Algorithms for Rank and Tensor Structured Matrices

P. Benner and T. Mach

Max Planck Institute for Dynamics of Complex Technical Systems Magdeburg,
Research Group Computational Methods in Systems and Control Theory

MAX-PLANCK-GESELLSCHAFT

MAX–PLANCK–INSTITUT
FÜR DYNAMIK KOMPLEXER
TECHNISCHER SYSTEME
MAGDEBURG

Motivation

The standard algorithms for dense matrices become expensive for large matrices, since the number of floating point operations often grows like n^3. Therefore it is necessary to have data-sparse algorithms that use the problem-inherent structure to reduce the computational complexity. Data-sparse means that the matrix can be represented with much fewer than n^2 storage. This is for instance the case for rank or tensor structured matrices.
We have worked on new highly efficient algorithms for hierarchical matrices and matrices in tensor train matrix format. Here we will focus on one eigenvalue algorithm for symmetric matrices of both classes of data-sparse matrices.
Let $M \in \mathbb{C}^{n\times n}$. Then the pair (λ, v) is called eigenpair if it fulfills the equation:

$$Mv = \lambda v.$$

If M is real and symmetric, i.e., $M = M^T \in \mathbb{R}^{n\times n}$, then $\lambda \in \mathbb{R}$ and $v \in \mathbb{R}^n$.

Hierarchical Matrices

Hierarchical ($\mathcal{H}$-)matrices, see [2], are a data-sparse representation of dense matrices. With the help of a hierarchical structure the matrix is partitioned into submatrices. Most of these submatrices have good low rank approximations and are represented by those, but a few do not. These submatrices have to be stored as dense matrices. This leads to an efficient hierarchical rank structure:

- adaptive rank $k(\varepsilon)$
- storage $N_{St,\mathcal{H}}(T,k) = \mathcal{O}(n \log n\, k(\varepsilon))$
- approximate arithmetic operations

$M_{\mathcal{H}} v$	$\mathcal{O}(n \log n\, k(\varepsilon))$
$+_{\mathcal{H}}, -_{\mathcal{H}}$	$\mathcal{O}(n \log n\, k(\varepsilon)^2)$
$*_{\mathcal{H}}, \mathcal{H}LU(\cdot), (\cdot)^{-1}_{\mathcal{H}}, \mathcal{H}LDL^T$	$\mathcal{O}(n (\log n)^2\, k(\varepsilon)^2)$

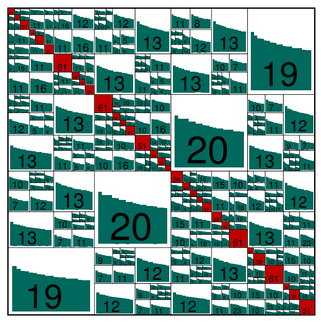

Figure 1: $\mathcal{H}$-matrix structure diagram.

FEM and boundary element matrices and their inverses can be approximated by hierarchical matrices. In [4] $\mathcal{H}$-matrices are used for modeling population dynamics.

Tensor-Train Matrices

Matrices in tensor-train matrix (ttm) format, see [6], can be represented by

$$M = \sum_{\alpha_1,\ldots,\alpha_d} M_1(i_1, j_1, \alpha_1) M_2(\alpha_1, i_2, j_2, \alpha_2) \cdots M_d(\alpha_{d-1}, i_d, j_d).$$

If we split the row index i of M into $(i_1, i_2, \ldots, i_d)$ and the column index j into $(j_1, j_2, \ldots, j_d)$. The coupling indices α_i run from 1 to $r_i < r$, the local ranks of the tensor train.

$$G_1(i_1, j_1, \alpha_1) - \alpha_1 - G_2(\alpha_1, i_2, j_2, \alpha_2) - \alpha_2 - \cdots - \alpha_{d-1} - G_d(\alpha_{d-1}, i_d, j_d)$$

The matrix $M \in \mathbb{R}^{n\times n}$, with $n = \ell^d$ can be stored in the ttm format with only $(d-2)\ell^2 r^2 + 2\,\ell^2 r$ entries, so that the storage complexity is logarithmic in n. Vectors $v \in \mathbb{R}^n$ can be stored in the related tensor-train format.

