

# Inexact Gauss-Seidel Coarse Solvers for AMG and s-step CG

Stephen Thomas<sup>[0000-0001-8007-0057]</sup>,  
Pasqua D'Ambra<sup>[0000-0003-2047-4986]</sup>

## 1 Introduction

Sparse iterative solvers for large-scale partial differential equations encounter a fundamental exascale bottleneck due to the latency of global synchronizations. In Krylov methods such as Conjugate Gradient (CG), each iteration requires at least two global reductions to compute inner products; at processor counts reaching hundreds of thousands, synchronization costs can dominate the overall runtime, particularly under weak scaling. The  $s$ -step, or communication-avoiding, formulation of CG mitigates this limitation by generating  $s$  Krylov basis vectors per outer iteration, thereby reducing synchronization frequency by a factor of  $s$ . This reduction is especially critical for domain decomposition preconditioners, where the boundary-to-interior ratio deteriorates with increasing processor counts, amplifying synchronization overhead.

The  $s$ -step methodology originates with Chronopoulos and Gear [2, 1], who observed that several matrix-vector products could be performed prior to synchronization. Subsequent works [7, 3] developed this idea within the broader framework of communication-avoiding algorithms. A central challenge of the approach is that each outer iteration requires solving the dense Gram system:  $G\alpha = P^T r$ ,  $G = P^T A P \in \mathbb{R}^{s \times s}$ , for which standard methods, such as direct Cholesky factorization ( $O(s^3)$ ) or iterative schemes such as LSQR, become increasingly burdensome as  $s$  grows.

The principal contribution of this work is to demonstrate that Gram matrices arising from Chebyshev polynomial bases admit highly efficient solution via FGS. A moderate number of sweeps, typically  $\nu = 20$ – $30$ , suffices for convergence of

---

Stephen Thomas  
Lehigh University, 27 Memorial Drive West, Bethlehem, PA 18015 USA, e-mail:  
sjt223@lehigh.com

Pasqua D'Ambra  
Institute for Applied Computing of the National Research Council of Italy, Via P. Castellino, 111,  
80131, Naples, Italy, e-mail: pasqua.dambra@cnr.it

the outer iteration, reducing complexity from  $O(s^3)$  to  $O(\nu s^2)$  while improving numerical stability and enabling efficient hybrid CPU/GPU implementations.

Our analysis distinguishes *conditioning* from *decay*. Although the bound  $\kappa(G) = O(s^2)$  [8] governs worst-case behavior, the structural estimate  $\|L\|_F = O(\sqrt{s})$  explains the effectiveness of FGS: because  $r^{(1)} = -L^T \alpha^{(1)}$ , convergence depends on  $\|L\|$  rather than  $\kappa(G)$ , clarifying why  $\nu = 20\text{--}30$  sweeps suffice despite the conditioning barrier. We further establish the algebraic equivalence between FGS and MGS orthogonalization and extend the framework to AMG coarse-grid streaming.

The chapter is organized as follows. Section 2 introduces notations and preliminaries. Sections 3–5 analyze the FGS iteration, its stability, conditioning, and its equivalence with MGS. Section 6 extends the approach to AMG coarse-grid solves, and Section 7 studies inexact Gram solves. Numerical results and conclusions are given in Sections 8 and 9.

## 2 Preliminaries

We consider solving the SPD linear system  $Ax = b$ , where  $A \in \mathbb{R}^{n \times n}$  is sparse. The  $s$ -step CG method operates on  $\mathcal{K}_s(M^{-1}A, r^{(k)})$  with preconditioner  $M$  and residual  $r^{(k)}$ , forming a basis  $P \in \mathbb{R}^{n \times s}$ , solving the Gram system  $G\alpha = P^T r^{(k)}$ ,  $G = P^T A P$ , and updating  $x^{(k+1)} = x^{(k)} + P\alpha$ . For monomial bases,  $\kappa(G) = O(\kappa(M^{-1}A)^{s-1})$  [5], which typically limits  $s$  to  $s \leq 4$ .

