arXiv:2002.04989v4 [cs.PF] 16 Jun 2020

# Eigenvector Component Calculation Speedup over *NumPy* for High-Performance Computing

Shrey Dabhi[1,2][0000−0002−4364−9892] and Manojkumar Parmar[1,3][0000−0002−1183−4399]

[1] Robert Bosch Engineering and Business Solutions Private Limited, Bengaluru, India
manojkumar.parmar@bosch.com
[2] Department of Computer Science and Engineering, Institute of Technology, Nirma University, Ahmedabad, India
16bit039@nirmauni.ac.in
[3] HEC Paris, Jouy-en-Josas Cedex, France
manojkumar.parmar@hec.edu

**Abstract.** Applications related to artificial intelligence, machine learning, and system identification simulations essentially use eigenvectors. Calculating eigenvectors for very large matrices using conventional methods is compute-intensive and renders the applications slow. Recently, the *Eigenvector-Eigenvalue Identity* formula promising significant speedup was identified. We study the algorithmic implementation of the formula against the existing state-of-the-art algorithms and their implementations to evaluate the performance gains. We provide a first of its kind systematic study of the implementation of the formula. We demonstrate further improvements using high-performance computing concepts over native *NumPy* eigenvector implementation which uses LAPACK and BLAS.

**Keywords:** Eigenvalues · Eigenvectors · Eigenspace · Symmetric matrix · Hermitian matrix · NumPy · LAPACK · BLAS · High-performance computation · Vectorization · Parallelization

## 1 Introduction

In mathematics, eigenvectors are fundamental to many matrix operations. They form the base of all the dimensionality reduction operations in Artificial Intelligence and Machine Learning. Moreover, eigenvectors are of extreme importance in system simulation and calculation of many real-world phenomena. They are the key to analyzing the physics of rotating bodies, the stability of physical structures, oscillations of vibrating bodies, computational biology, economics, etc.

In linear algebra, an eigenvector is the characteristic vector of a linear transformation. It is a nonzero vector that changes at most by a scalar factor when that linear transformation is applied to it. The corresponding eigenvalue is the

Accepted at ICRTC-2020, to be published in Springer LNNS

factor by which the eigenvector is scaled. Geometrically, an eigenvector, corresponding to a real nonzero eigenvalue, points in a direction in which it is stretched by the transformation and the eigenvalue is the factor by which it is stretched. If the eigenvalue is negative, the direction is reversed. This complexity causes the computations and thereby the applications to slow down.

In the following sections, we discuss our approach and the processes we applied to conclude. In Section 2 we explain the concept of eigenvalues and eigenvectors in brief. We also discuss the drawbacks of conventional methods and the current state-of-the-art algorithms. We then review the existing implementations of the *Eigenvector-Eigenvalue Identity* formula to formulate our strategy for the systematic improvement and evaluation. In Section 3 we shed light on the actual experimental setup and present algorithms for the baseline (given by Denton et al. [1][2]) and the most performant variants of the implementations out of all the variants that we tested. We also present our thoughts on why exhibit Algorithm 2 gives the highest amount of speed up over *NumPy*'s implementation. In Section 4 we discuss the results of our experiments. In Section 5 we give the direction for future researchers to improve over our work.

## 2 Related Works

The eigenvalues of an $n \times n$ matrix $A$ are the roots of the characteristic equation:

$$det(A - \lambda I) = 0 \tag{1}$$

where $I$ is an $n \times n$ identity matrix Eq. (1) has $n$ solutions. The polynomial left-hand side of the characteristic equation is known as the characteristic polynomial. Hermitian matrices have real eigenvalues, i.e. the roots of Eq (1) are real. A real and symmetric matrix is simply a special case of a hermitian matrix. Traditionally, eigenvectors are calculated substituting the value of $\lambda$ in the characteristic polynomial. These polynomials are $n^{th}$-order polynomials and solving them directly to calculate the eigenvectors can be very compute-intensive.

