# Improving Coarsening and Interpolation for Algebraic Multigrid

Jeff Butler

Hans De Sterck

Department of Applied Mathematics

University of Waterloo

(In Collaboration with Ulrike Meier Yang, LLNL)

# Outline

- Introduction: Algebraic Multigrid (AMG)
- Modification 1: PMIS Greedy vs. PMIS
- Modification 2: PMIS restricted to finer grid levels
- Modification 3: FF and FF1 Interpolation

# Introduction

- Solve: Au = f
- A, f  from PDE discretization
  - → sparse
- Parallel
  - → large problems: $10^9$ degrees of freedom
- Unstructured grid problems

# AMG Structure

## Setup Phase

*On each level* (m): -Select coarse grid points (coarsen)

-Define interpolation operator, $P^{(m)}$

-Define restriction and coarse-grid operators

$R^{(m)} = P^{(m)\,T}$ $A^{(m+1)} = P^{(m)\,T} A^{(m)} P^{(m)}$

## Solve Phase

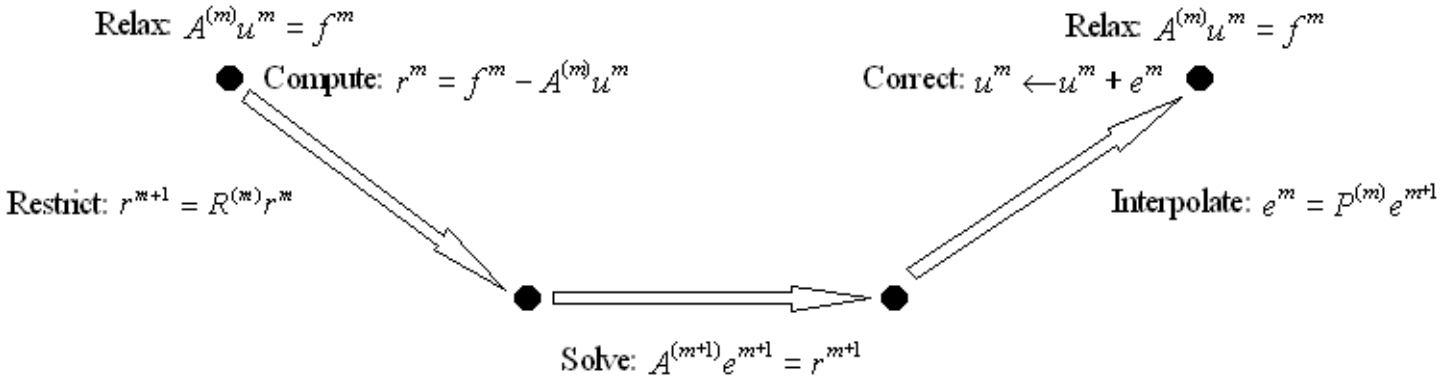

# AMG Complexity & Scalability

- Goal: Scalable Algorithm
  - O(n) operations per V-cycle
  - $\rho_{AMG}$ independent of n
- Operator Complexity: $$C_{op} = \frac{\sum_i nonzeros(A_i)}{nonzeros(A_0)}$$
- Measure of memory use, work in solve phase, and work required in coarsening and interpolation process
- E.g. 3D geometric multigrid:

$$C_{op} = 1 + \frac{1}{8} + \frac{1}{64} + \cdots < \frac{8}{7}$$

# Classical AMG Coarsening (Ruge, Stueben)

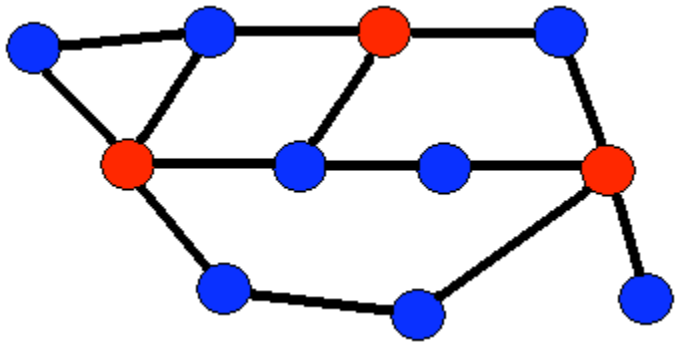

- (H1) All F-F connections require connection to a common C-point (good nearest neighbor interpolation)
- (H2) Maximal Independent Set:
  Independent: no two C-points are connected
  Maximal: If one more C-point is added, independence is lost

# Classical AMG Coarsening (Ruge, Stueben)

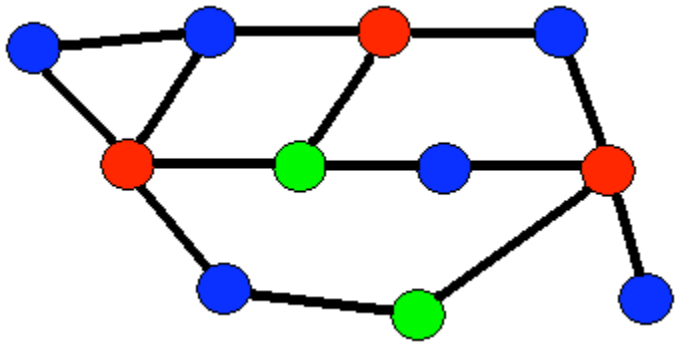

- (H1) All F-F connections require connection to a common C-point (good nearest neighbor interpolation)
- (H2) Maximal Independent Set:
  Independent: no two C-points are connected
  Maximal: If one more C-point is added, independence is lost

