Iterative Solver Selection Techniques for Sparse Linear Systems

INTERNATIONAL CONFERENCE ON PARALLEL PROCESSING

UNIVERSITY OF OREGON

Kanika Sood, Boyana Norris, Elizabeth Jessup
University of Oregon, University of Colorado Boulder

University of Colorado Boulder

PROBLEM STATEMENT

High-performance numerical frameworks rely on mature and highly optimized numerical packages (e.g.: PETSc [1]) for the numerical solution of nonlinear partial differential equations. Because of the large and growing number of valid solution choices, the selection of particular solver configuration is becoming increasingly challenging.

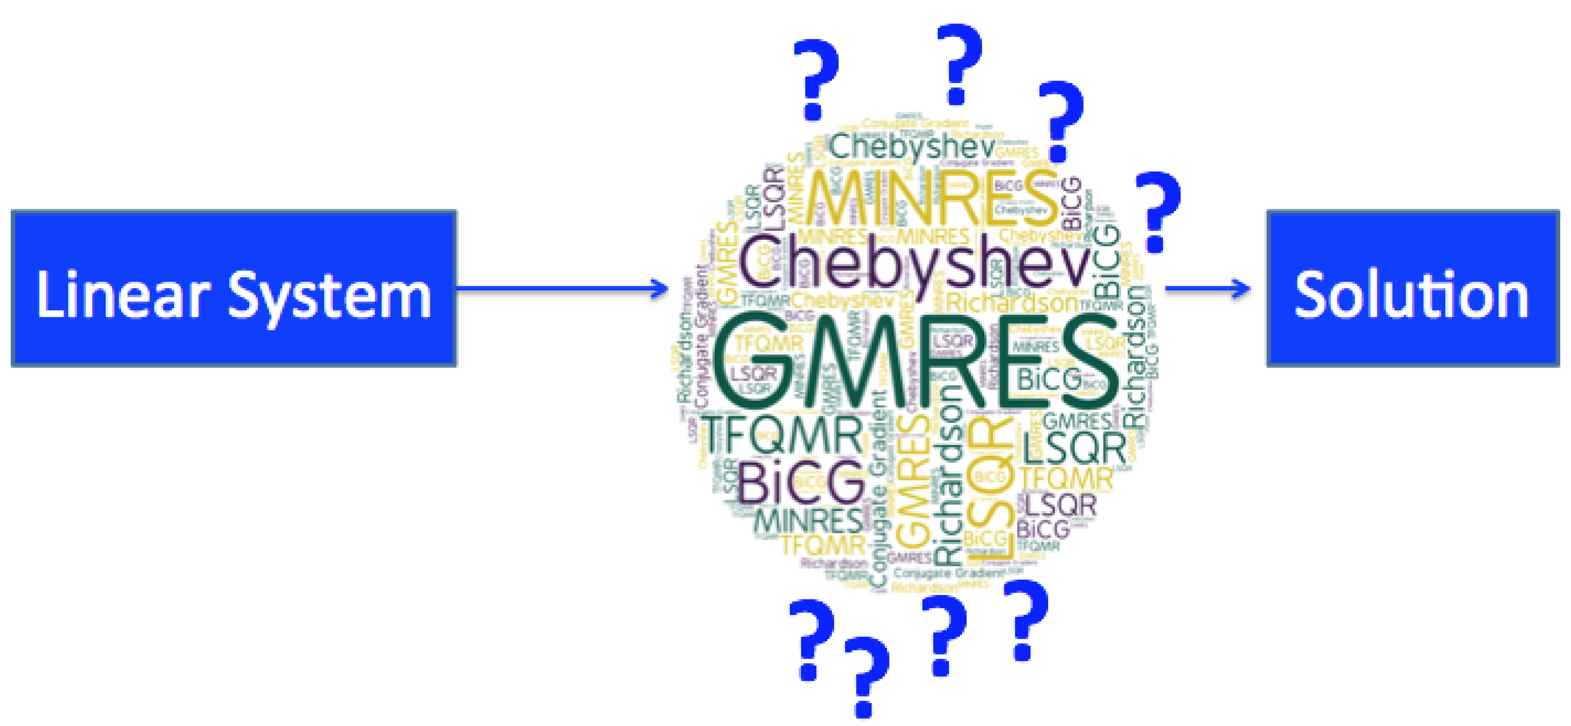

GOAL

Enable users to choose solver configurations that are likely to perform well for a given linear system.

CONTRIBUTIONS

- An accurate generalizable machine learning-based workflow for classifying arbitrary sparse linear systems.
- A parallel scaling model based on analytical communication estimates for systems that require large scale distributed memory resources.
- A comparatively less expensive set of features of the linear system, computed with and without matrix-free approach.
- A new approach of computing matrix features.

APPROACH

1. For small scales, use Machine Learning (ML) model [2] to suggest solvers.
2. Analyze the communication overheads of Krylov methods to generate a communication-based solver ranking.
3. For large scales, combine the ML model and the communication-based ranking [3].
4. Generate solver recommendations.