**Definition 1 (Chebyshev Polynomial Basis)** Let  $[\lambda_{\min}, \lambda_{\max}]$  contain the spectrum of  $M^{-1}A$ . The basis is generated by the three-term recurrence

$$p_0 = r^{(k)}, \quad p_1 = \theta(M^{-1}A - \sigma I)p_0, \quad p_{j+1} = 2\theta(M^{-1}A - \sigma I)p_j - p_{j-1},$$

for  $j \geq 1$ , with  $\theta = 2/(\lambda_{\max} - \lambda_{\min})$  and  $\sigma = (\lambda_{\max} + \lambda_{\min})/2$  [7].

The spectral map sends  $[\lambda_{\min}, \lambda_{\max}]$  to  $[-1, 1]$ , where Chebyshev polynomials minimize the maximum deviation from zero, yielding near  $A$ -orthogonality.

**Column scaling.** Following Ruhe [10], set  $d_i = (p_i^T A p_i)^{-1/2}$ , and let  $S = \text{diag}(d_i)$  so that  $\tilde{P} = PS$  satisfies  $\tilde{p}_i^T A \tilde{p}_i = 1$ . Redefining the Gram matrix as  $G = \tilde{P}^T A \tilde{P}$ , we then have  $\text{diag}(G) = I$ . Writing  $G = I + L + L^T$ , where  $L$  is strictly lower triangular, the FGS iteration  $(I + L)\alpha^{(\ell+1)} = \tilde{P}^T r^{(k)} - L^T \alpha^{(\ell)}$ ,  $\alpha^{(0)} = 0$ , requires  $O(s^2)$  work per sweep.

## 3 FGS Iteration Analysis

The following results characterize a single FGS sweep. Although some statements are given for the first iteration with  $\alpha^{(0)} = 0$ , they extend to later sweeps by replacing  $P^T r$  with the current right-hand side.

**Proposition 1 (FGS Residual Structure)** *Let  $\alpha^{(v+1)}$  satisfy  $(D+L)\alpha^{(v+1)} = P^T r - L^T \alpha^{(v)}$ . Then  $r^{(v+1)} = P^T r - G\alpha^{(v+1)} = -L^T \alpha^{(v+1)} + L^T \alpha^{(v)}$ , and for  $\alpha^{(0)} = 0$ ,  $r^{(1)} = -L^T \alpha^{(1)}$ .*

*Proof.* Because  $G = D + L + L^T$ ,

$$G\alpha^{(v+1)} = (D + L)\alpha^{(v+1)} + L^T \alpha^{(v+1)} = (P^T r - L^T \alpha^{(v)}) + L^T \alpha^{(v+1)},$$

and subtracting from  $P^T r$  gives the claim. □

This structure shows that FGS convergence depends primarily on the strictly upper-triangular part  $L^T$ , and by Theorem 2,  $\|L\|_F = O(\sqrt{s})$  for Chebyshev bases.

**Theorem 1 (Backward Stability)** *For any FGS sweep, the computed  $\tilde{\alpha}^{(v+1)}$  satisfies*

$$(D + L)\tilde{\alpha}^{(v+1)} = P^T r - L^T \alpha^{(v)} + \delta r,$$

with

$$\|\delta r\| \leq \epsilon_{mach} C(s) \left( \|P^T r\| + \|D + L\| \|\tilde{\alpha}^{(v+1)}\| \right), \quad C(s) = O(s)[\sigma].$$

*Proof.* Forward substitution for  $(D + L)\alpha = b$  computes

$$\alpha_j = b_j - \sum_{i < j} L_{ji} \alpha_i.$$

Floating-point evaluation yields  $\tilde{\alpha}_j = b_j - \sum_{i < j} L_{ji} \tilde{\alpha}_i + \epsilon_j$  with  $|\epsilon_j| \leq (j - 1)\epsilon_{mach}(|b_j| + \sum_{i < j} |L_{ji}| |\tilde{\alpha}_i|)$ ; summing over  $j$  and using norm inequalities gives the stated bound. □

For  $s \leq 20$  in double precision,  $C(s)\epsilon_{mach} \lesssim 10^{-14}$ , negligible relative to algorithmic error.