→ Enforce H1 rigorously with H2 as a guide (change F-points to C-points)

→ More C-points = higher complexity

# AMG Coarsenings

1. Ruge-Stueben (RS)
   - two passes: second pass to ensure that F-F have a common C
   - disadvantage: highly sequential

2. CLJP
   - based on parallel independent set algorithms developed by Luby and later by Jones & Plassman
   - ensures that F-F have a common C

3. PMIS (De Sterck, Yang)
   - Parallel Modified Independent Set (PMIS) (Luby)
   - do not enforce heuristic H1 (F-F without a common C)

# PMIS Selection Step 1

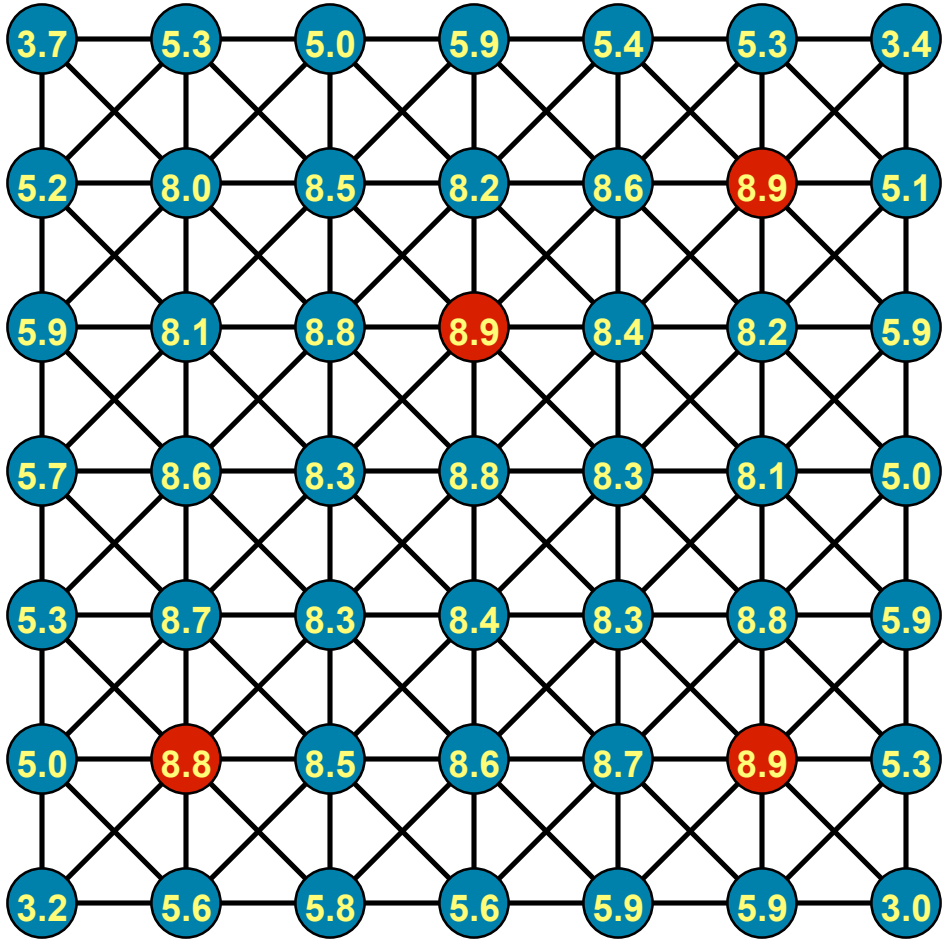


- select C-points with maximal local measure
- make neighbors F-points
- remove neighbor edges

