A block approach for approximating selected elements of the matrix inverse

Vasilis Georgiou, Efstratios Gallopoulos

Computer Engineering and Informatics Departement
University of Patras

July 03, 2018

1. Introduction
2. Extracting diagonal elements
3. Approximating diagonal blocks
4. Numerical experiments
5. Conclusions

1 Introduction

2 Extracting diagonal elements

3 Approximating diagonal blocks

4 Numerical experiments

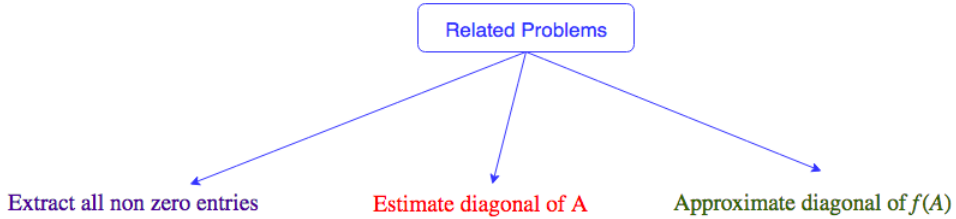
5 Conclusions

Extracting elements from matrices

Problem definition: Given sparse matrix $A \in \mathbb{R}^{n \times n}$, extract a specified set of entries in $\mathcal{O}(n)$ operations

Extracting elements from matrices

Problem definition: Given sparse matrix $A \in \mathbb{R}^{n \times n}$, extract a specified set of entries in $\mathcal{O}(n)$ operations



ESTIMATION OF SPARSE JACOBIAN MATRICES AND GRAPH COLORING PROBLEMS*

THOMAS F. COLEMAN† AND JORGE J. MORÉ‡

Abstract. Given a mapping with a sparse Jacobian matrix, we investigate the problem of minimizing the number of function evaluations needed to estimate the Jacobian matrix by differences. We show that this problem can be attacked as a graph coloring problem and that this approach leads to very efficient algorithms. The behavior of these algorithms is studied and, in particular, we prove that two of the algorithms are optimal for band graphs. We also present numerical evidence which indicates that these two algorithms are nearly optimal on practical problems.

An Estimator for the Diagonal of a Matrix *

C. Bekas, E. Kokiopoulou, Y. Saad

Computer Science & Engineering Dept.
University of Minnesota, Twin Cities
{beaks,kokobeh,saad}@cs.umn.edu

June 1, 2005

Abstract

A number of applications require to compute an approximation of the diagonal of a matrix when this matrix is not explicitly available but matrix-vector products with it are easy to evaluate. In some cases, it is the trace of the matrix rather than the diagonal that is needed. This paper describes methods for estimating diagonals and traces of matrices in these situations. The goal is to obtain a good estimate of the diagonal by applying only a small number of matrix-vector products, using selected vectors. We begin by considering the use of random test vectors and then explore special vectors obtained from Hadamard matrices. The methods are tested in the context of computational materials science to estimate the diagonal of the density matrix which holds the charge densities. Numerical experiments indicate that the diagonal estimator may offer an alternative method that in some cases can greatly reduce computational costs in electronic structures calculations.

Stochastic matrix-function estimators: Scalable Big-Data kernels with high performance

Peter W. J. Staar, Panagiotis Kl. Barkoutsos, Roxana Istrate, A. Cristiano I. Malossi, Ivano Tavernelli, Nikolaj Moll, Heiner Giefers, Christoph Hagleitner, Costas Bekas, Alessandro Curioni
IBM Research Zurich
Saümerstrasse 4
Rüschlikon, Switzerland
{tau,bpa,roi,acm,ita,nim,hgi,hle,bek,cur}@zurich.ibm.com

HIERARCHICAL PROBING FOR ESTIMATING THE TRACE OF THE MATRIX INVERSE ON TOROIDAL LATTICES

ANDREAS STATHOPOULOS †, JESSE LAEUCHLI †, AND KOSTAS ORGINOS ††

Improving stochastic estimates with inference methods: calculating matrix diagonals

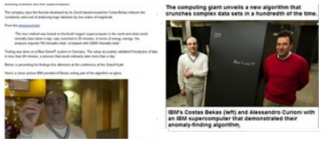
Marco Selig, Niels Oppermann, and Torsten A. Enßlin
Max-Planck-Institut für Astrophysik, Karl-Schwarzschild-Straße 1, 85741 Garching, Germany
(Received 2011 August 02; Accepted 2012 February 23)

Estimating the diagonal entries of a matrix, that is not directly accessible but only available as a linear operator in the form of a computer routine, is a common necessity in many computational applications, especially in image reconstruction and statistical inference. Here, methods of statistical inference are used to improve the accuracy or the computational costs of matrix probing methods to estimate matrix diagonals. In particular, the generalized Wiener filter methodology, as developed within information field theory, is shown to significantly improve estimates based on only a few sampling probes, in cases in which some form of continuity of the solution can be assumed. The strength, length scale, and precise functional form of the exploited autocorrelation function of the matrix diagonal is determined from the probes themselves. The developed algorithm is successfully applied to mock and real world problems. These performance tests show that, in situations where a matrix diagonal has to be calculated from only a small number of computationally expensive probes, a speedup by a factor of 2 to 10 is possible with the proposed method.

A Probing Method for Computing the Diagonal of a Matrix Inverse

Jok M. Tang and Yousef Saad

Department of Computer Science and Engineering, University of Minnesota, 200 Union Street S.E. Minneapolis, MN 55455, USA.



ESTIMATION OF SPARSE JACOBIAN MATRICES AND GRAPH COLORING PROBLEMS*

THOMAS F. COLEMAN† AND JORGE J. MORÉ‡

Abstract. Given a mapping with a sparse Jacobian matrix, we investigate the problem of minimizing the number of function evaluations needed to estimate the Jacobian matrix by differences. We show that this problem can be attacked as a graph coloring problem and that this approach leads to very efficient algorithms. The behavior of these algorithms is studied and, in particular, we prove that two of the algorithms are optimal for band graphs. We also present numerical evidence which indicates that these two algorithms are nearly optimal on practical problems.