**Corollary 1 (Error Bound)** *Let  $\delta_1 = \|r^{(1)}\|/\|P^T r\|$ . Then  $\delta_1 \lesssim \|L\| \kappa(G)$ , where  $\kappa(G) = \|G\| \|G^{-1}\|$  and the hidden constant reflects the closeness of  $G$  and  $D + L$ .*

**Lemma 1 (Multi-Sweep Analysis)** *Let  $T = -(D + L)^{-1} L^T$  be the iteration matrix and set  $\rho = \|T\|_2$ . After  $v$  sweeps,  $\|r^{(v)}\| \leq \rho^v \|P^T r\|$ , and geometric convergence holds when  $\rho < 1$ .*

*Proof.* The iteration  $\alpha^{(k+1)} = (D+L)^{-1}(P^T r - L^T \alpha^{(k)})$  gives  $e^{(k+1)} = \alpha^* - \alpha^{(k+1)} = T e^{(k)}$ , where  $\alpha^*$  is the exact solution. Thus  $\|e^{(v)}\| \leq \|T\|_2^v \|e^{(0)}\| = \rho^v \|\alpha^*\|$ . Because  $r^{(v)} = G(\alpha^* - \alpha^{(v)})$  and  $\|G\| \leq 1 + 2\|L\|$ , geometric convergence follows for well-conditioned  $G$  with  $\|L\| = O(\sqrt{s})$ . □

## 4 Chebyshev Gram Matrix Conditioning

The effectiveness of FGS as an inner solver depends on the spectral properties of the Gram matrix  $G$ . Classical analysis for monomial bases [5] predicts rapid growth of  $\kappa(G)$  with degree, but Chebyshev bases behave differently due to their near  $A$ -orthogonality [8]. This section shows how Chebyshev polynomials, combined with column scaling, induce off-diagonal decay in  $G$ , yielding  $\|L\|_F = O(\sqrt{s})$  and polynomial rather than exponential growth of  $\kappa(G)$ .

**Theorem 2 (Polynomial Growth)** *For the Chebyshev basis with column scaling,  $\kappa(G) \leq Cs^2$  with  $C$  weakly dependent on  $\kappa(M^{-1}A)$ , and  $\|L\|_F = O(\sqrt{s})$ .*

*Proof.* Chebyshev orthogonality on  $[-1, 1]$ ,

$$\int_{-1}^1 T_i(t)T_j(t) \frac{dt}{\sqrt{1-t^2}} = \begin{cases} 0 & i \neq j, \\ \pi/2 & i = j > 0, \end{cases}$$

transfers to the  $A$ -norm via the spectral decomposition of  $M^{-1}A$ . Let  $M^{-1}A = \sum_k \lambda_k v_k v_k^T$  and  $r_0 = \sum_k \beta_k v_k$ . Then

$$\hat{G}_{ij} = \sum_k \beta_k^2 \lambda_k T_i(\hat{\lambda}_k) T_j(\hat{\lambda}_k), \quad \hat{\lambda}_k = \theta(\lambda_k - \sigma),$$

and the approximate orthogonality of Chebyshev polynomials implies an off-diagonal decay  $|\hat{G}_{ij}| \lesssim 1/|i-j|$  for  $i \neq j$ . Column scaling by a diagonal matrix with bounded condition number preserves this decay up to constants, so  $G_{ij} = O(1/|i-j|)$ . Hence

$$\|L\|_F^2 \lesssim \sum_{i=2}^s \sum_{k=1}^{i-1} k^{-2} \leq s \pi^2/6 = O(s),$$

giving  $\|L\|_F = O(\sqrt{s})$ . Gershgorin bounds with  $|G_{ij}| \sim 1/|i-j|$  imply  $\lambda_{\max}(G) \lesssim 1 + \log s$ , so that  $\kappa(G)$  grows at most polynomially in  $s$ , and under mild assumptions on  $\lambda_{\min}(G)$  this yields  $\kappa(G) = O(s^2)$ .  $\square$

**Monomial comparison.** Monomial bases yield exponential growth  $\kappa(G^{\text{mono}}) = O(\kappa(M^{-1}A)^{s-1})$ , exceeding machine precision by  $s = 8$  for  $\kappa(M^{-1}A) = 100$ . By contrast, Chebyshev gives  $\kappa(G) \sim 100$  at  $s = 10$ , making FGS practical.

