Complexity of decomposing a symmetric matrix as a sum of a diagonal and low-rank matrix

Stephen Vavasis

University of Waterloo
Combinatorics & Optimization

Joint work with Levent Tunçel (Waterloo) and Jingye Xu (Georgia Tech)

Problems under consideration

P1 Given $A \in \mathbb{S}^n_+$ and r, write $A = D + R$, where $D \in \mathbb{S}^n_+$ is diagonal and $R \in \mathbb{S}^n_+$ has rank$\leq r$.

P2 Given $A \in \mathbb{S}^n$ and r, write $A = D + R$, where $D \in \mathbb{S}^n$ is diagonal and $R \in \mathbb{S}^n_+$ has rank$\leq r$.

Notation:

- $\mathbb{S}^n = n \times n$ symmetric matrices,
- $\mathbb{S}^n_+ =$ positive semidefinite symmetric matrices,
- $\mathbb{S}^n_{++} =$ positive definite symmetric matrices

Applications

- Problems P1 and P2 have appeared in various bodies of literature for almost a century.
- Example application of P1: Given an empirically measured covariance matrix A of n stock prices, seek interpretation of A as a sum of r market factors, $r \ll n$, that influence all stock prices plus n independent stock variabilities.
- Example application of P2: fit an ellipsoid in $\mathbb{R}^n$ through n given data points (Saunderson, Chandrasekaran, Parrilo, Willsky 2011)

Related work

- Albert (1944): Factor analysis
- Saunderson et al. (2011): SDP relaxation
- Wu et al. (2020): Block coordinate descent
- Gao & Absil (2022): Manifold optimization
- Recht & Ré (2023): Stochastic gradient descent

Our results

- P1 and P2 are solvable in time polynomial in n for fixed r. (Superexponential in r.)
- P1 and P2 are NP-hard. Furthermore, even if we allow approximate data and approximate solutions, they remain NP-hard.
- P2 is $\exists\mathbb{R}$-complete. (This result assumes exact data and exact solution.)

Solving P2 in the case $r = 1$ (I)

- To illustrate our algorithm, consider the special case of P2 when $r = 1$: Given $A \in \mathbb{S}^n$, find decomposition $A = D + \boldsymbol{u}\boldsymbol{u}^T$ for some $\boldsymbol{u} \in \mathbb{R}^n$.
- WLOG A is hollow, i.e., $A(i, i) = 0$ $\forall i = 1, \ldots, n$
- Clearly $A(i, j) = u_i u_j$ $\forall i \neq j$ and $D(i, i) = -u_i^2$ $\forall i$
- Assume $A(2 : n, 1) \neq \mathbf{0}$ (else reduce problem to $A(2 : n, 2 : n)$)

Solving P2 in the case $r = 1$ (II)

- Assumption implies $u_1 \neq 0$ and $D(1,1) < 0$
- $A(i,1) = u_1 u_i \ \forall j > 1$ and $A(i,j) = u_i u_j$ $\forall i > j > 1$.
- Yields equation: $-A(1,i)A(1,j)/D(1,1) = A(i,j) \ \forall i > j > 1$ (numerator is $u_1^2 u_i u_j$; denominator is $-u_1^2$).
- $\Rightarrow$ many linear equations for $x := 1/D(1,1)$
- Solve any one of these equations to obtain $D(1,1)$; let $u_1 = \sqrt{-D(1,1)}$.
- Obtain $D(i,i) = -u_i^2$, $i > 1$, via $-A(1,i)^2/D(1,1) = (u_1 u_i)^2/u_1^2 = u_i^2$.

Solving P2 in the case $r = 1$ (III)

- Algorithm on previous slide fails if all linear equations are $0 \cdot x = 0$.
- This happens iff all but one entry (say $A(1,2)$) of $A(1, 2:n)$ are zeroes.
- This can happen only if $\boldsymbol{u}(3:n) = \mathbf{0}$.
- In this case, problem reduces to 2×2 case, easily handled.
- The 2×2 case, though trivial, requires a solution of both equations and inequalities.