An Estimator for the Diagonal of a Matrix *

C. Bekas, E. Kokiopoulou, Y. Saad

Computer Science & Engineering Dept.
University of Minnesota, Twin Cities
{beaks,kokobeh,saad}@cs.umn.edu

June 1, 2005

Abstract

A number of applications require to compute an approximation of the diagonal of a matrix when this matrix is not explicitly available but matrix-vector products with it are easy to evaluate. In some cases, it is the trace of the matrix rather than the diagonal that is needed. This paper describes methods for estimating diagonals and traces of matrices in these situations. The goal is to obtain a good estimate of the diagonal by applying only a small number of matrix-vector products, using selected vectors. We begin by considering the use of random test vectors and then explore special vectors obtained from Hadamard matrices. The methods are tested in the context of computational materials science to estimate the diagonal of the density matrix which holds the charge densities. Numerical experiments indicate that the diagonal estimator may offer an alternative method that in some cases can greatly reduce computational costs in electronic structures calculations.

Stochastic matrix-function estimators: Scalable Big-Data kernels with high performance

Peter W. J. Staar, Panagiotis Kl. Barkoutsos, Roxana Istrate, A. Cristiano I. Malossi, Ivano Tavernelli, Nikolaj Moll, Heiner Giefers, Christoph Hagleitner, Costas Bekas, Alessandro Curioni
IBM Research Zurich
Saümerstrasse 4
Rüschlikon, Switzerland
{tau,bpa,roi,acm,ita,nim,hgi,hle,bek,cur}@zurich.ibm.com

HIERARCHICAL PROBING FOR ESTIMATING THE TRACE OF THE MATRIX INVERSE ON TOROIDAL LATTICES

ANDREAS STATHOPOULOS †, JESSE LAEUCHLI †, AND KOSTAS ORGINOS ††

Improving stochastic estimates with inference methods: calculating matrix diagonals

Marco Selig, Niels Oppermann, and Torsten A. Enßlin
Max-Planck-Institut für Astrophysik, Karl-Schwarzschild-Straße 1, 85741 Garching, Germany
(Received 2011 August 02; Accepted 2012 February 23)

Estimating the diagonal entries of a matrix, that is not directly accessible but only available as a linear operator in the form of a computer routine, is a common necessity in many computational applications, especially in image reconstruction and statistical inference. Here, methods of statistical inference are used to improve the accuracy or the computational costs of matrix probing methods to estimate matrix diagonals. In particular, the generalized Wiener filter methodology, as developed within information field theory, is shown to significantly improve estimates based on only a few sampling probes, in cases in which some form of continuity of the solution can be assumed. The strength, length scale, and precise functional form of the exploited autocorrelation function of the matrix diagonal is determined from the probes themselves. The developed algorithm is successfully applied to mock and real world problems. These performance tests show that, in situations where a matrix diagonal has to be calculated from only a small number of computationally expensive probes, a speedup by a factor of 2 to 10 is possible with the proposed method.

A Probing Method for Computing the Diagonal of a Matrix Inverse

Jok M. Tang and Yousef Saad

Department of Computer Science and Engineering, University of Minnesota, 200 Union Street S.E. Minneapolis, MN 55455, USA.



ESTIMATION OF SPARSE JACOBIAN MATRICES AND GRAPH COLORING PROBLEMS*

THOMAS F. COLEMAN† AND JORGE J. MORÉ‡

Abstract. Given a mapping with a sparse Jacobian matrix, we investigate the problem of minimizing the number of function evaluations needed to estimate the Jacobian matrix by differences. We show that this problem can be attacked as a graph coloring problem and that this approach leads to very efficient algorithms. The behavior of these algorithms is studied and, in particular, we prove that two of the algorithms are optimal for band graphs. We also present numerical evidence which indicates that these two algorithms are nearly optimal on practical problems.

An Estimator for the Diagonal of a Matrix *

C. Bekas, E. Kokiopoulou, Y. Saad

Computer Science & Engineering Dept.
University of Minnesota, Twin Cities
{beaks,kokobeh,saad}@cs.umn.edu

June 1, 2005

Abstract

A number of applications require to compute an approximation of the diagonal of a matrix when this matrix is not explicitly available but matrix-vector products with it are easy to evaluate. In some cases, it is the trace of the matrix rather than the diagonal that is needed. This paper describes methods for estimating diagonals and traces of matrices in these situations. The goal is to obtain a good estimate of the diagonal by applying only a small number of matrix-vector products, using selected vectors. We begin by considering the use of random test vectors and then explore special vectors obtained from Hadamard matrices. The methods are tested in the context of computational materials science to estimate the diagonal of the density matrix which holds the charge densities. Numerical experiments indicate that the diagonal estimator may offer an alternative method that in some cases can greatly reduce computational costs in electronic structures calculations.

Stochastic matrix-function estimators: Scalable Big-Data kernels with high performance

Peter W. J. Staar, Panagiotis Kl. Barkoutsos, Roxana Istrate, A. Cristiano I. Malossi, Ivano Tavernelli, Nikolaj Moll, Heiner Giefers, Christoph Hagleitner, Costas Bekas, Alessandro Curioni
IBM Research Zurich
Saümerstrasse 4
Rüschlikon, Switzerland
{tau,bpa,roi,acm,ita,nim,hgi,hle,bek,cur}@zurich.ibm.com

HIERARCHICAL PROBING FOR ESTIMATING THE TRACE OF THE MATRIX INVERSE ON TOROIDAL LATTICES

ANDREAS STATHOPOULOS [†], JESSE LAEUCHLI [†], AND KOSTAS ORGINOS [‡§]

Improving stochastic estimates with inference methods: calculating matrix diagonals

Marco Selig, Niels Oppermann, and Torsten A. Enßlin
Max-Planck-Institut für Astrophysik, Karl-Schwarzschild-Straße 1, 85741 Garching, Germany
(Received 2011 August 02; Accepted 2012 February 23)

Estimating the diagonal entries of a matrix, that is not directly accessible but only available as a linear operator in the form of a computer routine, is a common necessity in many computational applications, especially in image reconstruction and statistical inference. Here, methods of statistical inference are used to improve the accuracy or the computational costs of matrix probing methods to estimate matrix diagonals. In particular, the generalized Wiener filter methodology, as developed within information field theory, is shown to significantly improve estimates based on only a few sampling probes, in cases in which some form of continuity of the solution can be assumed. The strength, length scale, and precise functional form of the exploited autocorrelation function of the matrix diagonal is determined from the probes themselves. The developed algorithm is successfully applied to mock and real world problems. These performance tests show that, in situations where a matrix diagonal has to be calculated from only a small number of computationally expensive probes, a speedup by a factor of 2 to 10 is possible with the proposed method.