## 5 MGS–FGS Equivalence

A single FGS sweep corresponds algebraically to one step of MGS applied in the  $A$ -norm. This interpretation provides a compact view of FGS and helps explain its numerical behavior within the  $s$ -step framework.

**Theorem 3 (Algebraic Equivalence)** *One FGS sweep on  $G\alpha = \tilde{P}^T Ar$ , with  $G = \tilde{P}^T A \tilde{P}$  and  $\alpha^{(0)} = 0$ , is algebraically equivalent to one step of MGS orthogonalization of  $r$  against the columns of  $\tilde{P}$  in the  $A$ -norm.*

*Proof.* Consider MGS in the  $A$ -inner product with respect to the columns  $\{\tilde{p}_1, \dots, \tilde{p}_s\}$  of  $\tilde{P}$ . Initialize  $w_0 = r$  and, for  $j = 1, \dots, s$ , define  $\gamma_j = w_{j-1}^T A \tilde{p}_j$ ,  $w_j = w_{j-1} - \gamma_j \tilde{p}_j$ . FGS applied to  $G\alpha = \tilde{P}^T Ar$  with  $G = \tilde{P}^T A \tilde{P}$  and  $\alpha^{(0)} = 0$  produces, componentwise,

$$\alpha_j^{(1)} = \tilde{p}_j^T Ar - \sum_{i=1}^{j-1} (\tilde{p}_j^T A \tilde{p}_i) \alpha_i^{(1)}, \quad j = 1, \dots, s.$$

We prove by induction on  $j$  that  $\alpha_j^{(1)} = \gamma_j$  and  $w_{j-1} = r - \sum_{i=1}^{j-1} \alpha_i^{(1)} \tilde{p}_i$ .

For  $j = 1$ , we have  $\alpha_1^{(1)} = \tilde{p}_1^T Ar$ ,  $\gamma_1 = w_0^T A \tilde{p}_1 = r^T A \tilde{p}_1 = \tilde{p}_1^T Ar$ , so  $\alpha_1^{(1)} = \gamma_1$  and  $w_1 = r - \gamma_1 \tilde{p}_1$ . Assume the claim holds for all  $i < j$ , so that  $w_{j-1} = r - \sum_{i=1}^{j-1} \alpha_i^{(1)} \tilde{p}_i$ . Then

$$\begin{aligned} \gamma_j &= w_{j-1}^T A \tilde{p}_j = \left( r - \sum_{i=1}^{j-1} \alpha_i^{(1)} \tilde{p}_i \right)^T A \tilde{p}_j \\ &= \tilde{p}_j^T Ar - \sum_{i=1}^{j-1} \alpha_i^{(1)} (\tilde{p}_i^T A \tilde{p}_j) = \alpha_j^{(1)}, \end{aligned}$$

where we used the symmetry of  $A$  and the definition of the FGS update. Thus  $\gamma_j = \alpha_j^{(1)}$  for all  $j$ , and MGS and FGS compute the same projection coefficients. Hence one FGS sweep is algebraically equivalent to one MGS step in the  $A$ -norm.  $\square$

## 6 AMG Coarse-Grid Extension

In AMG, the coarse-grid operator  $A_c = P^T A_f P$  often becomes dense at the bottom levels of the hierarchy.

### Proposition 2 (AMG Convergence)

*Assume that the AMG prolongation operator  $P$  satisfies the weak approximation property [9]. Then the coarse-grid operator  $A_c = P^T A_f P$  is spectrally equivalent to the restriction of  $A_f$  to the coarse space, and its conditioning does not deteriorate with the number of fine-grid unknowns  $n_f$ . In particular,  $\kappa(A_c)$  remains bounded with respect to  $n_f$  for a fixed coarse level. Thus FGS converges uniformly with  $\rho_{FGS} < 1$ , requiring  $\nu = 10$ -30 iterations.*

Proposition 2 makes FGS an attractive alternative to direct solvers in the lower levels of AMG, where computational and memory costs are otherwise dominant. When needed, the method can also be implemented in a streaming (matrix-free)

fashion to avoid forming  $A_c$  explicitly, further reducing memory requirements on modern GPU architectures.