Solving P2 in the case $r > 1$

- Full algorithm in our paper proposes algorithm for rank-r case.
- For 'generic' data, all entries of D are found by solving overdetermined linear equations.
- But for nongeneric data (interesting cases that include hard instances), one obtains $O(n^r)$ polynomial systems each with $O(r^2)$ variables and $O(r^2)$ constraints.

NP-hardness of P1 and P2

- Our reductions are from the problem of testing graph 3-colorability.
- Recall: a graph is *3-colorable* if there is an assignment of colors red, green, blue to each node such that there is no monochromatic edge.
- Proved by Garey, Johnson & Stockmeyer (1976) that deciding whether a graph is 3-colorable is NP-complete.

Partially specified matrices

- Start from a *partially specified* matrix $A \in (\mathbb{R} \cup \{*\})^{m \times n}$.
- A *completion* of A is matrix $A^{\#} \in \mathbb{R}^{m \times n}$ such that $A^{\#}(i,j) = A(i,j)$ for all (i,j) such that $A(i,j) \neq *$.
- Given a partially specified A, let $\mathcal{C}(A)$ be the set of all its completions.
- P2 can be described as: Given a symmetric $A \in (\mathbb{R} \cup \{*\})^{n \times n}$ in which $A(i,j) = * \Leftrightarrow i = j$, find a low-rank semidefinite element of $\mathcal{C}(A)$.

Peeters result

- Let $\hat{*}$ denote an unspecified entry that must be filled in with a nonzero number.
- Peeters (1996) proved: Given a graph G, one can construct a partially specified symmetric matrix B all of whose diagonal entries are $\hat{*}$. Graph G is 3-colorable iff there exists a completion B whose rank is 3.
- Our NP-hardness proofs all rely on Peeters' construction.

$G =$ 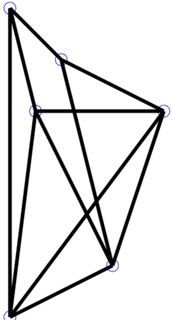

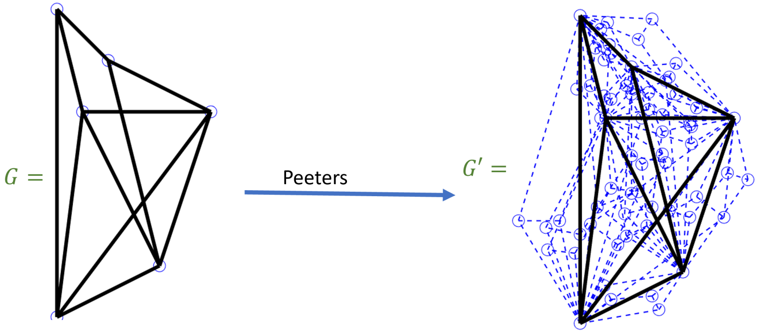

$G =$
Peeters
$G' =$

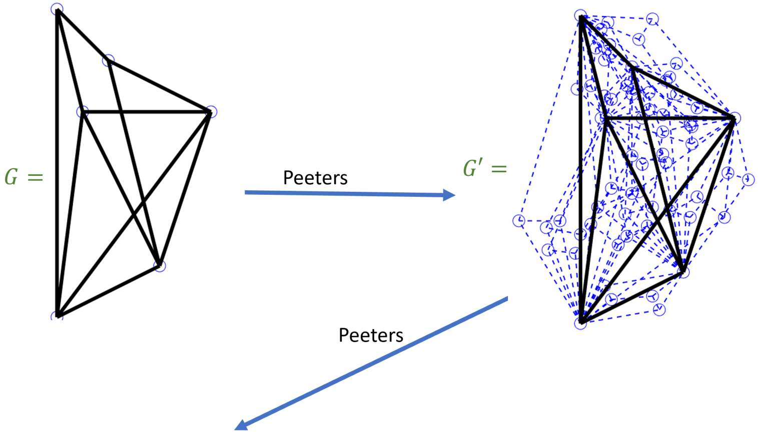

$$B = \begin{pmatrix} \hat{*} & 0 & * & \cdots & 0 \\ 0 & \hat{*} & 0 & \cdots & * \\ * & 0 & \hat{*} & & \\ \vdots & \vdots & & \ddots & \\ 0 & * & & & \hat{*} \end{pmatrix}$$