** This grid animation courtesy of Ulrike Yang, LLNL*

# PMIS Remove and Update Step 1

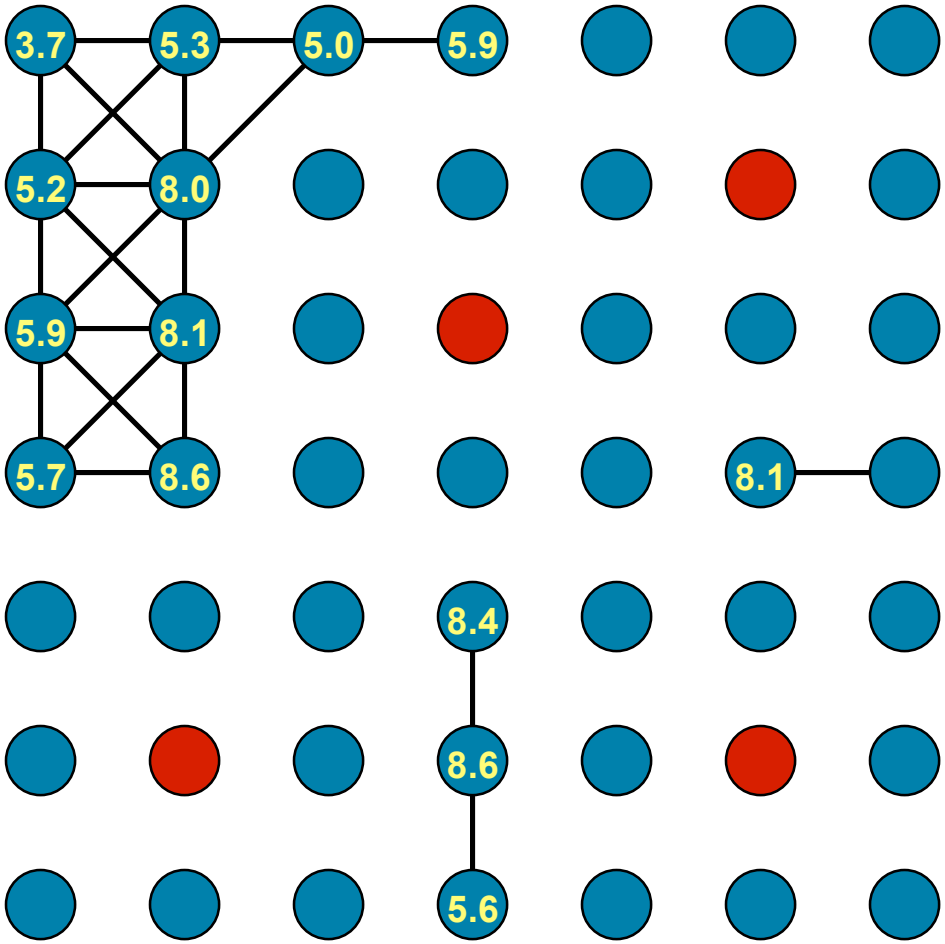


- select C-points with maximal local measure
- make neighbors F-points
- remove neighbor edges

** This grid animation courtesy of Ulrike Yang, LLNL*

# PMIS Selection Step 2

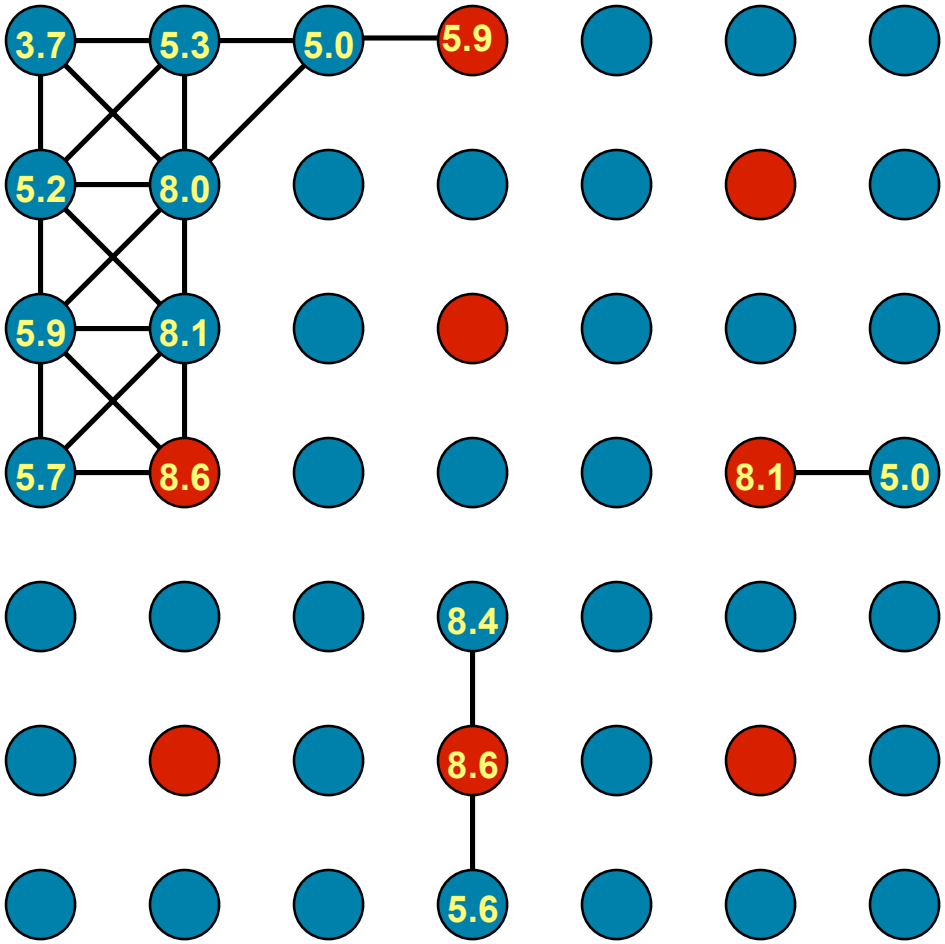


- select C-points with maximal local measure
- make neighbors F-points
- remove neighbor edges