A Probing Method for Computing the Diagonal of a Matrix Inverse

Jok M. Tang and Yousef Saad

Department of Computer Science and Engineering, University of Minnesota, 200 Union Street S.E. Minneapolis, MN 55455, USA.



ESTIMATION OF SPARSE JACOBIAN MATRICES AND GRAPH COLORING PROBLEMS*

THOMAS F. COLEMAN† AND JORGE J. MORÉ‡

Abstract. Given a mapping with a sparse Jacobian matrix, we investigate the problem of minimizing the number of function evaluations needed to estimate the Jacobian matrix by differences. We show that this problem can be attacked as a graph coloring problem and that this approach leads to very efficient algorithms. The behavior of these algorithms is studied and, in particular, we prove that two of the algorithms are optimal for band graphs. We also present numerical evidence which indicates that these two algorithms are nearly optimal on practical problems.

An Estimator for the Diagonal of a Matrix *

C. Bekas, E. Kokiopoulou, Y. Saad

Computer Science & Engineering Dept.
University of Minnesota, Twin Cities
{beaks,kokobeh,saad}@cs.umn.edu

June 1, 2005

Abstract

A number of applications require to compute an approximation of the diagonal of a matrix when this matrix is not explicitly available but matrix-vector products with it are easy to evaluate. In some cases, it is the trace of the matrix rather than the diagonal that is needed. This paper describes methods for estimating diagonals and traces of matrices in these situations. The goal is to obtain a good estimate of the diagonal by applying only a small number of matrix-vector products, using selected vectors. We begin by considering the use of random test vectors and then explore special vectors obtained from Hadamard matrices. The methods are tested in the context of computational materials science to estimate the diagonal of the density matrix which holds the charge densities. Numerical experiments indicate that the diagonal estimator may offer an alternative method that in some cases can greatly reduce computational costs in electronic structures calculations.

Stochastic matrix-function estimators: Scalable Big-Data kernels with high performance

Peter W. J. Staar, Panagiotis Kl. Barkoutsos, Roxana Istrate, A. Cristiano I. Malossi, Ivano Tavernelli, Nikolaj Moll, Heiner Giefers, Christoph Hagleitner, Costas Bekas, Alessandro Curioni
IBM Research Zurich
Saümerstrasse 4
Rüschlikon, Switzerland
{tau,bpa,roi,acm,ita,nim,hgi,ble,bek,cur}@zurich.ibm.com

HIERARCHICAL PROBING FOR ESTIMATING THE TRACE OF THE MATRIX INVERSE ON TOROIDAL LATTICES

ANDREAS STATHOPOULOS †, JESSE LAEUCHLI †, AND KOSTAS ORGINOS ‡§

Improving stochastic estimates with inference methods: calculating matrix diagonals

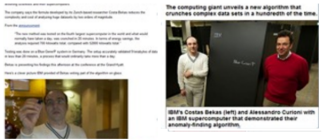
Marco Selig, Niels Oppermann, and Torsten A. Enßlin
Max-Planck-Institut für Astrophysik, Karl-Schwarzschild-Straße 1, 85741 Garching, Germany
(Received 2011 August 02; Accepted 2012 February 23)

Estimating the diagonal entries of a matrix, that is not directly accessible but only available as a linear operator in the form of a computer routine, is a common necessity in many computational applications, especially in image reconstruction and statistical inference. Here, methods of statistical inference are used to improve the accuracy or the computational costs of matrix probing methods to estimate matrix diagonals. In particular, the generalized Wiener filter methodology, as developed within information field theory, is shown to significantly improve estimates based on only a few sampling probes, in cases in which some form of continuity of the solution can be assumed. The strength, length scale, and precise functional form of the exploited autocorrelation function of the matrix diagonal is determined from the probes themselves. The developed algorithm is successfully applied to mock and real world problems. These performance tests show that, in situations where a matrix diagonal has to be calculated from only a small number of computationally expensive probes, a speedup by a factor of 2 to 10 is possible with the proposed method.

A Probing Method for Computing the Diagonal of a Matrix Inverse

Jok M. Tang and Yousef Saad

Department of Computer Science and Engineering, University of Minnesota, 200 Union Street S.E, Minneapolis, MN 55455, USA.



1 Introduction

2 Extracting diagonal elements

3 Approximating diagonal blocks

4 Numerical experiments

5 Conclusions

Extracting diagonal entries of A

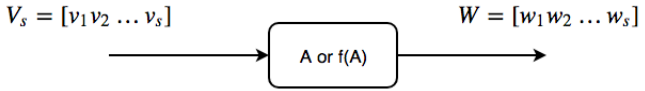
Proposed approaches:

- Stochastic estimation methods
- Approximations that rely on structure

Extracting diagonal entries of A

Proposed approaches:

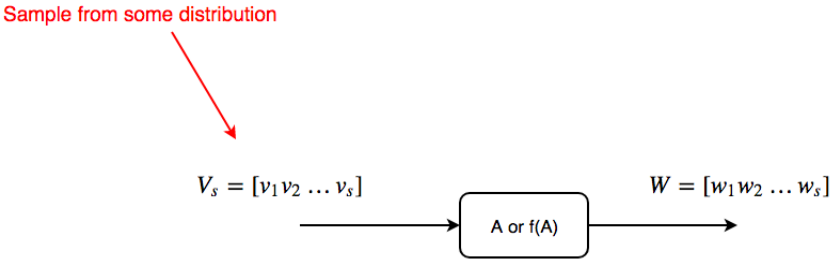
- Stochastic estimation methods
- Approximations that rely on structure