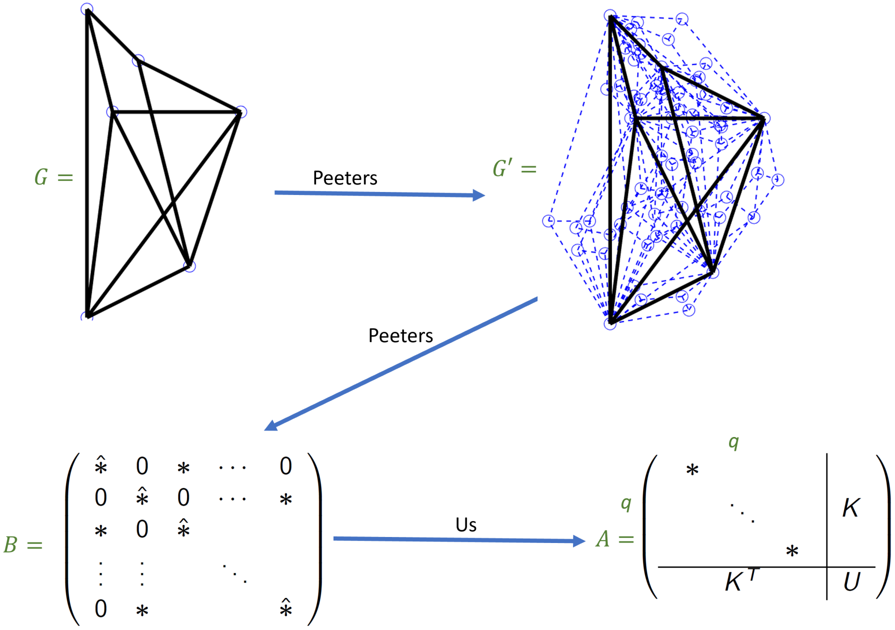

G =
Peeters
G′ =
Peeters
B =
Us
A =
q
q
K
K^T
U

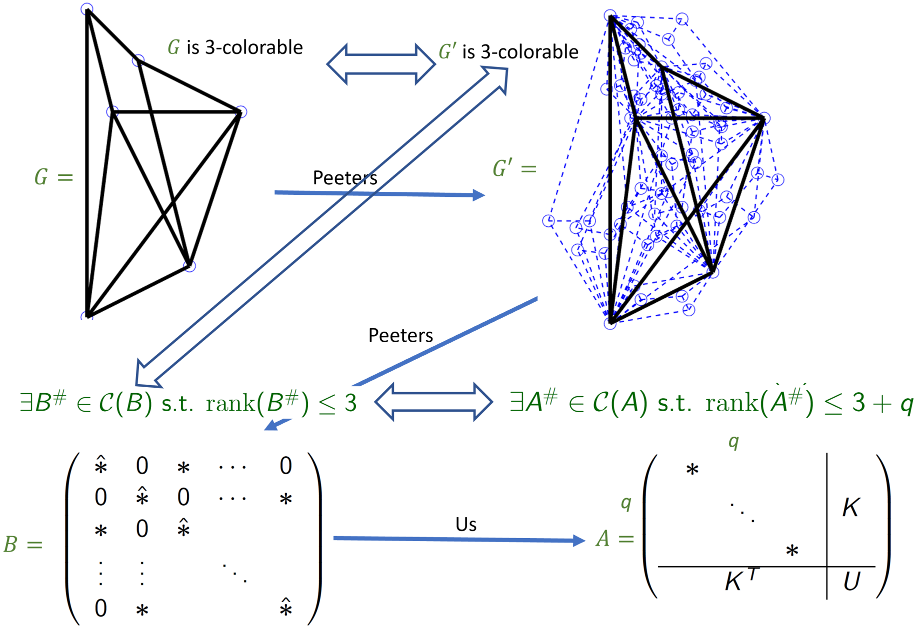

G is 3-colorable
G′ is 3-colorable
G =
Peeters
G′ =
Peeters
∃B# ∈ C(B) s.t. rank(B#) ≤ 3
∃A# ∈ C(A) s.t. rank(A#) ≤ 3 + q
B =
Us
A =
q
q
K
K^T
U