** This grid animation courtesy of Ulrike Yang, LLNL*

# PMIS Remove and Update Step 2

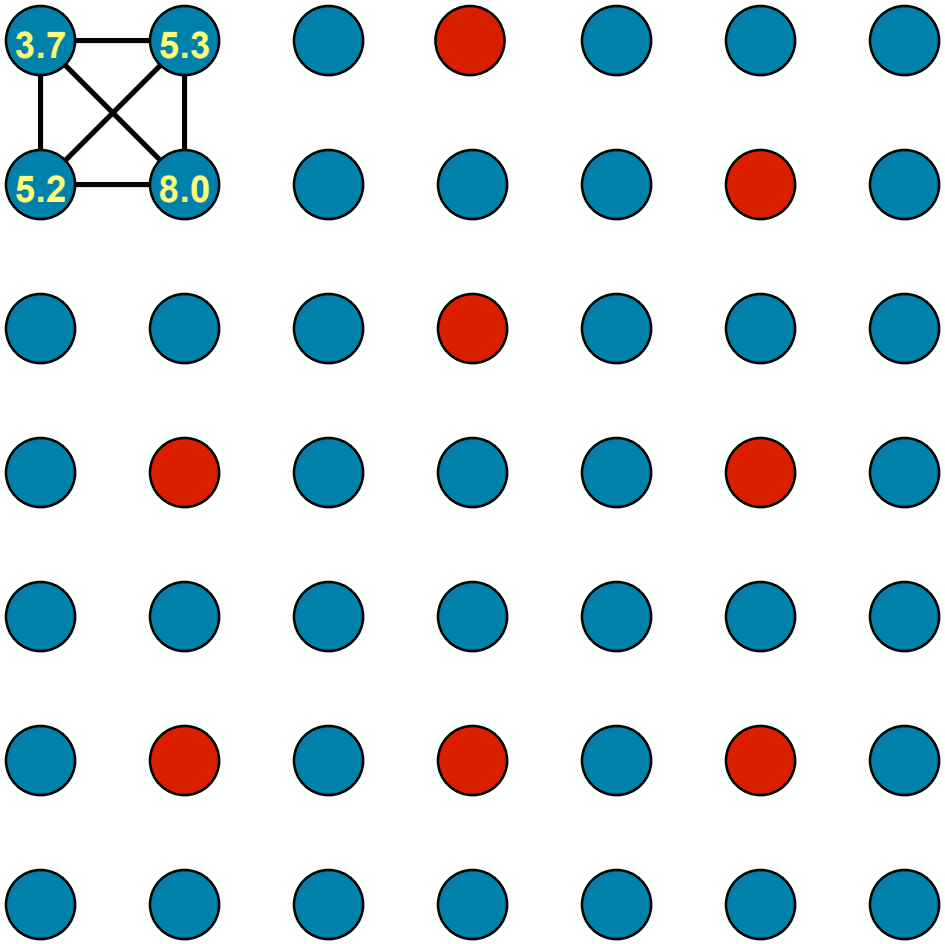


- select C-points with maximal local measure
- make neighbors F-points
- remove neighbor edges

** This grid animation courtesy of Ulrike Yang, LLNL*

# PMIS Final Coarsening

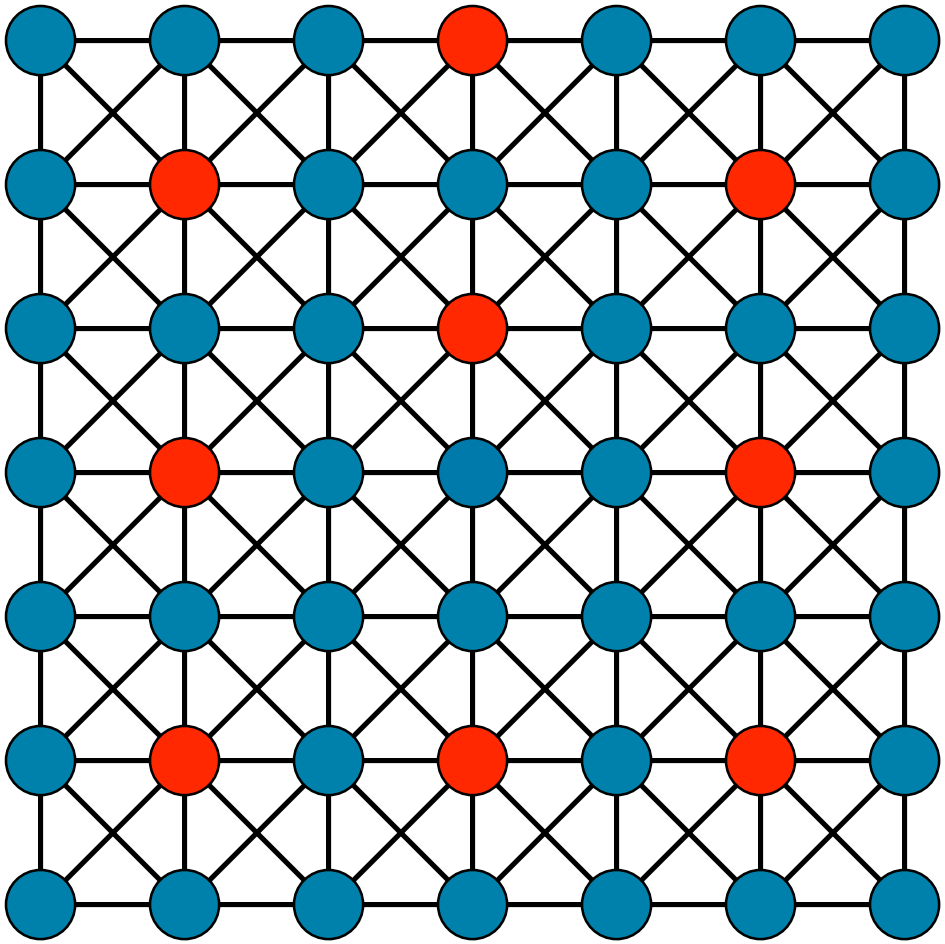

- select C-points with maximal local measure
- make neighbors F-points
- remove neighbor edges

** This grid animation courtesy of Ulrike Yang, LLNL*

# PMIS Summary (De Sterck, Yang & Heys: 2006)

- PMIS coarsening worked well for many problems (with GMRES)
- for some problems, convergence degraded compared to RS
  - decreased accuracy in interpolation due to an inadequate amount of C-points
- One solution: add C-points (RS, CLJP)
- Other possibilities:
  - modify PMIS to (hopefully) improve convergence (PMIS Greedy)
  - combine PMIS coarsening with RS/CLJP coarsening
  - use distance-two C-points as long-range interpolation for F-F connections without a common C-point (F-F interpolation)