Currently, the state-of-the-art algorithm for the computation of eigenvectors of hermitian matrices is the divide and conquer algorithm proposed by Cuppen et al. [3]. There are only 2 known implementations of this algorithm: one written in FORTRAN [4] and another one written in C language [5]. This algorithm first converts the hermitian or real symmetric matrix into a tridiagonal matrix, and then recursively divides the matrix into 2 tridiagonal matrices until the lowest possible size is reached. Then it calculates the eigenvectors for the smallest matrices and adds rank 1 corrections while accumulating the results to compensate for the loss of information incurred during the division of the matrix. The original recursive algorithm was later improved upon by various researchers for increasing the speed [6][7][8].

A very simple formula has been recently brought to light again after it was first discovered by Jacobi in 1834 and then subsequently rediscovered by many independent research teams throughout history. The new formula is proposed

by Denton et al. [1] and formally named as the "*Eigenvector-Eigenvalue Identity*". They have traced the appearance of identity (2) throughout history and to prevent further disappearance and rediscovery of the identity, decided to formalize it finally in the form mentioned above. Various mathematical proofs of the identity have been covered in [1], however, computational gains have not been investigated systematically. The formal definition of identity has been given below.

If $A$ is an $n \times n$ hermitian matrix with eigenvalues $\lambda_i(A), \ldots, \lambda_n(A)$ and $i, j = 1, \ldots, n$, then the $j^{th}$ component $v_{i,j}$ of a unit eigenvector $v_i$ associated to the eigenvalue $\lambda_i(A)$ is related to the eigenvalues $\lambda_1(M_j), \ldots, \lambda_{n-1}(M_j)$ of the minor $M_j$ of $A$ formed by removing the $j^{th}$ row and column by the formula

$$|v_{i,j}|^2 = \frac{\prod_{k=1}^{n-1}(\lambda_i(A) - \lambda_k(M_j))}{\prod_{k=1,k\neq i}^{n}(\lambda_i(A) - \lambda_k(A))} \tag{2}$$

In its current form identity (2) does not provide the direction of the components of the eigenvectors, but just the magnitude. This makes it unsuitable for applications requiring directions. But as mentioned in [1], it is possible to infer the directions of the components through various methods. It has been suggested by Ashok et al. [9] that for small matrices, the directions of eigenvectors can often be inferred by direct inspection of the eigenvector equation $Av_i = \lambda_i(A)v_i$. In general, we can apply Eq. (2) on multiple bases to retrieve the direction.

The basic implementation for identity formula (Eq. (2)) was provided by Denton et al. [1][2]. We improved upon this implementation by vectorizing, parallelizing and adding batched processing capabilities. Despite the observations of Amdahl's law, we did not achieve tangible speed-up in execution time in certain cases, due to the possible overhead of thread management.[4]

In this paper, we focus on implementing identity formula (2) as efficiently as possible and attempt to match the accuracy, speed, and robustness of the state-of-the-art implementations. The majority of the popular machine learning and linear algebra libraries use the low-level FORTRAN implementations from BLAS packages and LAPACK under the hood, for eigendecomposition. All the machine learning libraries just provide high-level wrappers which abstract out a lot of complexities and provide sensible defaults for the required hyper-parameters to achieve the desired speed and accuracy from the eigendecomposition algorithms from LAPACK.

## 3 Experiments

The baseline implementation of identity formula Eq (2) as provided by Denton et al. [2] is quite slow as it repeatedly computes eigenvalues of the same set of

[4] Amdahl's Law gives the theoretical speedup in the execution time of a routine at a fixed workload that can be expected out of a system whose resources are improved. In our case, we are increasing the computational power of the system by allowing our routine to utilize multiple CPU cores in parallel.

matrices, implying that it does not provide any advantage over other existing algorithms. The pseudo-code for the baseline implementation is provided in exhibit Algorithm 1. We used concepts of high-performance computing to improve on it. We conceived, designed and implemented 5 variations of the algorithm and the baseline version. The comparison of the performance of each version for 3 different tasks is discussed in the results section. The pseudo-code for the best performing variant is given in exhibit Algorithm 2.

