

# What makes block Jacobi slow to converge, and how low-rank changes can help

*List of authors:*

Felix Kwok<sup>1</sup>

For linear systems arising from the discretization of PDEs, one can often design naturally parallel preconditioners by decomposing the computational domain into subdomains and solving the local problems in parallel. While the basic block Jacobi/additive Schwarz preconditioner is an obvious first choice in many applications, convergence of the preconditioned system often becomes unacceptably slow as the grid is refined.

In this talk, we give two reasons that explain the slow convergence of the method. One reason is that the solution in one subdomain is too tightly coupled to its neighbours; thus, when errors are present in the neighbouring solutions, the method is incapable of reducing the induced error quickly. One can obtain much faster convergence by making low-rank changes to the preconditioner in order to incorporate Schur complement information between neighbouring subblocks.

The second reason for slow convergence is that the method does not allow global communication between far-away subdomains. For problems with many subdomains, a scalable method must be able to quickly reduce errors due to long-range effects; these errors typically reside in a low-dimensional subspace that can be represented by a coarse grid. Using a Schur complement-based approach, we derive an optimal method that converges in two iterations, regardless of the PDE, the geometry of the computational domain or the subdomain decomposition used. This method reveals the exact coarse-grid information that needs to be approximated in order to obtain a practical two-level preconditioner whose convergence rate is independent of the number of subdomains.

---

<sup>1</sup>Université de Genève