# Modification 1: PMIS Greedy vs. PMIS

- PMIS Greedy
  - same procedure as in PMIS
  - increase the measure of an unassigned point if it is strongly connected to a newly assigned F-point in each pass
  - helps to remove some randomness in the grid structure
  - interpolation has a chance to be more accurate

***5-pt Laplacian:*** *= finest grid* *= second finest grid* *=third finest grid*

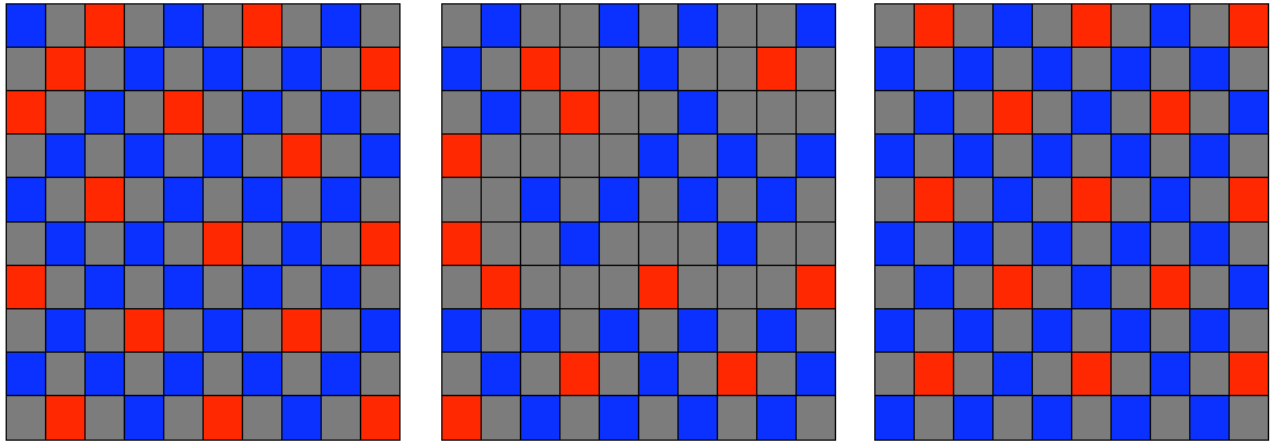

***RS*** ***PMIS*** ***PMIS Greedy***

# Results for an "Easy" Problem

- 27-pt laplacian, 1 processor, $128^3$ dof
- AMG

| | $C_{op}$ | # iterations | $t_{setup}$ (s) | $t_{solve}$ (s) | $t_{total}$ (s) |
|---|---|---|---|---|---|
| RS | 2.67 | 9 | 209.88 | 25.24 | 235.12 |
| PMIS | 1.10 | 47 | 33.67 | 65.08 | 98.75 |
| PMIS Greedy | 1.12 | 38 | 36.37 | 54.01 | 90.38 |

- AMG-GMRES(5)

| | $C_{op}$ | # iterations | $t_{setup}$ (s) | $t_{solve}$ (s) | $t_{total}$ (s) |
|---|---|---|---|---|---|
| RS | 2.67 | 6 | 208.94 | 74.72 | 283.66 |
| PMIS | 1.10 | 17 | 33.72 | 42.09 | 75.81 |
| PMIS Greedy | 1.12 | 15 | 36.39 | 36.99 | 73.38 |

# Results for a More Difficult Problem

- 3D elliptic PDE with jumps in the coefficient $a$

$$(au_x)_x + (au_y)_y + (au_z)_z = 1$$

- 1 processor, $80^3$ dof
- AMG-GMRES(5)

| | $C_{op}$ | # iterations | $t_{setup}$ (s) | $t_{solve}$ (s) | $t_{total}$ (s) |
|---|---|---|---|---|---|
| RS | 21.54 | | *Ran out* | *of Memory* | |
| PMIS | 2.46 | 188 | 4.06 | 77.34 | 81.40 |
| PMIS Greedy | 2.54 | 144 | 4.19 | 59.90 | 64.09 |

# PMIS vs. PMIS Greedy Summary

- better convergence and timing for some problems
- improvement not major
- for other problems, PMIS Greedy not much better (and sometimes even worse) than PMIS with respect to convergence and timing
- slightly higher operator complexities for PMIS Greedy
- PMIS Greedy will require more communication in parallel than PMIS

# Modification 2: Restrict PMIS to Finer Grid Levels

- Perform PMIS on first $g$ grid levels and RS on all remaining levels
- Recall:

  *3D geometric multigrid*: $C_{op} = 1 + \frac{1}{8} + \frac{1}{64} + \cdots < \frac{8}{7}$

- Advantage:
  - PMIS reduces operator complexity (compared to RS) where it makes the biggest difference (finer levels)
  - RS produces more "structured" grids leading to better interpolation on coarser levels where impact on operator complexity is reduced

# Results: 27-pt Laplacian Problem

- 1 processor, $128^3$ dof
- AMG

| | $C_{op}$ | # iterations | $t_{setup}$ (s) | $t_{solve}$ (s) | $t_{total}$ (s) |
|---|---|---|---|---|---|
| PMIS (all levels) | 1.10 | 47 | 33.67 | 65.08 | 98.75 |
| PMIS (first 2 levels) | 1.12 | 27 | 34.55 | 38.05 | 72.60 |