Charis et al. [10] provide an alternative implementation for the formula (2). They calculate only a particular eigenvector for the given matrix. A similar implementation in Python does not provide improvements over *NumPy* [11][12].[5] We then decided to test the speed up in the calculation of one of the components of a particular eigenvector of the given matrix.

**Algorithm 1:** Baseline pseudo-algorithm

```
Function EigenComponentBaseline(matrix, i, j):
    n ← matrix.shape
    minor ← del(matrix, i)
    matrixEV ← EigenValues(matrix)
    minorEV ← EigenValues(minor)
    numerator ← 1.0
    for k = 0; k < n − 1; k ← k + 1 do
        numerator ← numerator ∗ (matrixEV[j] − matrixEV[k])
    end
    denominator ← 1.0
    for k = 0; k < n; k ← k + 1 do
        if k ≠ j then
            denominator ← denominator ∗ (matrixEV[j] − minorEV[k])
        end
    end
    component ← numerator ÷ denominator
return component
```

The experiments to measure the speed-up were carried out on a high-performance workstation running Windows 10 operating system. It is equipped with 32 GiB of random access memory and Intel Xeon® CPU E3-1270 v6 @ 3.8 GHz with 4 physical cores and 8 logical cores.

We use Python 3.7.6 compiled for MSC v.1916 for a 64-bit processor. The computation was run 10 times for each matrix size using each variation of the implementation. The execution time was measured using cProfile utility provided in Python's standard library. The mean values for the 10 runs have been reported.

For calculating all the components of all the eigenvectors, the baseline implementation is the slowest. For an $n \times n$ matrix, the algorithm calls the function

[5] One of the reasons can be the lack of native parallel for loops in Python.

**Algorithm 2:** Optimized pseudo-algorithm

```
Function EigenComponentOptimized(matrix, i, j, batchSize):
    n ← matrix.shape
    minor ← del(matrix, i)
    matrixEV ← EigenValues(matrix)
    eigenValue ← matrixEV[j]
    matrixEV ← del(matrixEV, j)
    minorEV ← EigenValues(minor)
    batches, nBatch ← PrepareBatches(matrixEV, minorEV, batchSize)
    for k = 1; k < nBatch; k ← k + 1 do
        tNumer[k], tDenom[k] ← dispatch(BatchProcessor(batches[k]))
    end
    component ← 1.0
    for k = 0; k < nBatch; k ← k + 1 do
        component ← component ∗ (join(tNumer[k]) ÷ join(tDenom[k]))
    end
return component
```

$2n^3$ times as shown in exhibit Algorithm 1. By vectorizing the algorithm we can reduce the number of calls, therefore providing speed up.

We expected that parallel computation would reduce the execution time, but the possible overhead of thread creation and management leads to an increase in the time complexity.

The baseline and vectorized versions have limited performance for very large matrices (of the order of $150 \times 150$ and greater). One of the reasons could be the intermediate calculations leading to a datatype underflow or overflow. We batched all the calculations to combat this issue and also checked if computing all the batches in parallel would lead to any significant performance gains as well.

LAPACK's implementation of Cuppen's algorithm [3] used by *NumPy* [11][12] only returns the complete set of all the eigenvectors. Hence, we further explored if we could benefit from the calculation of single vectors or a single component of any one vector.

## 4 Results

In Fig. 1(a) and 1(b), identity refers to the batched vectorized implementation of the Eq. (2) and identity parallelized refers to exhibit Algorithm 2. In all the figures, X-axis represents the size of the matrix for which we are calculating the eigendecomposition and the Y-axis represents the time taken in seconds.

