

A New Exact State Reconstruction Strategy for Conjugate Gradient Methods with Arbitrary Preconditioners



universität
wien

Faculty of
Computer Science

Viktoria Mayer¹, Wilfried N. Gansterer²

¹ University of Vienna, Doctoral School Computer Science DoCS, Faculty of Computer Science, Vienna, Austria

² University of Vienna, Faculty of Computer Science, Vienna, Austria

{viktoria.mayer, wilfried.gansterer}@univie.ac.at

Abstract

We present **resilience extensions** to tolerate node failures in the Preconditioned Conjugate Gradient (PCG) method and in **communication-hiding** pipelined PCG variants

- Extension of an existing ESR strategy
- Can be used with arbitrary preconditioners
- Very small communication overhead due to exploiting algorithm-specific properties
- We consider communication-hiding variants, which overlap global communication with other operations
 - Pipelined PCG (PPCG)
 - Stable l -length Pipelined $p(l)$ -PCG

Resilience extensions

Existing ESR strategies

- Exploit data redundancy of MV products
- PCG: store search direction $\mathbf{p}^{(i)}$ redundantly while computing $\mathbf{A}\mathbf{p}^{(i)}$ in iteration i
- Recovery is not possible with implicit preconditioners (e.g., Multigrid)

New approach

- Introduce redundant local vector operations
- PCG: redundantly compute $\mathbf{x}^{(i+1)}$ in iteration i with redundant copies of $\mathbf{p}^{(i)}$ using

$$\mathbf{x}^{(i+1)} = \mathbf{x}^{(i)} + \alpha^{(i)} \mathbf{p}^{(i)}$$

- PPCG, $p(l)$ -PCG: additional redundant local vector operations (negligible on large-scale parallel computers)
- Overall, the same communication overhead as existing ESR strategies