## 7 Inexact s-Step CG Convergence

In the  $s$ -step CG method, the accuracy with which the small Gram systems  $G_k \alpha_k = P_k^T r_k$  are solved at each iteration  $k$  plays a crucial role in determining the overall convergence. Because FGS performs only a fixed number of sweeps, each Gram solve is inherently inexact, and the resulting perturbations may accumulate across outer iterations. Here we provide conditions under which the method remains stable and convergent. The analysis follows the framework of inexact Krylov methods and yields practical bounds on the admissible Gram solve error.

**Theorem 4 (Inexact Convergence)** *Consider  $s$ -step CG with Gram solves satisfying*

$$\|P_k^T r_k - G_k \alpha_k\| \leq \delta_k \|P_k^T r_k\|$$

*at each outer iteration  $k$ . Convergence is ensured if  $\sum_{k=1}^{N_{outer}} \delta_k \|A\| \|P_k\| \lesssim \epsilon_{tot}$ .*

*Proof sketch.* Following [4], the inexact Gram solve induces a perturbation  $e_k^{Gram}$  satisfying

$$\|e_k^{Gram}\| \lesssim \kappa(G_k) \lambda_{\min}(G_k)^{-1} \delta_k \|P_k^T r_k\|.$$

The resulting error in the update obeys  $\|P_k e_k^{Gram}\|_A \lesssim \delta_k \|A\| \|P_k\| \|r_k\|$ . Accumulating these contributions across iterations and using the geometric decay of  $\|r_k\|$  in exact CG yields the stated condition.  $\square$

CG's self-correcting property typically allows much larger inexactness in practice (e.g.,  $\delta \sim 10^{-4}$ ) than predicted by conservative theory (e.g.,  $\delta \sim 10^{-8}$ ).

**Proposition 3 (FGS Rate)** *For SPD  $G_k$ , the FGS iteration for  $G_k \alpha_k = P_k^T r_k$  converges linearly. Let  $T_{FGS}$  denote the corresponding iteration matrix. Then  $\rho_{FGS} = \rho(T_{FGS}) < 1$ , and standard condition-number bounds give  $\rho_{FGS} \lesssim 1 - c/\kappa(G_k)$  for some constant  $c > 0$ .*

For Chebyshev bases with  $\kappa(G_k) = O(s^2)$ , such bounds may suggest large  $\nu$ . However, experiments indicate that  $\nu = 20$ – $30$  sweeps are sufficient in practice, owing to the conservativeness of the bounds and the self-correcting behavior of CG.

## 8 Numerical Experiments

This section presents a set of very preliminar numerical experiments designed to assess the practical performance of the proposed FGS-based Gram solvers within

the s-step CG and AMG frameworks. We consider 3D Poisson problems  $-\Delta u = f$  on the unit cube  $\Omega = [0, 1]^3$  with homogeneous Dirichlet BCs, discretized using a 27-point finite difference stencil yielding  $\sim 12.2\text{M}$  degrees of freedom (DOFs) per GPU, on AMD MI250X GPUs (64 GB High-Bandwidth-Memory 2e, 1.6 TB/s bandwidth per die) using MPI and HIP. We evaluate both algorithmic behavior, such as the impact of the number of FGS sweeps  $\nu$  and the block size  $s$ , and the resulting weak scaling of the outer iterations. Chebyshev basis with adaptive parameters estimated via Lanczos (10 iterations compute extreme eigenvalues  $\lambda_{\min}, \lambda_{\max}$  with 10% safety margin) are used. The outer CG iteration is stopped when the relative residual satisfies  $\|r_k\|/\|r_0\| < 10^{-6}$ . The Gram systems are solved using  $\nu \in \{6, 15, 30\}$  FGS sweeps, allowing us to assess how different levels of inner accuracy affect overall convergence.