Fig. 1(a) shows the speed up over *NumPy*'s implementation for the calculation of a single component of a particular eigenvector for the given matrix. Fig. 1(b) shows the performance of our implementation against the current state-of-the-art, i.e. *NumPy*, for calculating a complete eigenvector for the given matrix.

Fig. 1(c) and Fig. 1(d) show how we step-by-step implemented the concepts of high-performance computing to improve the performance of the formula. We achieved speedup with each iteration of optimization and ended up with exhibit Algorithm 2 as the most optimized and performant implementation of the formula (2).

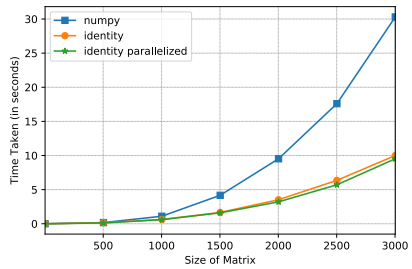


(a) For one of the components of an eigenvector

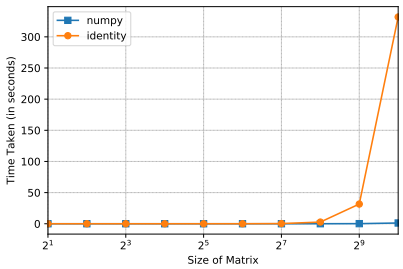


(b) For all the components of an eigenvector

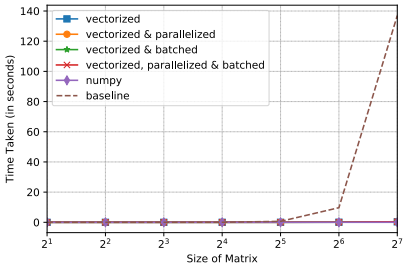


(c) For all the eigenvectors

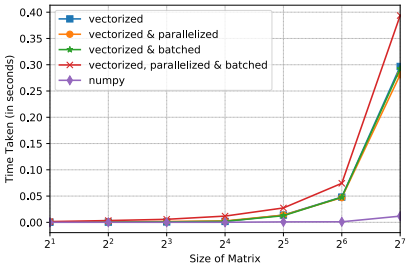


(d) For all the eigenvectors

Fig. 1: Comparison of the execution time of all the variants of the algorithm and LAPACK's implementation of Cuppen's algorithm

From Table 1 we can infer that by using Eq. (2) and exhibit Algorithm 2 with an increase in the size of the matrix, we can achieve up to 4.5× speed up over *NumPy*.

The formula is more efficient than the state-of-the-art algorithms only when we need a few components of a particular eigenvector. Currently, only the applications such as web indexing, web search, signal pre-processing, etc. which only require partial eigenvectors, can benefit from the implementation of this identity.

[6] Always computes the entire set

Table 1: Time in seconds for calculating 1 eigenvector component

| Size of matrix | NumPy[6] | Exhibit Alg. 2 |
|---|---|---|
| 2 | 0.000057 | 0.000233 |
| 502 | 0.170962 | 0.119722 |
| 1002 | 1.100800 | 0.595935 |
| 1502 | 4.165680 | 1.603579 |
| 2002 | 9.522480 | 3.212001 |
| 2502 | 17.632952 | 5.707374 |
| 3002 | 30.342656 | 9.522745 |
| 3502 | 47.404600 | 11.886900 |
| 4002 | 69.955400 | 16.775200 |
| 4502 | 98.680800 | 22.943500 |
| 5002 | 134.324200 | 30.547200 |
| 5502 | 177.629000 | 39.741700 |
| 6002 | 229.338600 | 50.682400 |

## 5 Future Work

It might be possible to further improve the performance if the formula is implemented in a lower-level language like C or FORTRAN to reduce overheads. LAPACK has also been written in FORTRAN, and therefore it can achieve such high performance. LAPACK algorithms also benefit from clever machine level multi-threading, which is extremely limited in a higher level language like Python, despite that we are getting a glimpse of the potential improvements as seen in Fig. 1 (a).