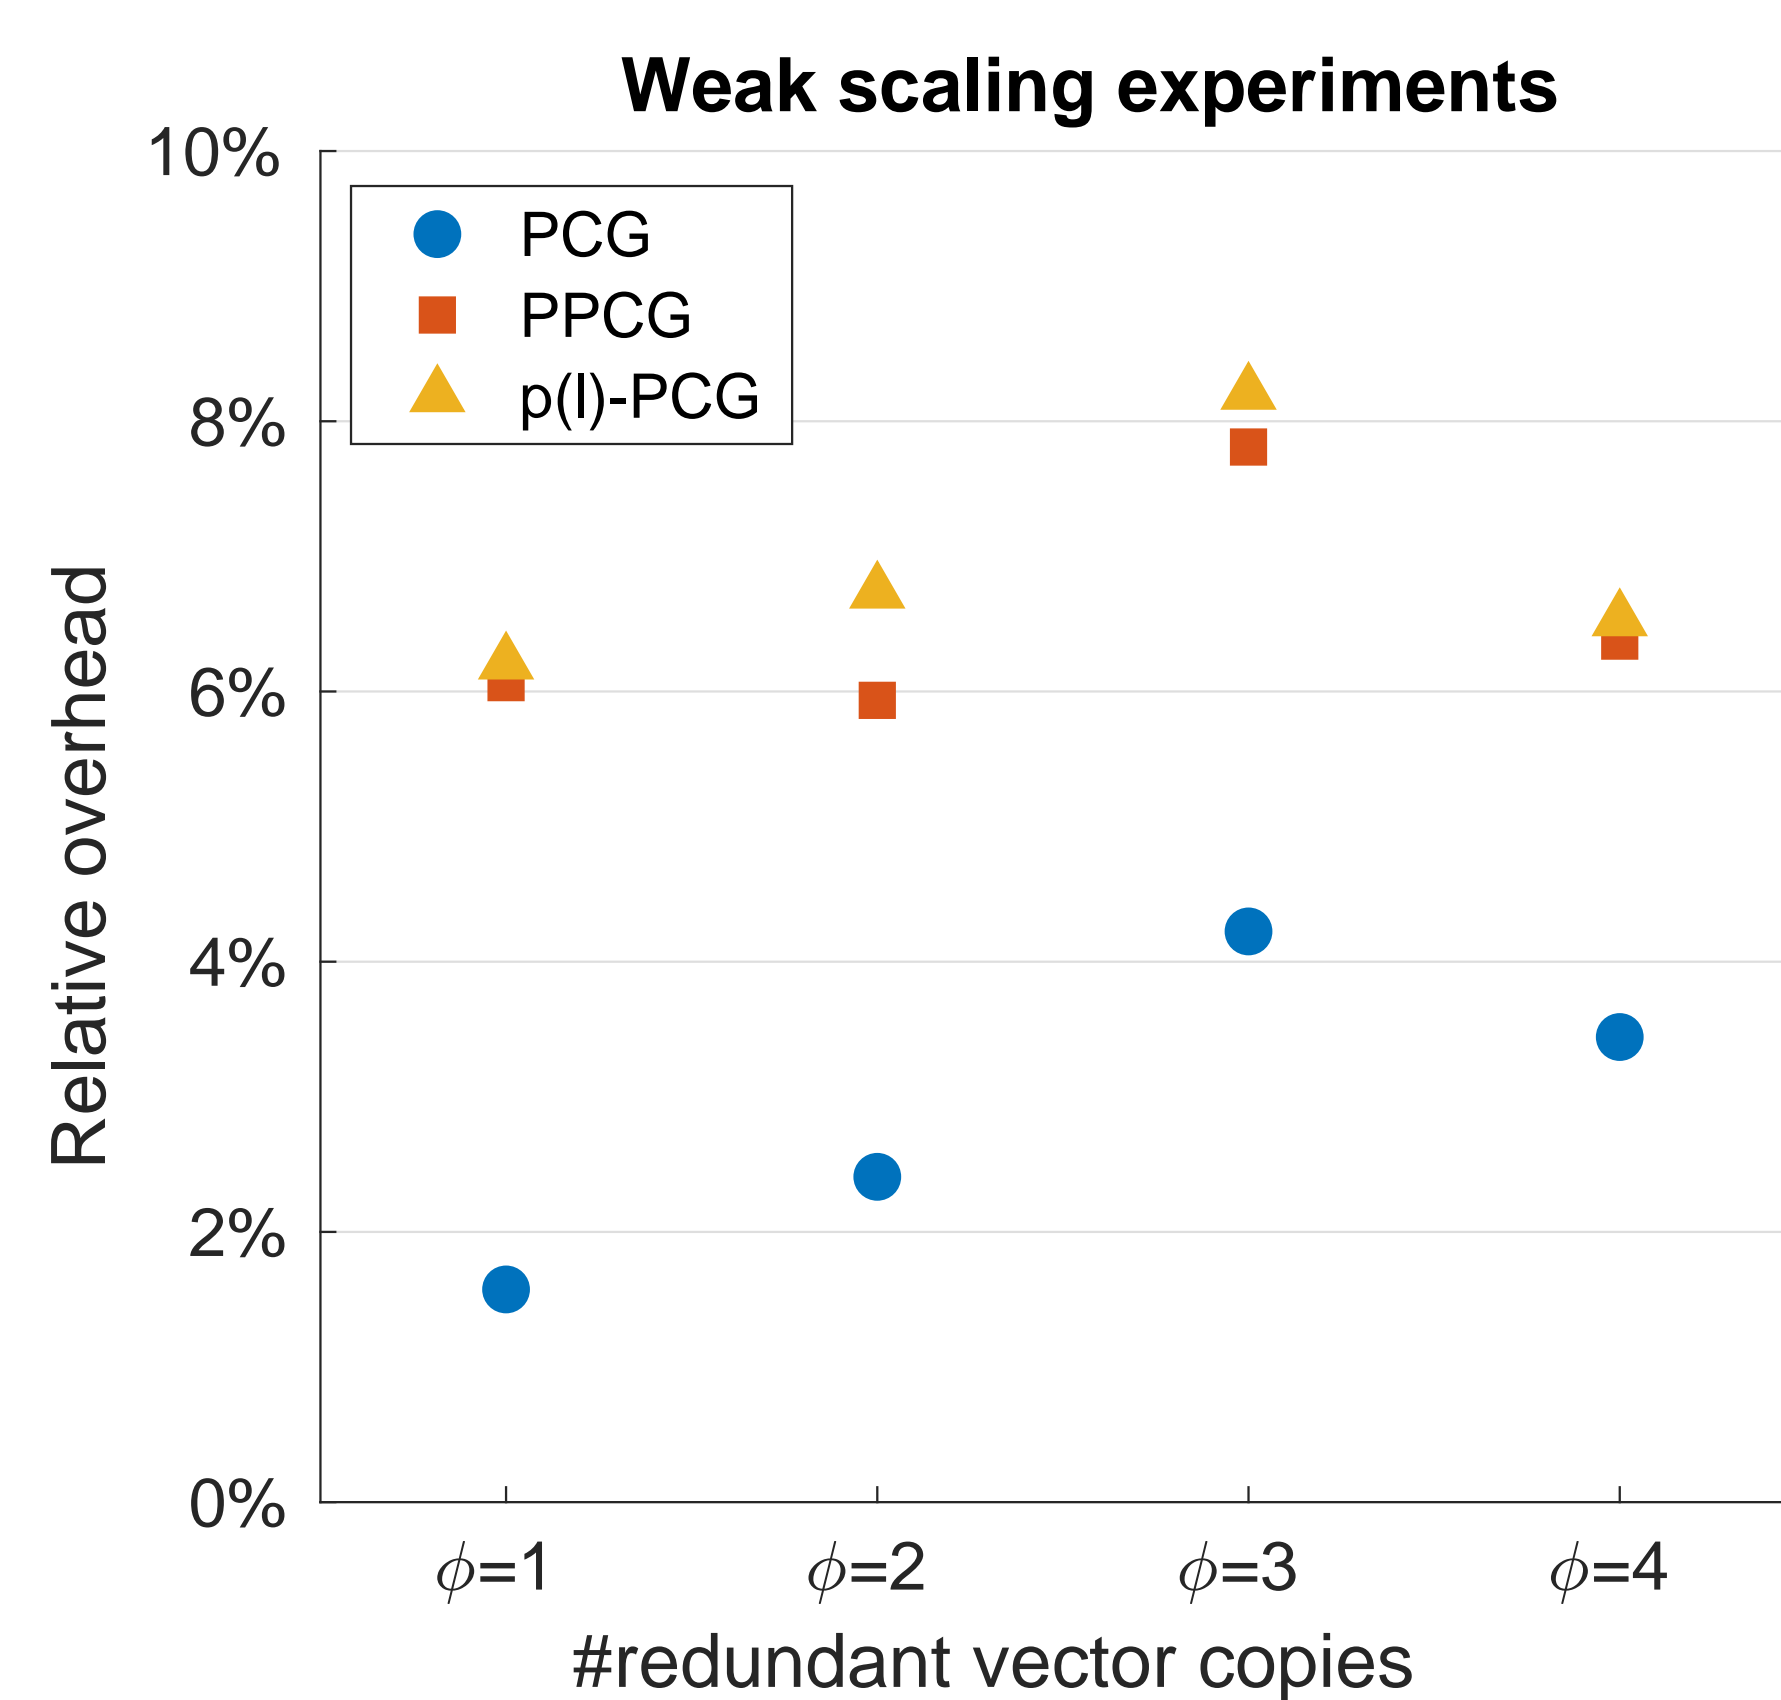
Recovery after node failures

- Reconstruction of lost data based on redundantly stored data
- Possible with any preconditioner
- Only involves operations which also occur in a normal solver iteration
- At most one preconditioner application

Contributions

- We developed PCG methods suitable for large-scale parallel computers in terms of scalability and cost-efficient tolerance of node failures
- Our algorithms can be used with any preconditioner, while existing ESR assumes that the preconditioner matrix is available explicitly
- Numerical experiments on the Vienna Scientific Cluster (VSC-5) illustrate very low runtime overhead for fault tolerance