Extracting diagonal entries of A

Proposed approaches:

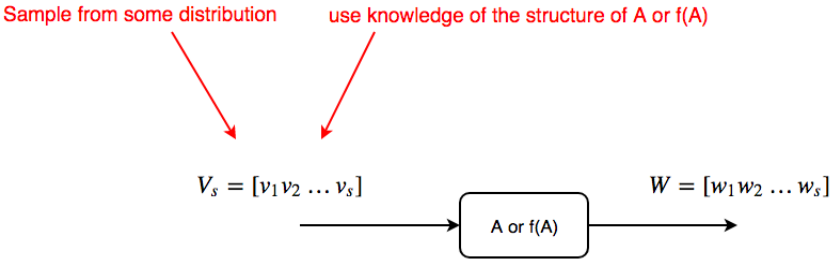
- Stochastic estimation methods
- Approximations that rely on structure



Extracting diagonal entries of A

Proposed approaches:

- Stochastic estimation methods
- Approximations that rely on structure



Extracting diagonal entries of A

$$\begin{pmatrix} \alpha_{11} & 0 & \alpha_{13} & 0 \\ 0 & \alpha_{22} & 0 & \alpha_{24} \\ \alpha_{31} & 0 & \alpha_{33} & \alpha_{34} \\ 0 & \alpha_{42} & \alpha_{43} & \alpha_{44} \end{pmatrix}$$

Extracting diagonal entries of A

$$\begin{pmatrix} \alpha_{11} & 0 & \alpha_{13} & 0 \\ 0 & \alpha_{22} & 0 & \alpha_{24} \\ \alpha_{31} & 0 & \alpha_{33} & \alpha_{34} \\ 0 & \alpha_{42} & \alpha_{43} & \alpha_{44} \end{pmatrix} \odot \underbrace{\begin{pmatrix} 1 & & 0 & \\ & 1 & & 0 \\ 0 & & 1 & 0 \\ & 0 & 0 & 1 \end{pmatrix}}_{C_s}$$

Extracting diagonal entries of A

$$\begin{pmatrix} \alpha_{11} & 0 & \alpha_{13} & 0 \\ 0 & \alpha_{22} & 0 & \alpha_{24} \\ \alpha_{31} & 0 & \alpha_{33} & \alpha_{34} \\ 0 & \alpha_{42} & \alpha_{43} & \alpha_{44} \end{pmatrix} \odot \underbrace{\begin{pmatrix} 1 & 0 & \\ & 1 & 0 \\ 0 & & 1 & 0 \\ & 0 & 0 & 1 \end{pmatrix}}_{C_s}$$

Extracting diagonal entries of A

$$\begin{pmatrix} \alpha_{11} & 0 & \alpha_{13} & 0 \\ 0 & \alpha_{22} & 0 & \alpha_{24} \\ \alpha_{31} & 0 & \alpha_{33} & \alpha_{34} \\ 0 & \alpha_{42} & \alpha_{43} & \alpha_{44} \end{pmatrix} \odot \underbrace{\begin{pmatrix} 1 & x & 0 & x \\ x & 1 & x & 0 \\ 0 & x & 1 & 0 \\ x & 0 & 0 & 1 \end{pmatrix}}_{C_s}$$

Extracting diagonal entries of A

$$\delta(A) = \left[\begin{pmatrix} \alpha_{11} & 0 & \alpha_{13} & 0 \\ 0 & \alpha_{22} & 0 & \alpha_{24} \\ \alpha_{31} & 0 & \alpha_{33} & \alpha_{34} \\ 0 & \alpha_{42} & \alpha_{43} & \alpha_{44} \end{pmatrix} \odot \underbrace{\begin{pmatrix} 1 & x & 0 & x \\ x & 1 & x & 0 \\ 0 & x & 1 & 0 \\ x & 0 & 0 & 1 \end{pmatrix}}_{C_s} \right] \begin{pmatrix} 1 \\ 1 \\ 1 \\ 1 \end{pmatrix}$$

Extracting diagonal entries of A

$$\delta(A) = \left[\begin{pmatrix} \alpha_{11} & 0 & \alpha_{13} & 0 \\ 0 & \alpha_{22} & 0 & \alpha_{24} \\ \alpha_{31} & 0 & \alpha_{33} & \alpha_{34} \\ 0 & \alpha_{42} & \alpha_{43} & \alpha_{44} \end{pmatrix} \odot \underbrace{\begin{pmatrix} 1 & x & 0 & x \\ x & 1 & x & 0 \\ 0 & x & 1 & 0 \\ x & 0 & 0 & 1 \end{pmatrix}}_{C_s} \right] \begin{pmatrix} 1 \\ 1 \\ 1 \\ 1 \end{pmatrix}$$

By selecting the "don't care" values accordingly

$$C_s = V_s V_s^\top \text{ in this case } V_s = \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ 0 & 1 \\ 1 & 0 \end{pmatrix}$$

Extracting diagonal entries of A

Proposition (Bekas, Kokiopoulou, Saad [BKS07])

Let $V_s \in \mathbb{R}^{n \times s}$ a matrix the columns of which are the probing vectors. If the i-th row of V_s is orthogonal to all rows j of V_s such that $\alpha_{ij} \neq 0$, then the diagonal entries of A are exactly

$$\delta(A) = [A \odot (VV^T)e] \oslash [(V_s^T V_s)e] = [(V_s \odot AV_s)e] \oslash [(V_s^\top V_s)e]$$

Extracting diagonal entries of A

Proposition (Bekas, Kokiopoulou, Saad [BKS07])