Preconditioned Inverse Iteration

The preconditioned inverse iteration (PINVIT), see [3], is an efficient method for the computation of the smallest eigenvalue(s) of symmetric, positive definite matrices. The method needs a number of steps independent of the matrix dimension n.
The smallest eigenpair is also the minimum of the Rayleigh quotient

$$\mu(x) := \mu(x, M) := \frac{x^T M x}{x^T x}.$$

PINVIT now minimizes this function by applying an inexact Newton method leading to the update equation

$$x_{i+1} := x_i - B^{-1}\left(Mx_i - x_i \mu(x_i)\right),$$

where B is a preconditioner for M fulfilling

$$\|\mathcal{I} - B^{-1}M\|_M \leq c < 1.$$

The convergence rate of the method depends on c, so c should be small, say 0.2. Since one has to compute the residual $r(x_i) := Mx_i - x_i\mu(x_i)$ in each step anyway, one can use the norm of the residual $\|r(x_i)\|_2$ as stopping criterion. If the residual is smaller than ϵ, then we often have computed the eigenvalue already to an accuracy of approximately ϵ^2.

Folded Spectrum Method

The computation of inner eigenvalues is usually realized by shifting the matrix M, so that the sought eigenvalue becomes the smallest magnitude eigenvalue of $M - \sigma I$. That is not possible here, since PINVIT requires M to be positive definite. So we have to apply the folded spectrum method, see [7]. We use $M_\sigma = (M - \sigma I)^2$ instead of M. The advantage is that M_σ is positive definite and has the same eigenvectors as M. We then compute the smallest eigenpairs of M_σ. Finally we get the sought inner eigenvalues by computing the Rayleigh quotient for the eigenvectors and M.

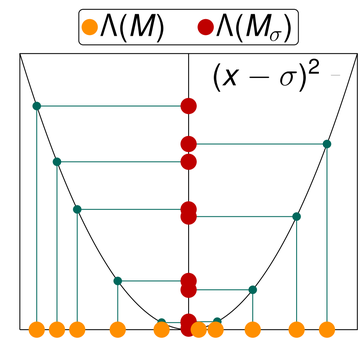

PINVIT for Data-Sparse Matrices

The preconditioned inverse iteration is ideal for data-sparse matrix formats, since it requires only

- the computation of a preconditioner (e.g., $X_k = X_{k-1}(2I - MX_{k-1})$, $X_k \to M^{-1}$) and
- matrix-vector products.

Both formats here provide these operations in an efficient way. In the future we will extend these examinations to other data-sparse matrix formats and to other algorithms.

Numerical Results

$\mathcal{H}$-Matrix 3D Laplace [1]

n	t_{chol} in s	t_{PINVIT} in s	$t_{\text{MATLAB eigs}}$ in s
512	0.07	0.03	0.05
4 096	1.92	0 45	0.36
32 768	75.83	7.40	28.43
262 144	2 597.82	194.89	1 932.03

Tensor-Train Matrix 4D Laplace [5]

n	t_{inv} in s	t_{PINVIT} in s	error
4 096	0.873	23.993	1.7104 e−07
65 536	1.436	119.348	5.9876 e−08
1 048 576	5.975	497.812	2.2100 e−08
16 777 216	12.655	1 710.326	7.9299 e−09
268 435 456	23.628	7 898.374	5.1963 e−10

Acknowledgments

Parts of this work were funded by a grant from the Free State of Saxony at Chemnitz University of Technology.

References

[1] P. Benner and T. Mach, *The preconditioned inverse iteration for hierarchical matrices*, Numer. Lin. Alg. Appl., (2012). 17 pages. Published online.

[2] W. Hackbusch, *Hierarchische Matrizen. Algorithmen und Analysis*, Springer-Verlag, Berlin, 2009.

[3] A. V. Knyazev, *Preconditioned eigensolvers — an oxymoron?*, Electr. Trans. Num. Anal., 7 (1998), pp. 104–123.

[4] J. Koch, W. Hackbusch, and K. Sundmacher, *$\mathcal{H}$-matrix methods for linear and quasi-linear integral operators appearing in population balances*, Comp. & Chem. Eng., 31 (2007), pp. 745–759.

[5] T. Mach, *Computing inner eigenvalues of matrices in tensor train matrix format*, in ENUMATH 2011 Proceedings Volume, 2011. Submitted, available as Max Planck Institute Magdeburg Preprint MPIMD/11-09, 11 pages.

[6] I. V. Oseledets and E. E. Tyrtyshnikov, *Breaking the curse of dimensionality or how to use SVD in many dimensions*, SIAM J. Sci. Comput., 31 (2009), pp. 3744–3759.

[7] L.-W. Wang and A. Zunger, *Electronic structure pseudopotential calculations of large (∼1000 atoms) Si quantum dots*, J. Phys. Chem., 98 (1994), pp. 2158–2165.