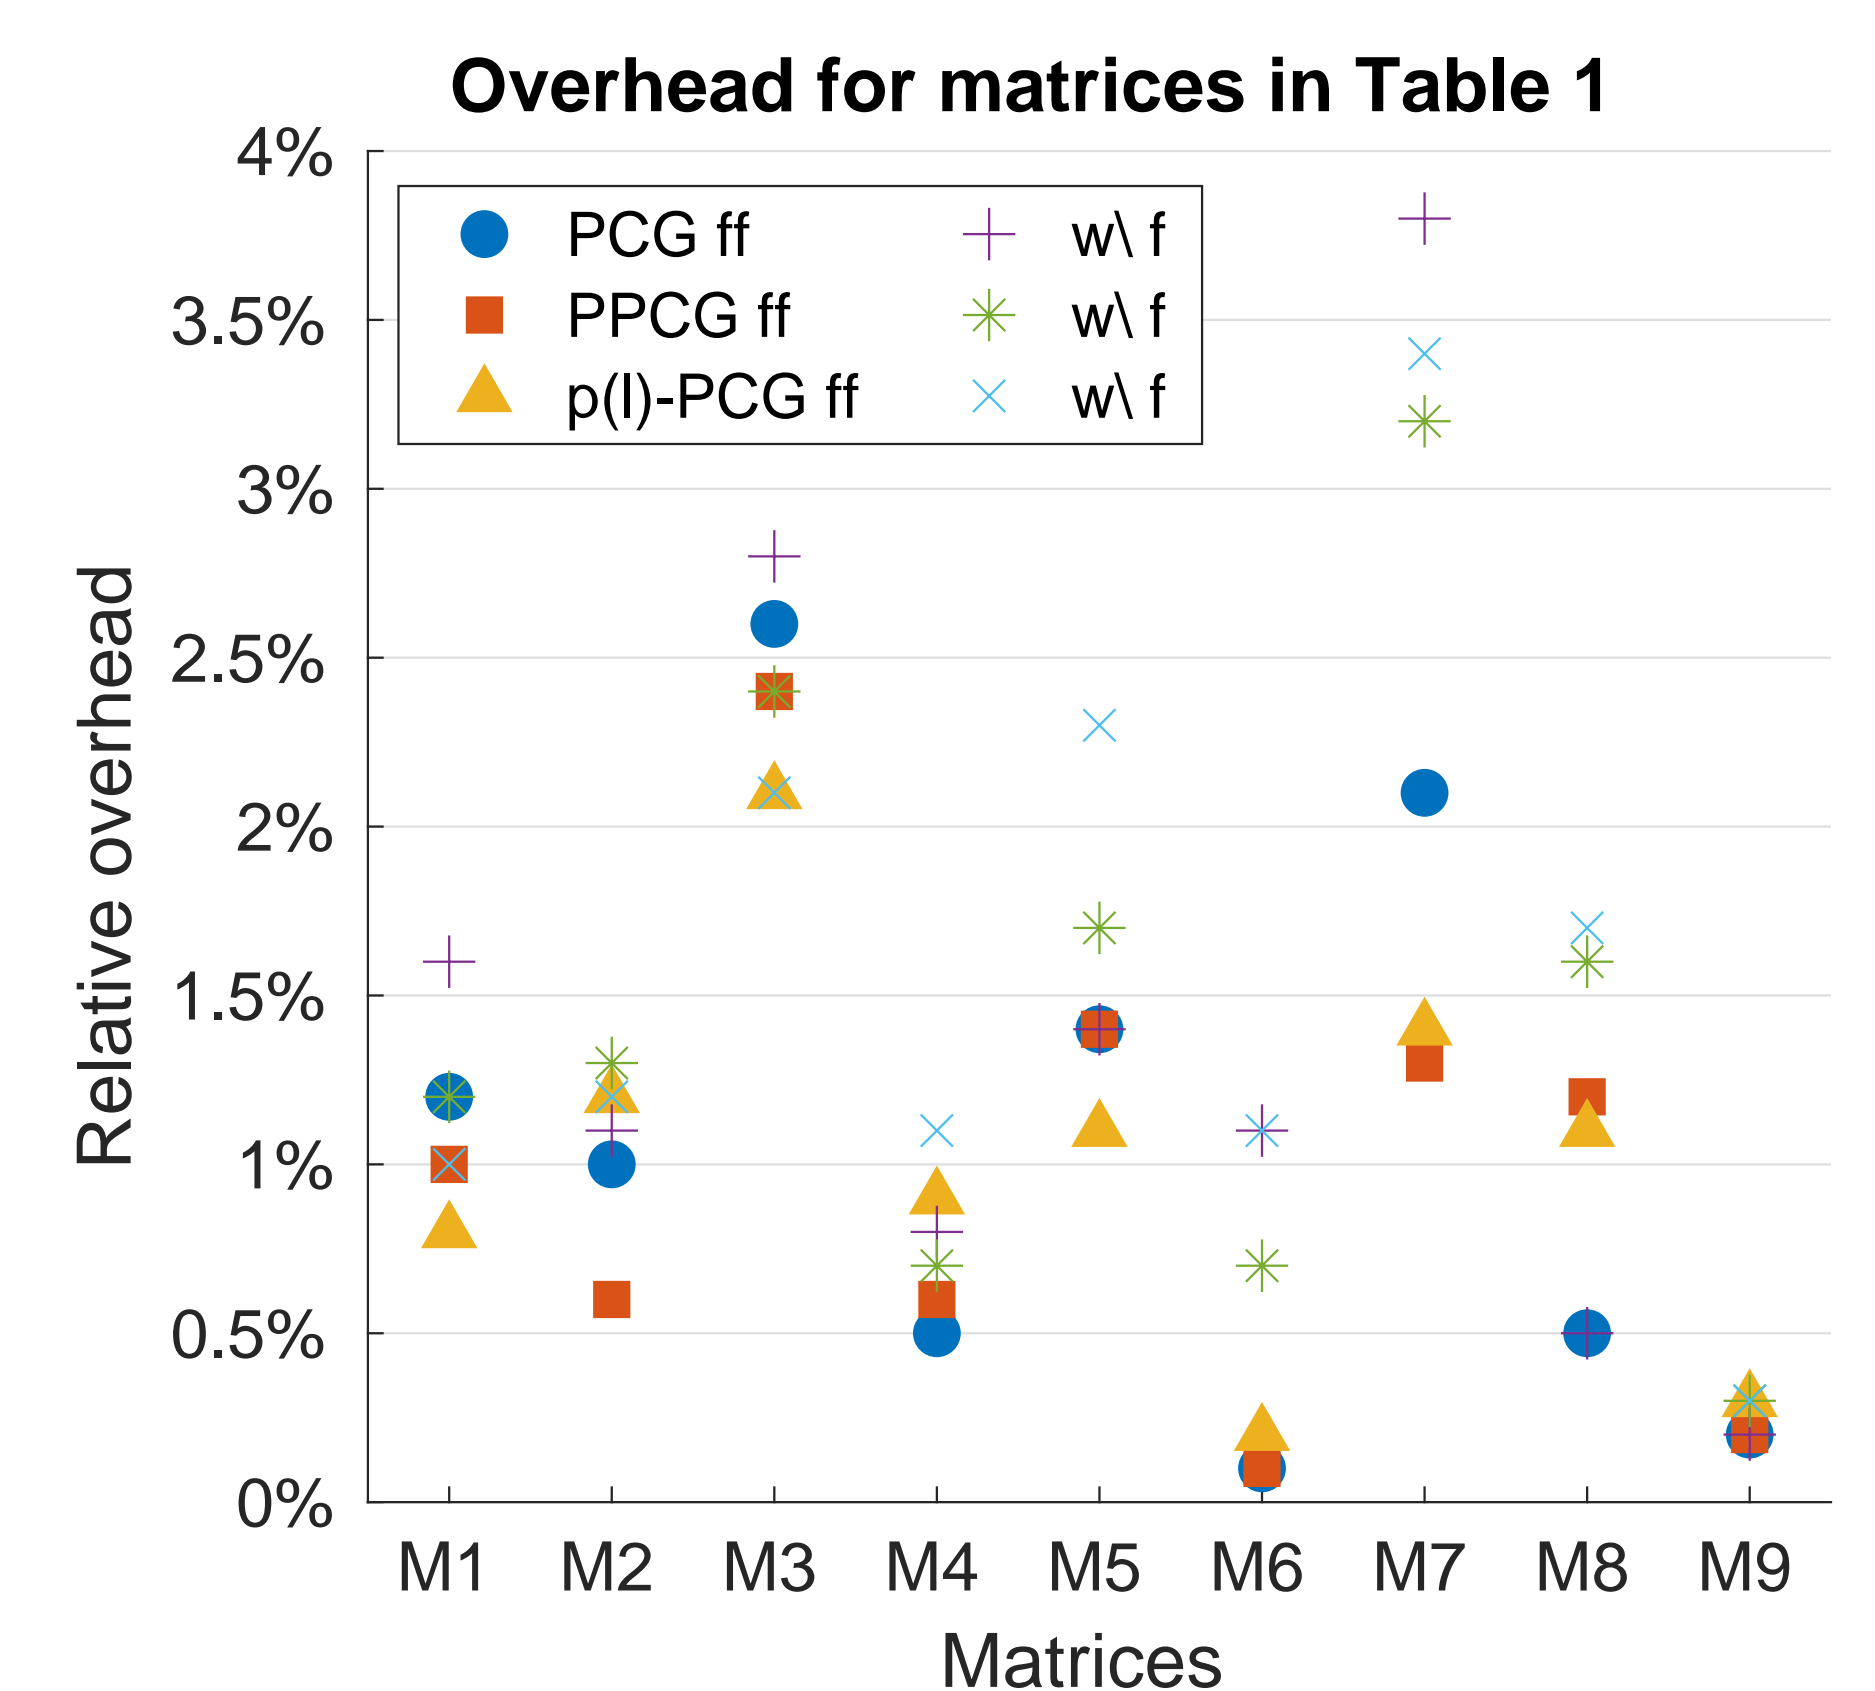
Experiments



Median runtime overheads over 20 test runs. ϕ : Number of redundant copies per redundantly stored vector. Pipeline length $l = 1$ for $p(l)$ -PCG.

Table 1: Test matrices from the SuiteSparse Matrix Collection. NNZ: Number of non-zeros.

Matrix	Id	Size	NNZ
af_shell3	M1	504 855	18e6
parabolic_fem	M2	525 825	4e6
apache2	M3	715 176	5e6
Emilia_923	M4	923 136	40e6
audikw_1	M5	943 695	78e6
ldoor	M6	952 203	42e6
Geo_1438	M7	1 437 960	60e6
StocF-1465	M8	1 465 137	21e6
Hook_1498	M9	1 498 023	59e6



Left figure: 3D Poisson matrices (27-point stencil) of sizes $128 \times 128 \times (32\phi)$, executed on 1024ϕ processes. Runtime overheads for resilience with ϕ simulated simultaneous process failures (after $\sim 50\%$ of solver iterations). Geometric Multigrid preconditioner with three levels and a Block-Jacobi-preconditioned Chebyshev iteration as both the bottom solver (10 iterations) and the smoother (1 iteration), with blocksize 10.

Right figure: Runtime overheads for matrices in Table 1, executed on 1024 processes. $\phi = 3$. ff: failure-free execution. w\ f: execution with 3 simulated simultaneous process failures (after $\sim 50\%$ of solver iterations). Preconditioner: Block-Jacobi-preconditioned Chebyshev iteration, 10 iterations per solver iteration.

Conclusion

- Novel ESR approach can efficiently be used with arbitrary preconditioners: same communication overhead as existing ESR, at most one preconditioner application during recovery
- Very small resilience overheads for nine real-world test cases (below 4%).
- Weak scaling experiments showed that the overheads vary only slightly with the problem size or with the number of simultaneous failures

Link

eprints.cs.univie.ac.at/8088/