- AMG-GMRES(5)

| | $C_{op}$ | # iterations | $t_{setup}$ (s) | $t_{solve}$ (s) | $t_{total}$ (s) |
|---|---|---|---|---|---|
| PMIS (all levels) | 1.10 | 17 | 33.72 | 42.09 | 75.81 |
| PMIS (first 2 levels) | 1.12 | 12 | 34.58 | 30.55 | 65.13 |

# Results: 3D elliptic PDE with jumps in *a*

- $(au_x)_x + (au_y)_y + (au_z)_z = 1$
- AMG, 1 processor, $120^3$ dof

| | $C_{op}$ | # iterations | $t_{setup}$ (s) | $t_{solve}$ (s) | $t_{total}$ (s) |
|---|---|---|---|---|---|
| PMIS (all levels) | 2.44 | >> 200 | *Slow to* | *converge* | |
| PMIS (first level only) | 7.2131 | 23 | 47.93 | 44.32 | 92.25 |
| PMIS (first 2 levels) | 3.49 | 168 | 20.58 | 190.16 | 210.74 |

- AMG-GMRES(5), 1 processor, $80^3$ dof

| | $C_{op}$ | # iterations | $t_{setup}$ (s) | $t_{solve}$ (s) | $t_{total}$ (s) |
|---|---|---|---|---|---|
| PMIS (all levels) | 2.46 | 188 | 4.06 | 77.34 | 81.40 |
| PMIS (first level only) | 6.97 | 11 | 12.76 | 8.38 | 21.14 |
| PMIS (first 3 levels) | 2.56 | 105 | 4.24 | 44.20 | 48.44 |

* $C_{op}$(RS) = 21.54

# PMIS on Higher Levels Only: Summary

- Improvement in convergence properties for a variety of problems (although some only small)
- How many levels PMIS?
  - problem dependent
- Trade-off between memory (increased operator complexity) and speed of convergence/execution
- Similar approach for parallel case is possible
  - CLJP has similar performance characteristics as RS

# Modification 3: Modified FF Interpolation

- FF Interpolation* (De Sterck, Yang: Copper 2005):

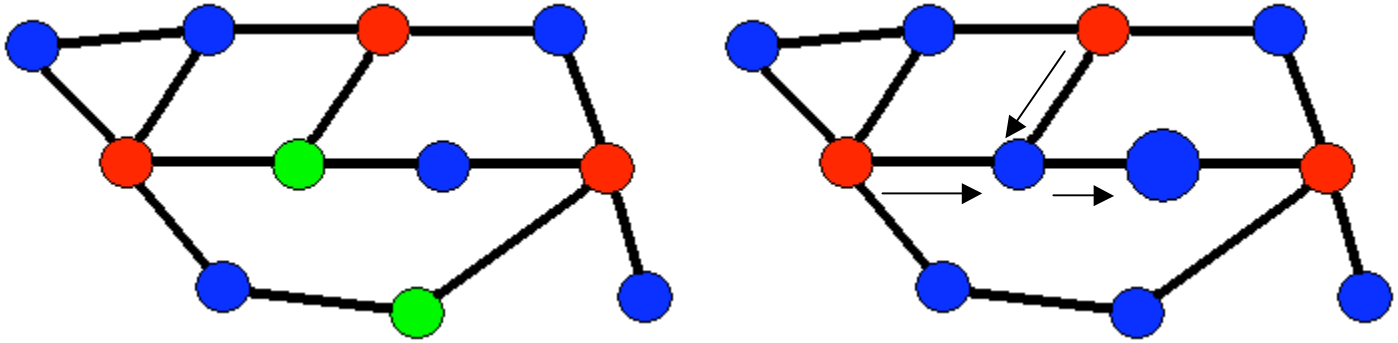

- When a strong F-F connection is encountered, do not add a C-point, but extend interpolation to distance-two C-points
- No C-points added, but get larger interpolation stencils (and therefore somewhat increased operator complexity)

*Closely related to Stueben's Standard Interpolation

# FF1 Interpolation

- Modified FF Interpolation (FF1)

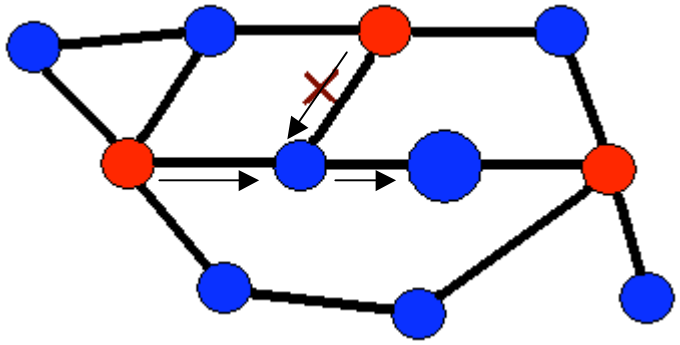

- To reduce operator complexity, only include one distance-two C-point when a strong FF connection is encountered
- Setup time, complexity are reduced

# Results: 7-pt Laplacian Problem

- PMIS coarsening, 1 processor, $128^3$ dof
- AMG

| | $C_{op}$ | # iterations | $t_{setup}$ (s) | $t_{solve}$ (s) | $t_{total}$ (s) |
|---|---|---|---|---|---|
| Regular | 2.36 | 77 | 16.63 | 85.93 | 102.56 |
| FF | 4.80 | 13 | 83.81 | 22.86 | 106.67 |
| FF1 | 3.68 | 15 | 44.22 | 22.07 | 66.29 |