Let $V_s \in \mathbb{R}^{n\times s}$ a matrix the columns of which are the probing vectors. If the i-th row of V_s is orthogonal to all rows j of V_s such that $\alpha_{ij} \neq 0$, then the diagonal entries of A are exactly

$$\delta(A) = [A \odot (VV^T)e] \oslash [(V_s^T V_s)e] = [(V_s \odot AV_s)e] \oslash [(V_s^\top V_s)e]$$

[TS] proposed algorithm for creating V_s by coloring graph $G(A)$

$$V_s = \begin{pmatrix} & \\ & \end{pmatrix}$$

Extracting diagonal entries of A

Proposition (Bekas, Kokiopoulou, Saad [BKS07])

Let $V_s \in \mathbb{R}^{n\times s}$ a matrix the columns of which are the probing vectors. If the i-th row of V_s is orthogonal to all rows j of V_s such that $\alpha_{ij} \neq 0$, then the diagonal entries of A are exactly

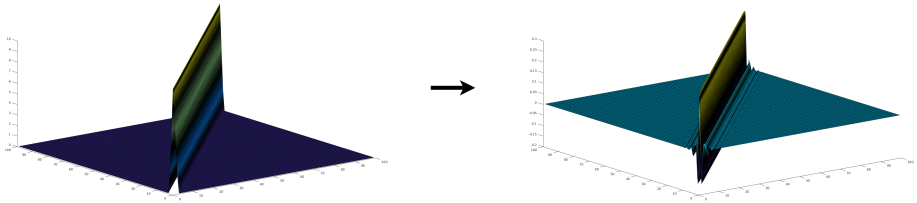
$$\delta(A) = [A \odot (VV^T)e] \oslash [(V_s^T V_s)e] = [(V_s \odot AV_s)e] \oslash [(V_s^\top V_s)e]$$

[TS] proposed algorithm for creating V_s by coloring graph $G(A)$

$$V_s = \begin{pmatrix} & \\ & \end{pmatrix} \quad \text{number of colors} = s$$

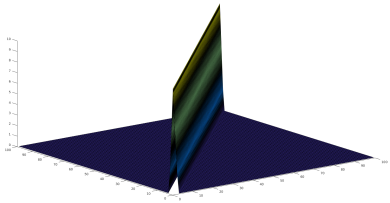
Approximating the sparsity pattern of inverse

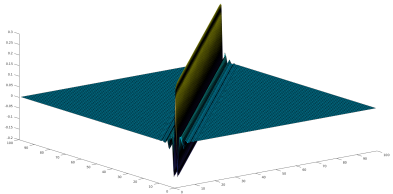
A $\quad\rightarrow\quad$ A^{-1}

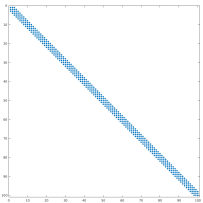


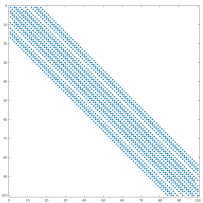
Approximating the sparsity pattern of inverse

A A^{-1}









Computing diagonal entries of the inverse

Algorithm by [TS]

Algorithm 1: Approximate diagonal of A^{-1}

1. Approximate $S(A^{-1})$ // sparsity pattern of A^{-1}
2. Create V_s by coloring $S(A^{-1})$
3. $\delta(A^{-1}) = [V_s \odot (A^{-1}V_s)e] \oslash [(V^\top V)e]$

end

Computing diagonal entries of the inverse

Algorithm by [TS]

Algorithm 1: Approximate diagonal of A^{-1}

1. Approximate $S(A^{-1})$ // sparsity pattern of A^{-1}
2. Create V_s by coloring $S(A^{-1})$
3. $\delta(A^{-1}) = [V_s \odot (A^{-1}V_s)e] \oslash [(V^\top V)e]$

end

Analysis

- Performance depends heavily on approximating $S(A^{-1})$

Computing diagonal entries of the inverse

Algorithm by [TS]

Algorithm 1: Approximate diagonal of A^{-1}

1. Approximate $S(A^{-1})$ // sparsity pattern of A^{-1}
2. Create V_s by coloring $S(A^{-1})$
3. $\delta(A^{-1}) = [V_s \odot (A^{-1}V_s)e] \oslash [(V^\top V)e]$

end

Analysis

- Performance depends heavily on approximating $S(A^{-1})$ $\rightarrow$ [Cho00]

Computing diagonal entries of the inverse

Algorithm by [TS]

Algorithm 1: Approximate diagonal of A^{-1}

1. Approximate $S(A^{-1})$ // sparsity pattern of A^{-1}
2. Create V_s by coloring $S(A^{-1})$
3. $\delta(A^{-1}) = [V_s \odot (A^{-1}V_s)e] \oslash [(V^\top V)e]$

end

Analysis

- Performance depends heavily on approximating $S(A^{-1})$ $\rightarrow$ [Cho00]
- Greedy coloring leads to small number of colors s

Computing diagonal entries of the inverse

Algorithm by [TS]

Algorithm 1: Approximate diagonal of A^{-1}

1. Approximate $S(A^{-1})$ // sparsity pattern of A^{-1}
2. Create V_s by coloring $S(A^{-1})$
3. $\delta(A^{-1}) = [V_s \odot (A^{-1}V_s)e] \oslash [(V^\top V)e]$

end

Analysis

- Performance depends heavily on approximating $S(A^{-1})$ $\rightarrow$ [Cho00]
- Greedy coloring leads to small number of colors s $\rightarrow$ entry decay is neccesary

Computing diagonal entries of the inverse

Algorithm by [TS]

Algorithm 1: Approximate diagonal of A^{-1}

1. Approximate $S(A^{-1})$ // sparsity pattern of A^{-1}
2. Create V_s by coloring $S(A^{-1})$
3. $\delta(A^{-1}) = [V_s \odot (A^{-1}V_s)e] \oslash [(V^\top V)e]$

end

Analysis

- Performance depends heavily on approximating $S(A^{-1})$ $\rightarrow$ [Cho00]
- Greedy coloring leads to small number of colors s $\rightarrow$ entry decay is neccesary
- step 3 requires solving s sparse linear systems