**Table 1** Weak scaling: outer iterations vs GPUs and  $\nu$  for two values of  $s$ .

$s = 10$							
GPUs	1	2	4	8	16	32	64
DOFs (M)	12.2	24.3	48.7	97.3	194.7	389.3	778.7
$\nu = 6$	–	–	32	34	36	62	136
$\nu = 15$	10	15	14	16	16	32	54
$\nu = 30$	8	13	12	15	17	19	33
$s = 20$							
$\nu = 6$	–	–	21	29	34	51	96
$\nu = 15$	9	12	15	16	16	22	24
$\nu = 30$	8	9	9	12	13	17	17

**Key observations.** With  $\nu = 6$ , iterations explode (136 at 64 GPUs for  $s = 10$ ). With  $\nu = 30$ , weak scaling is restored:  $s = 20$  yields  $2.1\times$  growth (8 to 17 iterations) across  $64\times$  processor increase. The 32-to-64 GPU jump reflects reduced AMG preconditioner effectiveness at small subdomain sizes. Measured  $\kappa(G)$  confirms  $O(s^2)$  scaling: 78 for  $s = 10$ , 310 for  $s = 20$ .

**AMG and performance.** We also evaluated a streaming variant of the coarse-grid solve, where  $A_c$  is never formed explicitly and FGS is applied using on-the-fly products  $A_f p_j$ . With  $\nu_{\text{coarse}} = 20$  at the coarsest level ( $n_c \approx 5000$ ), the streaming approach closely matches the direct solve, requiring 18 versus 17 V-cycle iterations ( $< 6\%$  difference). It also provides substantial memory savings: only 200 MB are needed instead of 500 MB for storing the  $A_c$ . Eliminating CPU–GPU transfers and host-side factorization reduces V-cycle time by roughly 12%.

## 9 Conclusions and Future Work

We have shown that FGS is an effective inner solver for Gram systems in communication-avoiding Krylov methods. Its equivalence to MGS in the  $A$ -norm ensures accurate projection coefficients, and our inexact  $s$ -step CG analysis shows that convergence is maintained with only 20–30 sweeps. The same ideas extend to AMG, where avoiding coarse-matrix formation reduces memory and runtime without sacrificing accuracy. Replacing the  $O(s^3)$  Gram factorization with an  $O(vs^2)$  FGS iteration enables larger block sizes, improves scalability, and lowers inner-solve cost, as demonstrated on 64 GPUs with modest memory usage. Ongoing work evaluates this approach at large scale and on leadership-class systems.

**Acknowledgements** This work has been partially supported by the “Energy Oriented Center of Excellence (EoCoE III): Fostering the European Energy Transition with Exascale” EuroHPC Project N. 101144014, funded by European Commission (EC).

## References

1. Chronopoulos, A.T., Gear, C.W.: On the efficient implementation of preconditioned  $s$ -step conjugate gradient methods on multiprocessors with memory hierarchy. *Parallel Comput.* **11**, 37–53 (1989)
2. Chronopoulos, A.T., Gear, C.W.:  $s$ -step iterative methods for symmetric linear systems. *J. Comput. Appl. Math.* **25**(2), 153–168 (1989)
3. Demmel, J., Hoemmen, M., Mohiyuddin, M., Yelick, K.: Avoiding communication in computing Krylov subspaces. *SIAM J. Sci. Comput.* **35**(5), A1520–A1548 (2013)
4. van den Eshof, J., Sleijpen, G.L.G.: Inexact Krylov subspace methods. *SIAM J. Matrix Anal. Appl.* **26**(1), 125–153 (2004)
5. Gautschi, W.: The condition of polynomials in power form. *Math. Comput.* **33**, 343–352 (1979)
6. Higham, N.J.: *Accuracy and Stability of Numerical Algorithms*, 2 edn. SIAM (2002)
7. Hoemmen, M.: *Communication-avoiding Krylov subspace methods*. Ph.D. thesis, University of California, Berkeley (2010)
8. Philippe, B., Reichel, L.: On the generation of Krylov subspace bases. *Appl. Numer. Math.* **62**, 1171–1186 (2012)
9. Ruge, J.W., Stüben, K.: Algebraic multigrid. In: *Multigrid Methods*, pp. 73–130. SIAM (1987)
10. Ruhe, A.: Numerical aspects of Gram-Schmidt orthogonalization. *Linear Algebra Its Appl.* **52**, 591–601 (1983)