- AMG-GMRES(5)

| | $C_{op}$ | # iterations | $t_{setup}$ (s) | $t_{solve}$ (s) | $t_{total}$ (s) |
|---|---|---|---|---|---|
| Regular | 2.36 | 20 | 16.67 | 35.26 | 51.93 |
| FF | 4.80 | 9 | 83.87 | 22.97 | 106.84 |
| FF1 | 3.68 | 9 | 43.85 | 19.85 | 63.70 |

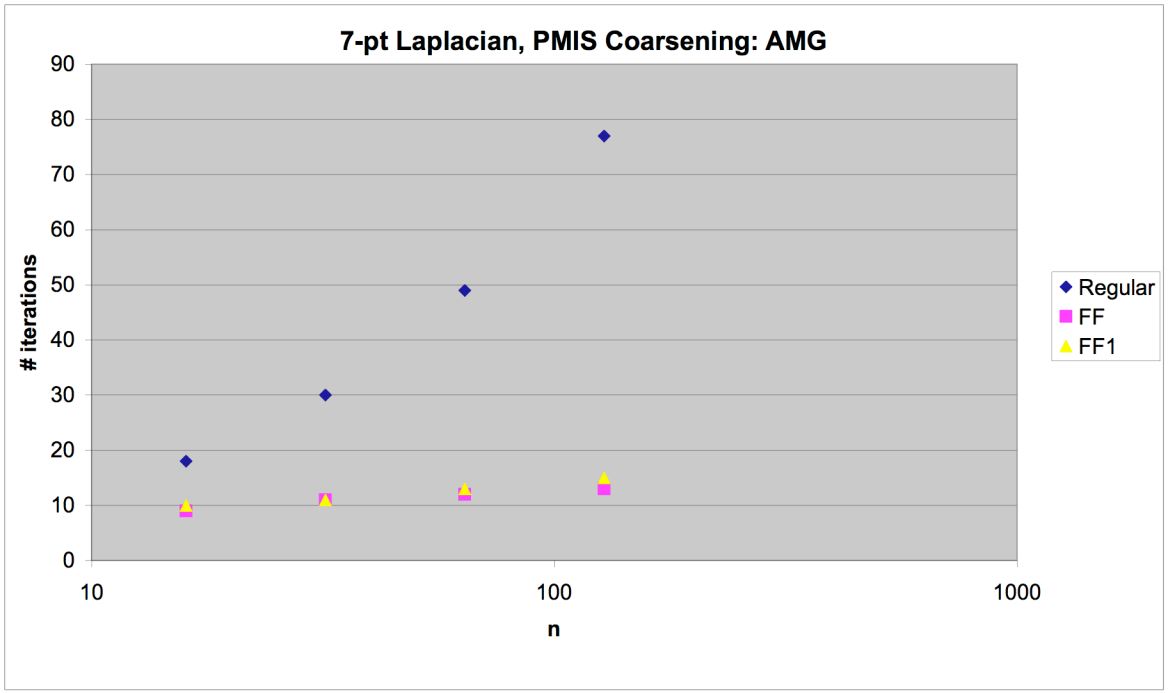
7-pt Laplacian, PMIS Coarsening: AMG
# iterations
90
80
70
60
50
40
30
20
10
0
10
100
1000
n
Regular
FF
FF1

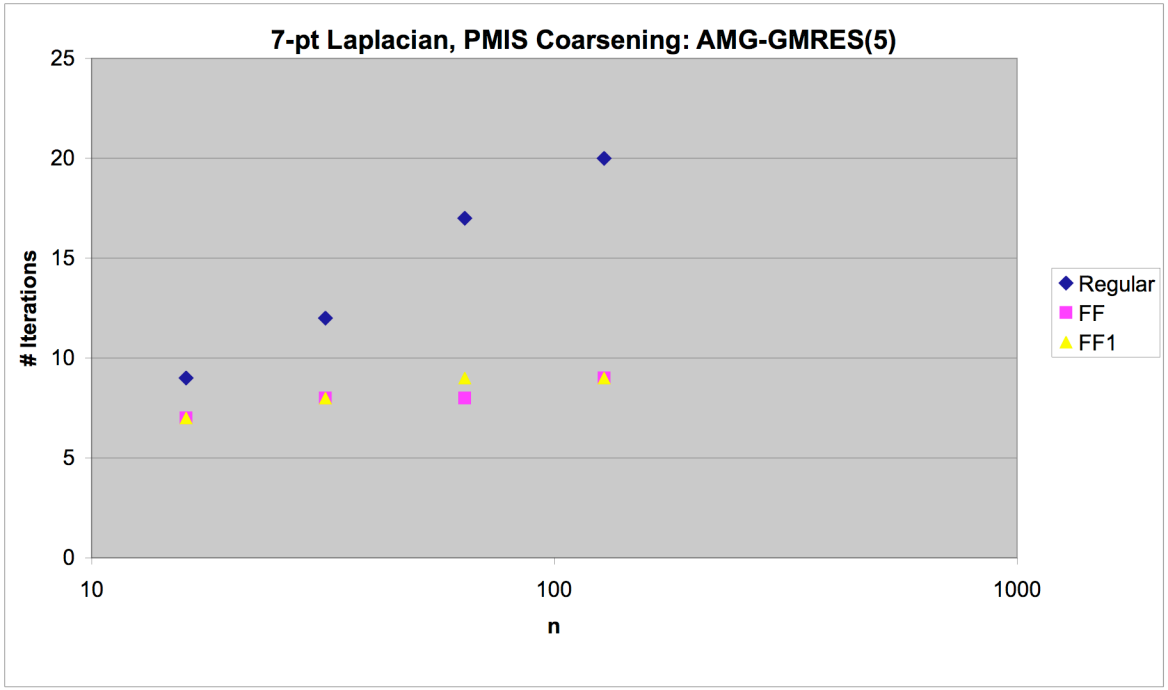
7-pt Laplacian, PMIS Coarsening: AMG-GMRES(5)
# Iterations
25
20
15
10
5
0
10
100
1000
n
Regular
FF
FF1