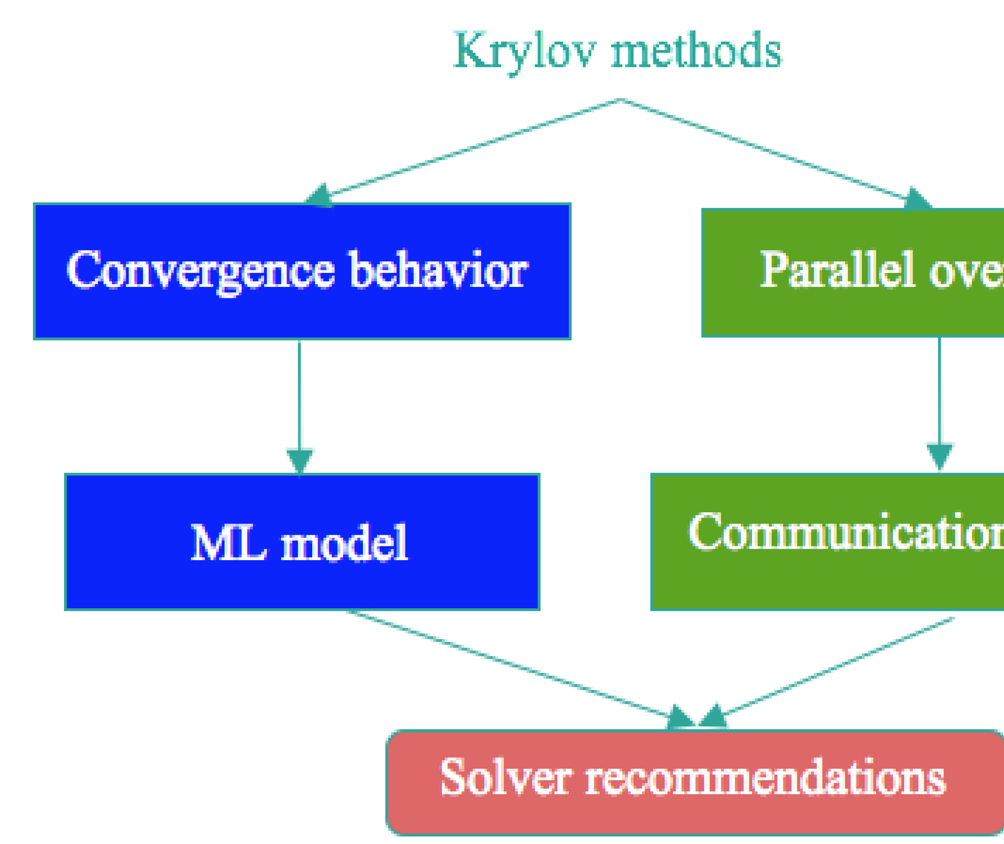

MATRIX FEATURES

- Characteristics of the linear systems
- Categories: Structural, size-based, eigenvalues.
- Examples: No. of non-zeros, matrix norm.
- Input for the convergence model.

Libraries used for computing features: Anamod, PETSc.

FEATURE COMPUTATION

- Matrix-full: Entire matrix is stored.
- Matrix-free: Matrix is not stored explicitly, instead matrix-vector product approximations are used.

REFERENCES

[1] Balay, Satish, et al. PETSc web page, http://www.mcs.anl.gov/petsc/ 2018.

[2] Jessup, Elizabeth, et al. *Performance-based numerical solver selection in the Lighthouse framework*. SIAM Journal on Scientific Computing, 2016.

[3] Sood, Kanika et al. *comparative Performance Modeling of Parallel Preconditioned Krylov Methods*, IEEE HPCC 2017.