Computing diagonal entries of the inverse

Algorithm by [TS]

Algorithm 1: Approximate diagonal of A^{-1}

1. Approximate $S(A^{-1})$ // sparsity pattern of A^{-1}
2. Create V_s by coloring $S(A^{-1})$
3. $\delta(A^{-1}) = [V_s \odot (A^{-1}V_s)e] \oslash [(V^\top V)e]$

end

Analysis

- Performance depends heavily on approximating $S(A^{-1})$ $\rightarrow$ [Cho00]
- Greedy coloring leads to small number of colors s $\rightarrow$ entry decay is neccesary
- step 3 requires solving s sparse linear systems $\rightarrow$ use iterative MRHS solvers

1 Introduction

2 Extracting diagonal elements

3 Approximating diagonal blocks

4 Numerical experiments

5 Conclusions

Goal: Propose a computational framework for approximating diagonal blocks of the inverse

Goal: Propose a computational framework for approximating diagonal blocks of the inverse

Motivation:

- Expand approximation methods that rely on sparsity pattern
- Exploit decay in matrix functions by blocks
- Enabling better preconditioners

Goal: Propose a computational framework for approximating diagonal blocks of the inverse

Motivation:

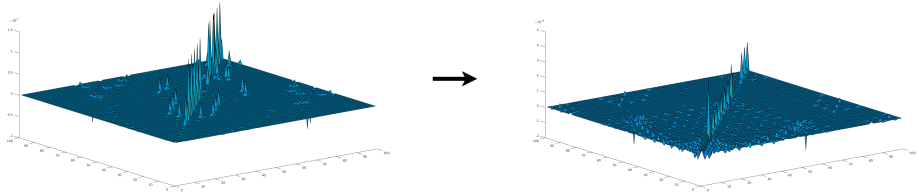
- Expand approximation methods that rely on sparsity pattern
- Exploit decay in matrix functions by blocks
- Enabling better preconditioners

Setting:

1. Matrix $A \in \mathbb{R}^{n \times n}$ and is invertible
2. A is symmetric but that condition is not necessary
3. Diagonal blocks of A are invertible

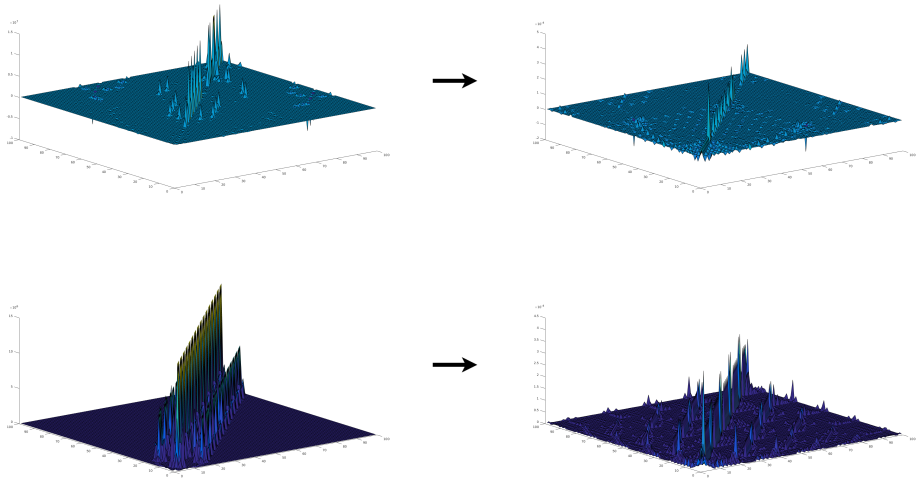
Block decay

A $\quad\rightarrow\quad$ A^{-1}



Block decay

A A^{-1}



Sparsity pattern of inverse

From Cayley-Hamilton theorem

$$A^{-1} = \frac{(-1)^{n-1}}{\det(A)}(\gamma_1 I + \gamma_2 A + \cdots + A^{n-1})$$

where γ_i for $i = 1, \ldots, n$ are the eigenvalues of the characteristic polynomial

Sparsity pattern of inverse

From Cayley-Hamilton theorem

$$A^{-1} = \frac{(-1)^{n-1}}{\det(A)}(\gamma_1 I + \gamma_2 A + \cdots + A^{n-1})$$

where γ_i for $i = 1, \ldots, n$ are the eigenvalues of the characteristic polynomial

$$\begin{aligned} S(A^{-1}) &= S(\frac{(-1)^{n-1}}{\det(A)}(\gamma_1 I + \gamma_2 A + \cdots + A^{n-1})) \\ &\approx S(\frac{(-1)^{n-1}}{\det(A)}(\gamma_1 I + \gamma_2 A + \cdots + A^{k})) \qquad \text{for } k < n-1 \end{aligned}$$

Sparsity pattern of inverse

From Cayley-Hamilton theorem

$$A^{-1} = \frac{(-1)^{n-1}}{\det(A)}(\gamma_1 I + \gamma_2 A + \cdots + A^{n-1})$$

where γ_i for $i = 1, \ldots, n$ are the eigenvalues of the characteristic polynomial

$$\begin{aligned} S(A^{-1}) &= S(\frac{(-1)^{n-1}}{\det(A)}(\gamma_1 I + \gamma_2 A + \cdots + A^{n-1})) \\ &\approx S(\frac{(-1)^{n-1}}{\det(A)}(\gamma_1 I + \gamma_2 A + \cdots + A^{k})) \qquad \text{for } k < n-1 \end{aligned}$$

Performance improvements:

1. $S(A^{-1}) \approx S(\sum_{i=0}^{k} A^i) \rightarrow$ good approximations when A is diagonally dominant [TS]

Sparsity pattern of inverse

From Cayley-Hamilton theorem

$$A^{-1} = \frac{(-1)^{n-1}}{\det(A)}(\gamma_1 I + \gamma_2 A + \cdots + A^{n-1})$$

where γ_i for $i = 1, \ldots, n$ are the eigenvalues of the characteristic polynomial

$$\begin{aligned} S(A^{-1}) &= S(\frac{(-1)^{n-1}}{\det(A)}(\gamma_1 I + \gamma_2 A + \cdots + A^{n-1})) \\ &\approx S(\frac{(-1)^{n-1}}{\det(A)}(\gamma_1 I + \gamma_2 A + \cdots + A^{k})) \qquad \text{for } k < n-1 \end{aligned}$$

Performance improvements:

1. $S(A^{-1}) \approx S(\sum_{i=0}^{k} A^i) \rightarrow$ good approximations when A is diagonally dominant [TS]
2. $S(A^{-1}) \approx S(A^k)$

Sparsity pattern of blocks

We process blocks in terms of their block norms. Similar with [Mol97] Assuming partitioning in blocks with the same size.