# Results: 3D elliptic PDE with jumps in *a*

- $(au_x)_x + (au_y)_y + (au_z)_z = 1$
- AMG, 1 processor, $120^3$ dof

| | $C_{op}$ | # iterations | $t_{setup}$ (s) | $t_{solve}$ (s) | $t_{total}$ (s) |
|---|---|---|---|---|---|
| Regular | 2.44 | >> 200 | *Slow to* | *converge* | |
| FF | 4.94 | 14 | 62.95 | 20.54 | 83.49 |
| FF1 | 3.84 | 18 | 35.36 | 22.47 | 57.83 |

- AMG-GMRES(5), 1 processor, $80^3$ dof

| | $C_{op}$ | # iterations | $t_{setup}$ (s) | $t_{solve}$ (s) | $t_{total}$ (s) |
|---|---|---|---|---|---|
| Regular | 2.46 | 188 | 4.06 | 77.34 | 81.40 |
| FF | 4.90 | 9 | 16.46 | 5.34 | 21.80 |
| FF1 | 3.85 | 10 | 9.67 | 4.99 | 14.66 |

# Modified FF Interpolation (FF1) Summary

- PMIS with FF or FF1 is scalable! (comparable to RS with classical interpolation, but lower complexity; no need for GMRES)
- FF1 reduces setup time compared to FF
- FF1 iteration number kept low as with FF
- FF1 operator complexity reduced compared to FF
- PMIS with FF1 seems a good parallel alternative for RS with a second pass and classical interpolation

# Thanks…


- Ulrike Meier Yang and the Center for Applied Scientific Computing, Lawrence Livermore National Laboratory

- travel support University of Waterloo Graduate Studies

# Algebraic Multigrid (AMG)

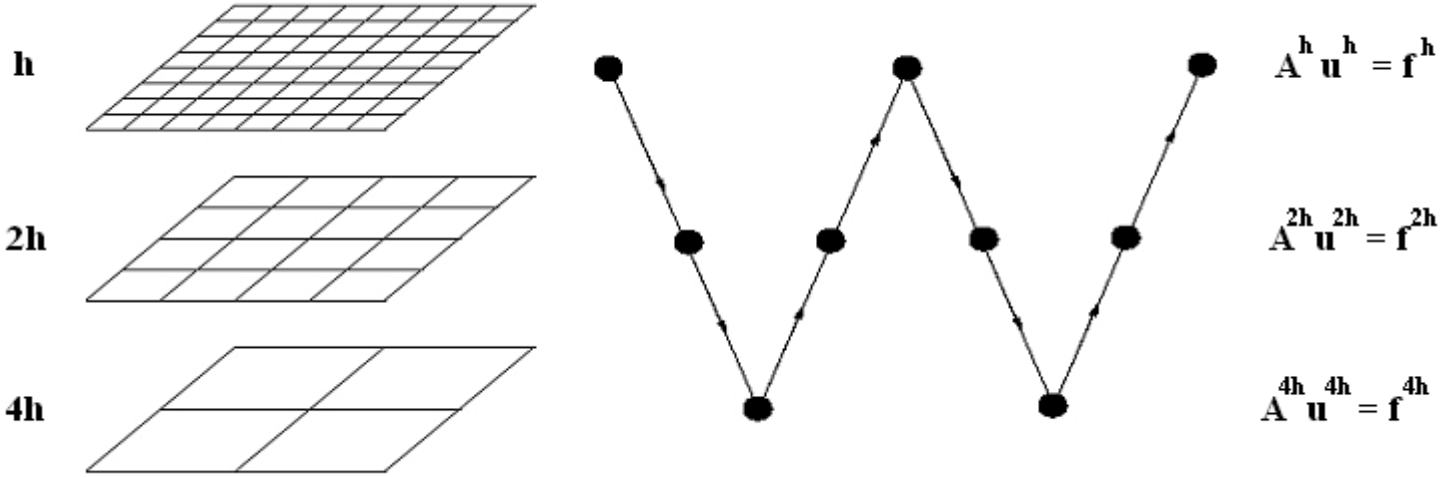


- Multiple levels (coarsen to some minimum)
- Iterative
- AMG suitable for unstructured grids

# PMIS Coarsening

- weighted independent set algorithm:
  - points, i, that influence many equations ($\lambda_i$ large) are good candidates for C-points
- add random number between 0 and 1 to $\lambda_i$ to break ties
- pick C-points with maximal measures (like in CLJP), then make all their neighbors fine points (like in RS)
- proceed until all points have been assigned as fine or coarse

# Conclusions

- PMIS vs. PMIS Greedy coarsening
  - PMIS Greedy does not provide a considerable advantage
  - any advantage will be degraded in a parallel implementation
- Restricting PMIS to finer grid levels
  - considerable improvement over uniform RS and PMIS coarsening possible
  - scalability possible
  - trade-off between memory use and convergence speed
- FF1 Interpolation (with PMIS coarsening)
  - comparable performance to FF (both better than regular interpolation)
  - scalability possible
  - better suited to parallel implementation than FF