, Cullis, Brian R \& Thompson, Robin (2001). Analyzing variety by environment data using multiplicative mixed models and adjustments for spatial field trend. Biometrics 57(4), 1138-1147.

Van Rossum, Guido \& Others (2007). Python Programming Language. In USENIX annual technical conference. 41, 36.

Vandenberghe, Lieven \& Andersen, Martin S (2015). Chordal Graphs and Semidefinite Optimization. Foundations and Trends® in Optimization 1(4), 241-433.
URL http://www. nowpublishers.com/article/Details/0PT-006

## A. Permutation matrices

A brief notation introduction to the concept of a permutation matrix $\boldsymbol{P}^{*}$ (not to be confused with $\boldsymbol{P}$ which is not a permutation matrix). A permutation matrix is simply the identity matrix with some of the rows or columns rearranged. In other words, it is a square matrix with exactly one 1 in each row and in each column with the rest of the elements being 0 . Pre-multiplying by a permutation matrix permutes the rows of a matrix for example $\boldsymbol{P}^{*} \boldsymbol{A}$ will have the rows in a permuted order as compared to $\boldsymbol{A}$ while $\boldsymbol{A} \boldsymbol{P}^{*}$ will have the columns rearranged (except in the special case where $\boldsymbol{P}^{*}=\boldsymbol{I}$ ).An important property is that $\boldsymbol{P}^{*}\left(\boldsymbol{P}^{*}\right)^{\top}=\boldsymbol{P}^{*} \boldsymbol{I}\left(\boldsymbol{P}^{*}\right)^{\top}=\boldsymbol{I}$. For the remainder of this report the superscript is dropped from $\boldsymbol{P}^{*}$ where it is clear from context that $\boldsymbol{P}$ is a permutation matrix.

A special case of a permutation matrix is one that is anti-diagonal, i.e. it has 1 s along the diagonal from the bottom left to the top right and 0s elsewhere. Let this be denoted as $\boldsymbol{P}_{\text {Rev }}$. A symmetric permutation using $\boldsymbol{P}_{\boldsymbol{R e v}}$ produces a matrix that is reversely permuted, i.e. the first diagonal element becomes the last, second diagonal element becomes second last etc. Also $\boldsymbol{P}_{\text {Rev }}=\boldsymbol{P}_{\text {Rev }}^{-1}=$ $\boldsymbol{P}_{\text {Rev }}^{\top}$. These are relevant because ASReml (See Gilmour et al. (2015) for a user guide and Gilmour (1998) for technical details) implements absorption with $k$ iterating from $n$ to 2 while conventional sparse and dense packages, including the BLAS, expect the factorizations to occur with $k$ iterating from 1 to $n$ (or to $m$ in a supernodal setting).

## B. Calculation of FLOPs for BLAS routines

In the following a supernode $k$ is considered such that it has rows represented by the set $\left\{\mathbb{B}^{(k)}\right\}$ meaning that it has $\left|\mathbb{B}^{(\mathbb{k})}\right|$ rows. It has columns in the off-diagonal part represented by the set $\left\{\mathscr{S}^{(k)}\right\}$ and thus has $\left|\mathbb{S}^{(k)}\right|$ such columns. The FLOPs performed by various subroutines are presented, in alphabetical order.

LAUUM-3 is used to multiply two matrices, the first upper triangular of order $\left|\mathbb{B}^{(k)}\right|$ and the second the transpose of the former, forming the symmetric result matrix row by row. The number of multiplications and additions for the $j^{\text {th }}$ element in the $i^{\text {th }}$ row of the resultant matrix are summed:

$$
\begin{aligned}
& \sum_{i=1}^{\left|\mathbb{B}^{(k)}\right|} \sum_{j=1}^{i}\left(\min \left(\left|\mathbb{B}^{(\mathbb{k})}\right|-i+1,\left|\mathbb{B}^{(\mathbb{k})}\right|-j+1\right)+\left(\min \left(\left|\mathbb{B}^{(\mathbb{k})}\right|-i+1,\left|\mathbb{B}^{(\mathbb{k})}\right|-j+1\right)-1\right)\right) \\
& =\sum_{i=1}^{\left|\mathbb{B}^{(\mathbb{k})}\right|} \sum_{j=1}^{i}\left(\left(\left|\mathbb{B}^{(\mathbb{k})}\right|-i+1\right)+\left(\left|\mathbb{B}^{(\mathbb{k})}\right|-i+1-1\right)\right) \\
& =\sum_{i=1}^{\left|\mathbb{B}^{(\mathbb{k})}\right|} \sum_{j=1}^{i}\left(2\left|\mathbb{B}^{(\mathbb{k})}\right|-2 i+1\right) \\
& =\sum_{i=1}^{\left|\mathbb{B}^{(\mathbb{k})}\right|}\left(2\left|\mathbb{B}^{(\mathbb{k})}\right| i-2 i^{2}+i\right) \\
& =2\left|\mathbb{B}^{(\mathbb{k})}\right|\left(\left|\mathbb{B}^{(\mathbb{k})}\right|\left(\left|\mathbb{B}^{(k)}\right|+1\right) / 2\right)-2\left(\left(2\left|\mathbb{B}^{(\mathbb{k})}\right|+1\right)\left|\mathbb{B}^{(\mathbb{k})}\right|\left(\left|\mathbb{B}^{(\mathbb{k})}\right|+1\right) / 6\right)+\left|\mathbb{B}^{(\mathbb{k})}\right|\left(\left|\mathbb{B}^{(\mathbb{k})}\right|+1\right) / 2 \\
& =\left|\mathbb{B}^{(\mathbb{k})}\right|^{2}\left(\left|\mathbb{B}^{(\mathbb{k})}\right|+1\right)-\left(2\left|\mathbb{B}^{(\mathbb{k})}\right|+1\right)\left|\mathbb{B}^{(\mathbb{k})}\right|\left(\left|\mathbb{B}^{(\mathbb{k})}\right|+1\right) / 3+\left|\mathbb{B}^{(\mathbb{k})}\right|\left(\left|\mathbb{B}^{(\mathbb{k})}\right|+1\right) / 2 \\
& \approx\left|\mathbb{B}^{(\mathbb{k})}\right|^{3}-2\left|\mathbb{B}^{(\mathbb{k})}\right|^{3} / 3 \\
& =\left|\mathbb{B}^{(\mathbb{k})}\right|^{3} / 3 .
\end{aligned}
$$

POTRF-3 Dense Cholesky factorization of the diagonal block of size $\left|\mathbb{B}^{(k)}\right|$. The number of multiplications and additions (Equation ??), divisions (Equation ??) and square roots (Equation ??) for the $i^{\text {th }}$ row respectively are
summed below:

$$
\begin{aligned}
& \sum_{i=1}^{\left|\mathbb{B}^{(k)}\right|}(i(i-1) / 2+i(i-1) / 2+(i-1)+1) \\
& =\left(\left|\mathbb{B}^{(k)}\right|-1\right)\left|\mathbb{B}^{(k)}\right|\left(\left|\mathbb{B}^{(k)}\right|+1\right) / 6+\left(\left|\mathbb{B}^{(k)}\right|-1\right)\left|\mathbb{B}^{(\mathbb{k})}\right|\left(\left|\mathbb{B}^{(\mathbb{k})}\right|+1\right) / 6+\left(\left|\mathbb{B}^{(k)}\right|-1\right)\left|\mathbb{B}^{(\mathbb{k})}\right| / 2+\left|\mathbb{B}^{(k)}\right| \\
& =\left(\left|\mathbb{B}^{(\mathbb{k})}\right|-1\right)\left|\mathbb{B}^{(k)}\right|\left(\left|\mathbb{B}^{(k)}\right|+1\right) / 3+\left(\left|\mathbb{B}^{(k)}\right|+1\right)\left|\mathbb{B}^{(\mathbb{k})}\right| / 2 \\
& \approx\left|\mathbb{B}^{(k)}\right|^{3} / 3 .
\end{aligned}
$$

PSTRF-3 As above, except that in the singular case for the diagonal block the location of the singularities are not yet known, they are discovered in this phase. If one were to assume that they are discovered last (worst possible case) then the final $\zeta$ steps are not applied. The number of multiplications, additions, divisions and square roots in the $i^{\text {th }}$ row not applied as a result of this are summed below:

$$
\begin{aligned}
& \sum_{i=1}^{\zeta_{\left\{\mathbb{B}^{(k)}\right\}}}(i(i-1) / 2+i(i-1) / 2+i-1+1) \\
& =\left(\zeta_{\left\{\mathbb{B}^{(k)}\right\}}-1\right) \zeta_{\left\{\mathbb{B}^{(k)}\right\}}\left(\zeta_{\left\{\mathbb{B}^{(k)}\right\}}+1\right) / 6+\left(\zeta_{\left\{\mathbb{B}^{(k)}\right\}}-1\right) \zeta_{\left\{\mathbb{B}^{(k)}\right\}}\left(\zeta_{\left\{\mathbb{B}^{(k)}\right\}}+1\right) / 6+\left(\zeta_{\left\{\mathbb{B}^{(k)}\right\}}-1\right) \zeta_{\left\{\mathbb{B}^{(k)}\right\}} / 2+\zeta_{\left\{\mathbb{B}^{(k)}\right\}} \\
& =\left(\zeta_{\left\{\mathbb{B}^{(k)}\right\}}-1\right) \zeta_{\left\{\mathbb{B}^{(k)}\right\}}\left(\zeta_{\left\{\mathbb{B}^{(k)}\right\}}+1\right) / 3+\zeta_{\left\{\mathbb{B}^{(k)}\right\}}\left(\zeta_{\left\{\mathbb{B}^{(k)}\right\}}+1\right) / 2 \\
& \approx \zeta_{\left\{\mathbb{B}^{(k)}\right\}}^{3} / 3 .
\end{aligned}
$$

So the final FLOPs for PSTRF in the singular case are $\approx\left(\left|\mathbb{B}^{(k)}\right|^{3}-\right.$ $\left.\zeta_{\{\mathbb{B}(k)\}}^{3}\right) / 3$.

TRSM-3 The number of multiplications and subtractions (Equation ??) and divisions (Equation ??) for the $i^{\text {th }}$ row to solve for the corresponding elements of the right hand side are summed below:

$$
\begin{aligned}
& \sum_{i=1}^{\left|\mathbb{B}^{(k)}\right|}((i-1)+(i-1)+1) \\
& =\left(\left|\mathbb{B}^{(k)}\right|-1\right)\left|\mathbb{B}^{(k)}\right| / 2+\left(\left|\mathbb{B}^{(k)}\right|-1\right)\left|\mathbb{B}^{(\mathbb{k})}\right| / 2+\left|\mathbb{B}^{(k)}\right| \\
& =\left(\left|\mathbb{B}^{(k)}\right|-1\right)\left|\mathbb{B}^{(\mathbb{k})}\right|+\left|\mathbb{B}^{(\mathbb{k})}\right| \\
& =\left|\mathbb{B}^{(\mathbb{k})}\right|^{2} .
\end{aligned}
$$

TRTRI-3 This can be considered as using the triangular part of the supernode (post $C F$ ) of order $\left|\mathbb{B}^{(\mathbb{k})}\right|$ to solve for the identity matrix order $\left|\mathbb{B}^{(k)}\right|$ taking advantage of the fact that the inverse of a lower triangular matrix is lower triangular. The number of multiplications, subtractions and divisions (Equation ??) in the $i^{\text {th }}$ row for the $j^{\text {th }}$ right hand side are
summed below:

$$
\begin{aligned}
& \sum_{j=1}^{\left|\mathbb{B}^{(k)}\right|\left|\mathbb{B}^{(\mathbb{k})}\right|}((i-1)+(i-1)+1) \\
& =\sum_{j=1}^{\left|\mathbb{B}^{(\mathbb{k})}\right|} \sum_{i=j}^{\mathbb{B}^{(\mathbb{k})} \mid}(2 i-1) \\
& =\sum_{j=1}^{\left|\mathbb{B}^{(\mathbb{R})}\right|}\left(\sum_{i=1}^{\left|\mathbb{B}^{(k)}\right|}(2 i-1)-\sum_{i=1}^{j-1}(2 i-1)\right) \\
& =\sum_{j=1}^{\left|\mathbb{B}^{(\mathbb{k})}\right|}\left(2\left(\left|\mathbb{B}^{(\mathbb{k})}\right|\left(\left|\mathbb{B}^{(\mathbb{k})}\right|+1\right) / 2\right)-\left|\mathbb{B}^{(\mathbb{k})}\right|-(2(j(j-1) / 2)-j+1)\right) \\
& =\sum_{j=1}^{\left|\mathbb{B}^{(\mathbb{R})}\right|}\left(\left|\mathbb{B}^{(\mathbb{k})}\right|\left(\left|\mathbb{B}^{(\mathbb{k})}\right|+1\right)-\left|\mathbb{B}^{(\mathbb{k})}\right|-2 j(j-1)+j-1\right) \\
& =\sum_{j=1}^{\left|\mathbb{B}^{(\mathbb{R})}\right|}\left(\left|\mathbb{B}^{(\mathbb{R})}\right|^{2}-2 j^{2}+3 j-1\right) \\
& =\left|\mathbb{B}^{(\mathbb{k})}\right|^{3}-2\left(\left(2\left|\mathbb{B}^{(\mathbb{k})}\right|+1\right)\left|\mathbb{B}^{(\mathbb{R})}\right|\left(\left|\mathbb{B}^{(\mathbb{k})}\right|+1\right) / 6\right)+3\left(\left|\mathbb{B}^{(\mathbb{k})}\right|\left(\left|\mathbb{B}^{(\mathbb{k})}\right|+1\right) / 2\right)-\left|\mathbb{B}^{(\mathbb{k})}\right| \\
& =\left|\mathbb{B}^{(\mathbb{k})}\right|^{3}-\left(2\left|\mathbb{B}^{(\mathbb{k})}\right|+1\right)\left|\mathbb{B}^{(\mathbb{k})}\right|\left(\left|\mathbb{B}^{(\mathbb{k})}\right|+1\right) / 3+3\left|\mathbb{B}^{(\mathbb{k})}\right|\left(\left|\mathbb{B}^{(\mathbb{k})}\right|+1\right) / 2-\left|\mathbb{B}^{(\mathbb{k})}\right| \\
& \approx\left|\mathbb{B}^{(\mathbb{k})}\right|^{3}-2\left|\mathbb{B}^{(\mathbb{k})}\right|^{3} / 3 \\
& =\left|\mathbb{B}^{(\mathbb{k})}\right|^{3} / 3 .
\end{aligned}
$$

TWOSIDED Poulson et al. (2012) claim that their variant 4 is the best one, so the TRSM-3 calculations of FLOPs herein pertain to that variant. There is a recursive blocked algorithm at work such that the subscripts with a $\left\{\mathbb{C}^{(k)}\right\}$ refer to the current block, which splits the matrix into the top and bottom parts. Therefore those with a $\left\{\mathbb{T}^{*}\right\}$ refer to those prior to the current block, i.e. parts that have already been computed, while those with a $\left\{\mathbb{B}^{*}\right\}$ refer to those after the current block. It is helpful to realise that $\boldsymbol{L}_{\left\{\mathbb{C}^{(k)}\right\}\left\{\mathbb{C}^{(k)}\right\}}$ and $\boldsymbol{C}_{\left\{\mathbb{C}^{(k)}\right\}\left\{\mathbb{C}^{(k)}\right\}}$ are $|\mathbb{C}|$ by $|\mathbb{C}|$ matrices where $|\mathbb{C}|$ is the blocksize (and thus there are $\left|\mathbb{B}^{(\mathbb{k})}\right| /|\mathbb{C}|$ blocks, assuming that $|\mathbb{C}|$ evenly divides $\left|\mathbb{B}^{(\mathbb{k})}\right|$, for the sake of brevity), and so any $\left\{\mathbb{C}^{(k)}\right\}$ subscript refers to a dimension of size $|\mathbb{C}|$, and any variable with a $\left\{\mathbb{B}^{*}\right\}$ subscript has dimension $|\mathbb{C}|\left(\left|\mathbb{B}^{(\mathbb{k})}\right| /|\mathbb{C}|-i\right)=\left|\mathbb{B}^{(k)}\right|-i|\mathbb{C}|$ where $i$ refers to the current block. Similarly anything with a $\left\{\mathbb{T}^{*}\right\}$ subscript has a dimension of $(i-1)|\mathbb{C}|$. There are two key components which are shaded that are considered. The first of these:

$$
\boldsymbol{C}_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{T}^{*}\right\}}:=\boldsymbol{C}_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{T}^{*}\right\}}-\boldsymbol{L}_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{C}^{(k)}\right\}} \boldsymbol{C}_{\left\{\mathbb{C}^{(k)}\right\}\left\{\mathbb{T}^{*}\right\}},
$$

which becomes after ignoring subtraction (since this is lower order):

$$
\boldsymbol{C}_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{T}^{*}\right\}}:=\boldsymbol{L}_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{C}^{(k)}\right\}} \boldsymbol{C}_{\left\{\mathbb{C}^{(k)}\right\}\left\{\mathbb{T}^{*}\right\}}
$$