## 6 Summary

The *Eigenvector-Eigenvalue Identity* formula theoretically holds a promise to achieve speedup over the current state-of-the-art algorithms for the calculation of eigenvectors. We tried to realize this promise in an optimized algorithm. Our implementation can achieve up to 4.5× speed up over *NumPy* for applications requiring partial eigenvectors. We are of the opinion that this algorithm can further provide a speedup if implemented using lower-level languages.

## Acknowledgment

We want to thank Sri Krishnan V, Mohan B V, Palak Pareshkumar Sukharamwala and Tanya Motwani from Robert Bosch Engineering and Business Solutions Private Limited, India, for their valuable comments, contributions and continued support to the project. We are grateful to all experts for providing us with their valuable insights and informed opinions ensuring completeness of our study.

## References


1. Denton, P.B., Parke, S.J., Tao, T., Zhang, X.: Eigenvectors from Eigenvalues: a survey of a basic identity in linear algebra (2019), https://arxiv.org/abs/1908.03795
2. Denton, P.: Peterdenton/eigenvector-eigenvalue-identity (Jan 2020), https://github.com/PeterDenton/Eigenvector-Eigenvalue-Identity
3. Cuppen, J.J.M.: A divide and conquer method for the symmetric tridiagonal eigenproblem. Numerische Mathematik **36**(2), 177–195 (Jun 1980). https://doi.org/10.1007/BF01396757, https://doi.org/10.1007/BF01396757
4. Anderson, E., Bai, Z., Bischof, C., Blackford, S., Demmel, J., Dongarra, J., Du Croz, J., Greenbaum, A., Hammarling, S., McKenney, A., Sorensen, D.: LAPACK Users' Guide. Society for Industrial and Applied Mathematics, Philadelphia, PA, third edn. (1999)
5. Masyagin, M.: masyagin1998/divide-and-conquer-eigenvalues (Dec 2019), https://github.com/masyagin1998/divide-and-conquer-eigenvalues
6. Rutter, J.D.: A serial implementation of cuppen's divide and conquer algorithm for the symmetric eigenvalue problem. Tech. Rep. UCB/CSD-94-799, EECS Department, University of California, Berkeley (Feb 1994), http://www2.eecs.berkeley.edu/Pubs/TechRpts/1994/6315.html
7. Gu, M., Eisenstat, S.C.: A divide-and-conquer algorithm for the symmetric tridiagonal eigenproblem. SIAM Journal on Matrix Analysis and Applications **16**(1), 172–191 (1995). https://doi.org/10.1137/S0895479892241287, https://doi.org/10.1137/S0895479892241287
8. Tisseur, F., Dongarra, J.: A parallel divide and conquer algorithm for the symmetric eigenvalue problem on distributed memory architectures. SIAM Journal on Scientific Computing **20**(6), 2223–2236 (1999). https://doi.org/10.1137/S1064827598336951, https://doi.org/10.1137/S1064827598336951
9. Mukherjee, A.K., Datta, K.K.: Two new graph-theoretical methods for generation of eigenvectors of chemical graphs. Proceedings of the Indian Academy of Sciences - Chemical Sciences **101**(6), 499–517 (Dec 1989). https://doi.org/10.1007/BF02880817, https://doi.org/10.1007/BF02880817
10. Gyurgyik, C.: cgyurgyik/eigenvectors-from-eigenvalues (Dec 2019), https://github.com/cgyurgyik/eigenvectors-from-eigenvalues/
11. Oliphant, T.E.: A guide to NumPy, vol. 1. Trelgol Publishing USA (2006)
12. van der Walt, S., Colbert, S.C., Varoquaux, G.: The NumPy array: A structure for efficient numerical computation. Computing in Science Engineering **13**(2), 22–30 (March 2011). https://doi.org/10.1109/MCSE.2011.37