Sparsity pattern of blocks

We process blocks in terms of their block norms. Similar with [Mol97] Assuming partitioning in blocks with the same size.

Definition (Retriction function)

Let $\Phi_l : \mathbb{R}^{n\times n} \to \mathbb{R}^{N\times N}$ for $N = \frac{n}{l}$ where l is the block size of the partioning, such that

$$[\Phi_l(A)]_{ij} = \|A_{ij}\|_F$$

Sparsity pattern of blocks

We process blocks in terms of their block norms. Similar with [Mol97] Assuming partitioning in blocks with the same size.

Definition (Retriction function)

Let $\Phi_l : \mathbb{R}^{n\times n} \to \mathbb{R}^{N\times N}$ for $N = \frac{n}{l}$ where l is the block size of the partioning, such that

$$[\Phi_l(A)]_{ij} = \|A_{ij}\|_F$$

Can we generalize rapid decay phenomena for blocks?

Results on decay of blocks:

- [Mol97] Expressing the result of Demko [DMS84] in blocks
- [ELN17] Improvements on the decay of blocks

Block diagonal dominance

Let $A \in \mathbb{R}^{n \times n}$ the following partition in square blocks Ai,j of size B

$$A = \begin{pmatrix} A_{11} & A_{12} & \dots & A_{NN} \\ A_{21} & A_{22} & \dots & A_{NN} \\ \vdots & \vdots & \ddots & \vdots \\ A_{N1} & A_{N2} & \dots & A_{NN} \end{pmatrix}$$

Definition (Block diagonal dominance [FV62])

Let A be a partitioned matrix as shown above, then A is block diagonally dominant relative to the partitioning scheme if

$$(\|A_{ii}^{-1}\|)^{-1} \geq \sum_{\substack{j=1 \\ j \neq i}}^{N} \|A_{ij}\| \text{ for } i = 1, 2, \dots, N$$

Sparsity pattern of blocks

Approximate the sparsity pattern of $\Phi_l(A^k)$

$$S(\Phi_l(A)^k) \approx S(\Phi_l(A^k))$$

Matrix $\Phi_l(A)\mathbb{R}^{n/l \times n/l}$ and

$$[\Phi_l(A)]_{ij} \geq 0$$

Sparsity pattern of blocks

Approximate the sparsity pattern of $\Phi_l(A^k)$

$$S(\Phi_l(A)^k) \approx S(\Phi_l(A^k))$$

Matrix $\Phi_l(A)\mathbb{R}^{n/l \times n/l}$ and

$$[\Phi_l(A)]_{ij} \geq 0$$

We generate probing vectors $\{v_i\}$ for $\Phi(A)^k$ that translate to

$$U_i = v_i \otimes I_l$$

Sparsity pattern of blocks

Approximate the sparsity pattern of $\Phi_l(A^k)$

$$S(\Phi_l(A)^k) \approx S(\Phi_l(A^k))$$

Matrix $\Phi_l(A)\mathbb{R}^{n/l \times n/l}$ and

$$[\Phi_l(A)]_{ij} \geq 0$$

We generate probing vectors $\{v_i\}$ for $\Phi(A)^k$ that translate to

$$U_i = v_i \otimes I_l$$

Advantages:

1. $\Phi_l(A)$ has smaller size than A and is still sparse

Sparsity pattern of blocks

Approximate the sparsity pattern of $\Phi_l(A^k)$

$$S(\Phi_l(A)^k) \approx S(\Phi_l(A^k))$$

Matrix $\Phi_l(A)\mathbb{R}^{n/l \times n/l}$ and

$$[\Phi_l(A)]_{ij} \geq 0$$

We generate probing vectors $\{v_i\}$ for $\Phi(A)^k$ that translate to

$$U_i = v_i \otimes I_l$$

Advantages:

1. $\Phi_l(A)$ has smaller size than A and is still sparse
2. Good approximations for matrices that are close to block diagonally dominant

Sparsity pattern of blocks

Observations:

- Elements of $\Phi_l(A)$ are non negative $\rightarrow$ thresholding is important to ensure small number of probing vectors

Sparsity pattern of blocks

Observations:

- Elements of $\Phi_l(A)$ are non negative $\rightarrow$ thresholding is important to ensure small number of probing vectors

Having created sets of probing vectors $\{U_i\}$ we compute the diagonal blocks $\Delta(A^{-1})$ as

$$\Delta(A^{-1}) = \sum_{i=0}^{s} U_i \bullet (A^{-1} U_i)$$

We use the notation $\bullet$ to define multiplication with blocks of size $l \times l$

Sparsity pattern of blocks

Observations:

- Elements of $\Phi_l(A)$ are non negative $\rightarrow$ thresholding is important to ensure small number of probing vectors

Having created sets of probing vectors $\{U_i\}$ we compute the diagonal blocks $\Delta(A^{-1})$ as

$$\Delta(A^{-1}) = \sum_{i=0}^{s} U_i \bullet (A^{-1} U_i)$$

We use the notation $\bullet$ to define multiplication with blocks of size $l \times l$
For the case of diagonal approximation $l = 1$ and the known formula holds

$$\delta(A^{-1}) = \sum_{i=0}^{s} v_i \odot (A^{-1} v_i)$$

ϵ-heatmap

Definition (ϵ-heatmap)