SMALL SCALE SOLUTION

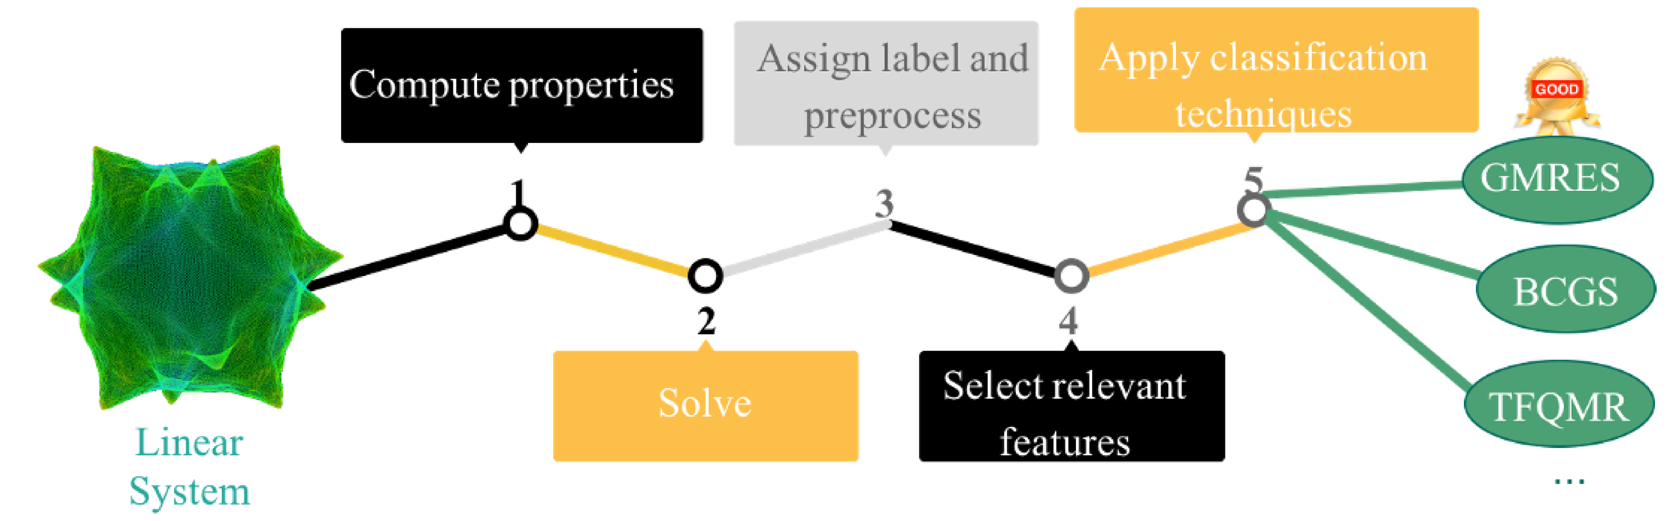

A low-overhead technique to select solution methods effectively, by training the ML model and suggesting based on the selected input problem characteristics.

LARGE SCALE PARALLEL SOLUTION

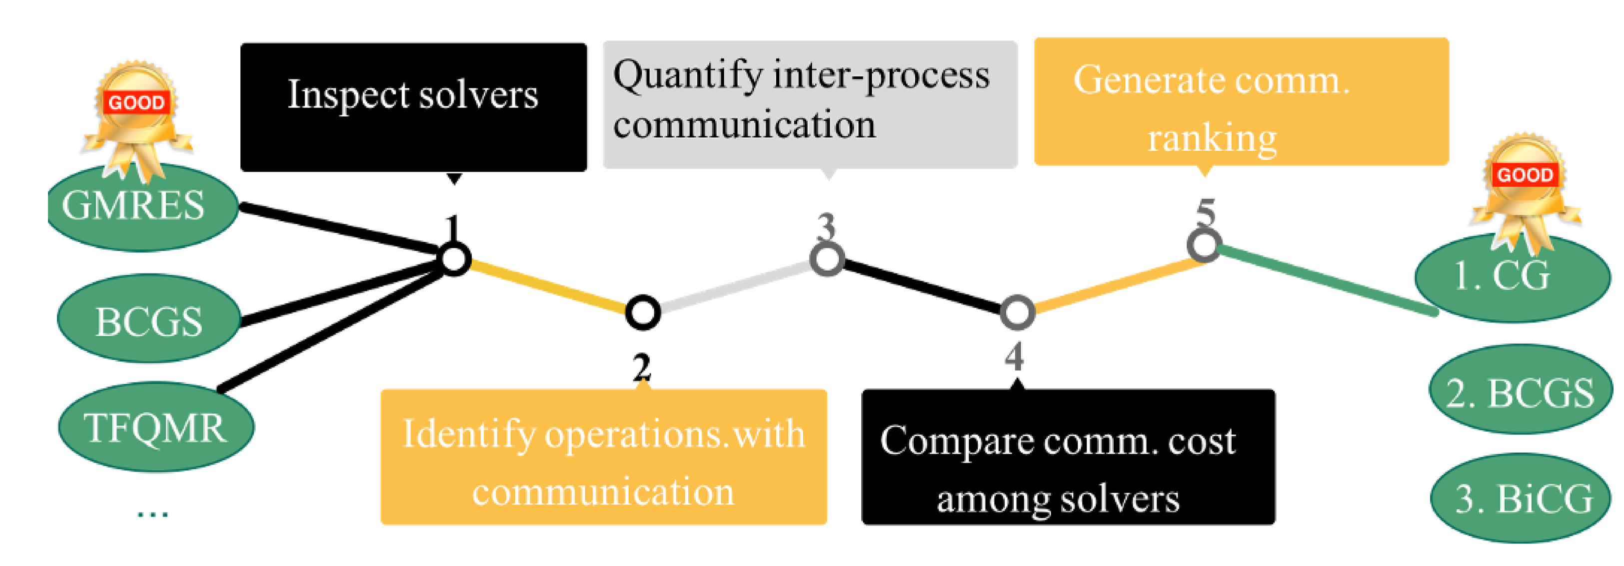

A new technique for modeling performance that captures performance variation at different parallelism scales and generating ranked list of solver suggestions.

CONCLUSION

- Capture the convergence behavior using ML model.
- Capture the parallel overhead based on communication.
- Demonstrate their effectiveness in PDE based applications.
- Matrix-free approach for feature computation.
- **Enable solver recommendations for sparse linear systems at different scales.**

CCF- 1219089, CNS-0821794

U.S. DEPARTMENT OF ENERGY
DE-AC02-06CH11357