The matrices $\boldsymbol{L}_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{C}^{(k)}\right\}}$ and $\boldsymbol{C}_{\left\{\mathbb{C}^{(k)}\right\}\left\{\mathbb{T}^{*}\right\}}$ are standard rectangular matrices of sizes $\left(\left|\mathbb{B}^{(\mathbb{k}}\right|-i|\mathbb{C}|\right) \times|\mathbb{C}|$ and $|\mathbb{C}| \times(i-1)|\mathbb{C}|$ respectively. So multiplying them together leads to the following number of multiplications and additions summed over blocks:

$$
\begin{aligned}
& \sum_{i=1}^{\left|\mathbb{B}^{(k)}\right|| | \mathbb{C} \mid}\left(\left(\left|\mathbb{B}^{(k)}\right|-i|\mathbb{C}|\right) \times|\mathbb{C}| \times(i-1)|\mathbb{C}|+\left(\left|\mathbb{B}^{(k)}\right|-i|\mathbb{C}|\right) \times|\mathbb{C}| \times(i-1)|\mathbb{C}|\right) \\
= & \sum_{i=1}^{\left|\mathbb{B}^{(k)}\right|| | \mathbb{C} \mid}\left(2|\mathbb{C}|^{2}\left(\left|\mathbb{B}^{(k)}\right|-i|\mathbb{C}|\right)(i-1)\right) \\
= & \sum_{i=1}^{\left|\mathbb{B}^{(k)}\right|| | \mathbb{C} \mid} \\
= & \sum_{i=1}^{\left|\mathbb{B}^{(k)}\right|| | \mathbb{C} \mid} \mid \\
= & \left.\left.2|\mathbb{C}|^{2}\left|\mathbb{B}^{2}\right| \mathbb{B}^{(k)}|(i-1)-2| \mathbb{C}\right|^{3} i(i-1)\right) \\
= & \left.\left.2\left|\mathbb{B}^{(k)}\right|\right|^{2}\left(\left|\mathbb{B}^{(k)}\right| /|\mathbb{C}|\left(\left|\mathbb{B}^{(k)}\right| /|\mathbb{C}|-1\right) / 2\right)-4 \mid \mathbb{C}^{3} i(i-1) / 2\right) \\
\approx & \mid \mathbb{B}^{(k)}\left(\left(\left|\mathbb{B}^{(k)}\right| /|\mathbb{C}|-1\right)\left|\mathbb{B}^{(k)}\right|\left|/|\mathbb{C}|\left(\left|\mathbb{B}^{(k)}\right| /|\mathbb{C}|+1\right) / 6\right)\right. \\
= & \left|\mathbb{B}^{(k)}\right|^{3} / 3 .
\end{aligned}
$$

It is now worth considering the second shaded area:

$$
\boldsymbol{C}_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{B}^{*}\right\}}:=\boldsymbol{C}_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{B}^{*}\right\}}-\left(\boldsymbol{L}_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{C}^{(k)}\right\}} C_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{C}^{(k)}\right\}}^{\top}+\boldsymbol{C}_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{C}^{(k)}\right\}} \boldsymbol{L}_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{C}^{(k)}\right\}}^{\top}\right),
$$

which is achieved using HER2K. The author assumes that since the first product in brackets is equal to the transpose of the second product that the second product comes for a lower order cost than the first product, and thus can be ignored. Similarly subtraction at the front of the equation will be ignored. This leaves the following:

$$
\boldsymbol{C}_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{B}^{*}\right\}}:=\left(\boldsymbol{L}_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{C}^{(k)}\right\}} \boldsymbol{C}_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{C}^{(k)}\right\}}^{\top}\right) .
$$

The matrices $\boldsymbol{L}_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{C}^{(k)}\right\}}$ and $\boldsymbol{C}_{\left\{\mathbb{B}^{*}\right\}\left\{\mathbb{C}^{(k)}\right\}}^{\top}$ are standard rectangular matrices of sizes $\left(\left|\mathbb{B}^{(k)}\right|-i|\mathbb{C}|\right) \times|\mathbb{C}|$ and $|\mathbb{C}| \times\left(\left|\mathbb{B}^{(k)}\right|-i|\mathbb{C}|\right)$ respectively. So multiplying them together leads to the following number of multiplica-
tions and additions summed over blocks:

$$
\begin{aligned}
& \sum_{i=1}^{\left|\mathbb{B}^{(k)}\right| /|\mathbb{C}|}\left(\left(\left|\mathbb{B}^{(k)}\right|-i|\mathbb{C}|\right) \times|\mathbb{C}| \times\left(\left|\mathbb{B}^{(k)}\right|-i|\mathbb{C}|\right)+\left(\left|\mathbb{B}^{(k)}\right|-i|\mathbb{C}|\right) \times|\mathbb{C}| \times\left(\left|\mathbb{B}^{(k)}\right|-i|\mathbb{C}|\right)\right) \\
& =\sum_{i=1}^{\left|\mathbb{B}^{(k)}\right| /|\mathbb{C}|}\left(2|\mathbb{C}|\left(\left|\mathbb{B}^{(k)}\right|-i|\mathbb{C}|\right)^{2}\right) \\
& =\sum_{i=1}^{|\mathbb{B}(\mathbb{B})| /|\mathbb{C}|}\left(2|\mathbb{C}|\left|\mathbb{B}^{(k)}\right|^{2}-4|\mathbb{C}|^{2}\left|\mathbb{B}^{(k)}\right| i+2|\mathbb{C}|^{3} i^{2}\right) \\
& =2|\mathbb{C}|\left|\mathbb{B}^{(k)}\right|^{2}\left(\left|\mathbb{B}^{(k)}\right| /|\mathbb{C}|\right)-4|\mathbb{C}|^{2}\left|\mathbb{B}^{(k)}\right|\left(\left|\mathbb{B}^{(k)}\right| /|\mathbb{C}|\left(\left|\mathbb{B}^{(k)}\right| /|\mathbb{C}|+1\right) / 2\right) \\
& +2|\mathbb{C}|^{3}\left(\left|\mathbb{B}^{(k)}\right| /|\mathbb{C}|\left(\left|\mathbb{B}^{(k)}\right| / / \mathbb{C} \mid+1\right)\left(2\left|\mathbb{B}^{(k)}\right| /|\mathbb{C}|+1\right)\right) \\
& \approx 2\left|\mathbb{B}^{(k)}\right|^{3}-2\left|\mathbb{B}^{(k)}\right|^{3}+4\left|\mathbb{B}^{(k)}\right|^{3} / 6 \\
& =2\left|\mathbb{B}^{(k)}\right|^{3} / 3 .
\end{aligned}
$$

Adding these two together (ignoring the non-shaded parts as those are lower order of magnitude) leads to $2\left|\mathbb{B}^{(\mathbb{k})}\right|^{3} / 3+\left|\mathbb{B}^{(k)}\right|^{3} / 3=\left|\mathbb{B}^{(k)}\right|^{3}$ which is in agreement with the aforementioned paper.

## C. C wrapper for implementation of forttwosidedtrsm

The following is the implementation of the $C$ wrapper to enable the use of twosidedtrsm via FORTRAN. First an environment is created as the Elemental library requires an Elemental environment to be created in order to function. Then options are specified, in this case a lower triangular input matrix that is not unit diagonal. Then Elemental matrices are created for both the Cholesky factor and for the output matrix, as this type is required to work with the functions of the Elemental library. Finally, the function twosidedtrsm is called to perform the invert cross-multiply.

```
void forttwosidedtrsm(long* numrows, long* numcols, long* ldl, double** l,
long* ldz, double** z, long* nbr)
{
int tempArgc = 0;
char** tempArgv = new char* [0];
El::Environment* env = new El::Environment( tempArgc, tempArgv );
char uplo = 'L';
El::UpperOrLowerNS::UpperOrLower ElUplo =
El::UpperOrLowerNS::CharToUpperOrLower(uplo);
char diag = 'N';
El::UnitOrNonUnitNS::UnitOrNonUnit ElDiag =
El::UnitOrNonUnitNS::CharToUnitOrNonUnit(diag);
long tempNBR = *nbr;
double* tempL = *l;
double* adjustedL = &tempL[tempNBR];
double* tempZ = *z;
double* adjustedZ = &tempZ[tempNBR];
```

```
El::Matrix<double> fortLMatrix( (El::Int ) (*numrows), (El::Int ) (*numcols),
adjustedL, (El::Int ) (*ldl));
El::Matrix<double> fortZMatrix( (El::Int ) (*numrows), (El::Int ) (*numcols),
adjustedZ, (El::Int ) (*ldz));
// All the matrices are square here so numrows = numcols
// Further all their dimensions must be the same so this one parameter is all that is nee
El::TwoSidedTrsm( ElUplo, ElDiag, fortZMatrix, fortLMatrix);
}
```