Construction of K and U

- K is a node-edge adjacency matrix, two 1's per column
- U constructed as follows:

$$B = \begin{pmatrix} \hat{*} & 0 & * & \cdots & 0 \\ 0 & \hat{*} & 0 & \cdots & * \\ * & 0 & \hat{*} & & \\ \vdots & \vdots & & \ddots & \\ 0 & * & & & \hat{*} \end{pmatrix} \to U = \begin{pmatrix} H & Z & S & \cdots & Z \\ Z & H & Z & \cdots & S \\ S & Z & H & & \\ \vdots & \vdots & & \ddots & \\ Z & S & & & H \end{pmatrix} \text{ where}$$

$$H = \begin{pmatrix} * & 1 & 1 \\ 1 & * & 1 \\ 1 & 1 & * \end{pmatrix}, Z = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, S = \begin{pmatrix} 2 & 2 & 2 \\ 2 & 2 & 2 \\ 2 & 2 & 2 \end{pmatrix}.$$

How this works

- The entries of the upper left diagonal block can be chosen so that, after its elimination, the '2' entries in U are decreased (e.g., to 0 or 1).
- The rank of the completion is equal to q plus the rank of the Schur complement after elimination of the diagonal block.
- In order for rank of the Schur complement to be ≤ 3:
  - The diagonal blocks of the Schur complement must be filled in with 1's.
  - The 0-1 pattern must correspond to three color classes.

NP-hardness of P1

- Similar construction as in P2 works, except we place large entries on the diagonal instead of 0's.
- This is because P1 can only subtract elements from the diagonal ($R = A - D$ in P1, where $D \in \mathbb{S}_{+}^{n}$)
- Same Schur complement argument applies.

Approximate P1 and P2

- Approximate P1: Given $A \in \mathbb{S}_+^n$ and a promise that there exists a rank-r semidefinite R_0 and positive definite diagonal D_0 such that $\|A - D_0 - R_0\| \le \epsilon$ (A, r, ϵ given; D_0, R_0 not given).
- Find positive definite diagonal D, semidefinite matrix R such that $\|A - D - R\| \le \epsilon c_n$, where c_n depends on n and can be arbitrary.
- We show: This problem is NP-hard. So is Approximate P2 (much more complicated argument).

$\exists\mathbb{R}$ complete problems

- The canonical $\exists\mathbb{R}$ problem: Given a sequence of multivariate polynomial equations and inequalities with integer coefficients, does the system have a real root?
- A general decision problem is in $\exists\mathbb{R}$ if it can be transformed to a question about polynomial equations as in the first bullet.
- It is known: $NP \subseteq \exists\mathbb{R} \subseteq PSPACE$ (Canny).

Matrix completion is $\exists\mathbb{R}$-complete

- Shitov showed: matrix completion is $\exists\mathbb{R}$ complete.
- Specifically, Shitov showed that given a polynomial system, one can construct from it a partly specified symmetric matrix that has a semidefinite rank-3 completion iff the system has a real root.

Our reduction

- In order to use Shitov's construction for P2, we need to overcome the same issues as mentioned earlier:
  - In P2, unspecified entries are confined to the diagonal, and
  - In P2, all diagonal entries must be unspecified.
- We reuse the similar techniques as before, namely,
  - We use take Schur complement of two diagonal matrices to transform a matrix-completion problem with an arbitrary pattern of unspecified entries to one in which all unspecified entries are on the diagonal.
  - We make multiple copies of certain rows/columns diagonal elements must have certain values in any rank-3 completion.

P1 is not known to be $\exists\mathbb{R}$-complete

- P1 NP-hardness proof: Replace $*$'s on the diagonal in P2 gadget A with big numbers to get a P1 instance.
- However, for general integer matrices, the norm of D in the P2 solution can be double exponentially larger than the norm of A.
- This follows because the $\exists\mathbb{R}$-completeness of P2 shows that the system: $x_0 = 2$, $x_1 = x_0^2$, $x_2 = x_1^2, \ldots, x_n = x_{n-1}^2$ can be encoded as an $O(n)$-sized matrix completion problem even though the solution has $x_n = 2^{(2^n)}$.
- These big diagonal entries cannot be written down in polynomial time, so not clear how to transform a polynomial system to P1.