Given $A \in \mathbb{R}^{n \times n}$ and restriction function $\Phi_L(A)$, for $L = [l_1, \ldots, l_p]$ and $\epsilon > 0$ we define ϵ-heatmap of A to be every $H_\epsilon \in \mathbb{R}^{n \times n}$ such that for

$$[E]_{ij} = \begin{cases} [\Phi_l(A - H_\epsilon)]_{ij} / [\Phi(A)]_{ij} & \text{where } [\Phi(A)]_{ij} \geq 0 \\ 0, & \text{otherwise} \end{cases}$$

the following is true

$$\max_{i,j=1,2,\ldots,l} \{[E]_{ij}\} \leq \epsilon$$

1 Introduction

2 Extracting diagonal elements

3 Approximating diagonal blocks

4 Numerical experiments

5 Conclusions

Numerical experiments

Matrices from sparse Suitesparse matrix collection

For $\Delta(A^{-1}), H_\epsilon \in \mathbb{R}^{n \times s*l}$ define the vector $r \in \mathbb{R}^{n/l}$

$$r_i = \frac{\Phi_l(A^{-1} - H_\epsilon)_i}{\Phi_l(\Delta(A^{-1}))} \quad \text{for i} = 1, \ldots, \text{n/l}$$

$$e = \|r\|_\infty$$

Numerical experiments

Matrices from sparse Suitesparse matrix collection

For $\Delta(A^{-1}), H_\epsilon \in \mathbb{R}^{n \times s*l}$ define the vector $r \in \mathbb{R}^{n/l}$

$$r_i = \frac{\Phi_l(A^{-1} - H_\epsilon)_i}{\Phi_l(\Delta(A^{-1}))} \quad \text{for i} = 1, \ldots, \text{n/l}$$

$$e = \|r\|_\infty$$

Apply the same approach to exp(A) using chebyshev approximation

$$r_i = \frac{\Phi_l(f(A) - H_\epsilon)_i}{\Phi_l(\Delta(f(A)))} \quad \text{for i} = 1, \ldots, \text{n/l}$$

$$e = \|r\|_\infty$$

Numerical experiments

matrix	block size	error	n	s
msc00726	10	6e-2	760	340
msc01050	3	5e-2	1050	36
msc01440	5	1e-4	1440	640
bcsstm12	3	3e-4	1473	123
bcsstk12	5	1e-1	1473	335
bcsstm11	3	1e-1	1473	190
bcsstk14	5	1e-5	1806	620
plat1919	10	8e-2	1919	760

Table: approximation of $\Delta(A^{-1})$

Numerical experiments

matrix	block size	error	n	s
msc00726	10	2e-4	760	340
msc01050	3	1e-2	1050	36
msc01440	5	1e-4	1440	240
bcsstm12	3	2e-6	1473	126
bcsstk12	3	1e-6	1473	159
bcsstm11	3	1e-5	1473	159
bcsstk14	5	6e-2	1806	620
plat1919	10	1e-3	1919	760

Table: approximation of $\Delta(e^A)$

Application to preconditioning

matrix	n	block size	approx	Jacobi	no precond
plat1919	1919	25	274	1919	1919
bcsstk11	1445	25	671	803	1445
bcsstk12	1445	5	800	1133	1445
msc01440	1440	25	273	446	1395
bcsstk19	817	25	158	259	816
bcsstk03	112	5	85	110	110
bcsstm12	1473	25	38	87	810
bcsstk13	1473	25	1217	1092	1031
bcsstk14	1806	25	299	59	1511

Figure: Number of iterations for approximate inverse, jacobi and no preconditioning

Applications to preconditioning

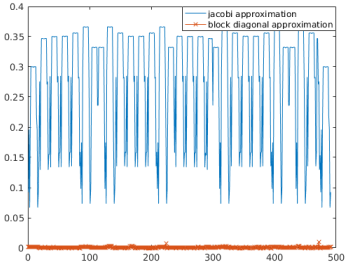


Figure: bcsstm

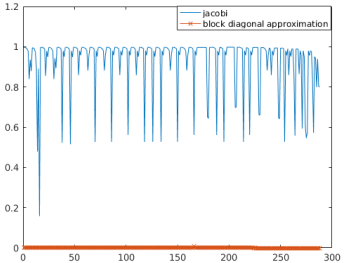


Figure: msc01440

1 Introduction

2 Extracting diagonal elements

3 Approximating diagonal blocks

4 Numerical experiments

5 Conclusions

Conclusions

- Developed an approach based on the sparsity pattern of blocks
- Defined the ϵ-heatmap matrix as a approximation of blocks
- On small size numerical test gained reasonably good approximation

Ongoing work

- Systematic approach for numerical dropping
- Coarsening based techniques for increasing accuracy
- Ways to automatically set block size depending on the application
- Generalizing to multiple block size case
- Experimental study on large matrices

Thank you!

References I

C. Bekas, E. Kokiopoulou, and Y. Saad, *An estimator for the diagonal of a matrix*, Appl. Numer. Math. **57** (2007), no. 11-12, 1214–1229.

E. Chow, *A priori sparsity patterns for parallel sparse approximate inverse preconditioners*, SIAM J. Scientific Computing **21** (2000), no. 5, 1804–1822.

S. Demko, W. F. Moss, and P. W. Smith, *Decay rates for inverses of band matrices*, Math. Comp. **43** (1984), 491–499.

C. Echeverría, J. Liesen, and R. Nabben, *Block diagonal dominance of matrices revisited: bounds for the norms of inverses and eigenvalue inclusion sets*, arXiv.org (2017).

D. G. Feingold and R. S. Varga, *Block diagonally dominant matrices and generalizations of the Gerschgorin circle theorem*, Pacific J. Math. **12** (1962), 1241–1250.

References II

- Luca Molinari, *Transfer matrices and tridiagonal-block hamiltonians with periodic and scattering boundary conditions*, Journal of Physics A: Mathematical and General **30** (1997), no. 3, 983.
- J. M. Tang and Y. Saad, *A probing method for computing the diagonal of a matrix inverse*, Numerical Linear Algebra with Applications **19**, no. 3.