Coarsening and Interpolation in Algebraic Multigrid: a Balancing Act

Hans De Sterck

Department of Applied Mathematics

University of Waterloo, Ontario, Canada

Ulrike Meier Yang

Center for Applied Scientific Computing

Lawrence Livermore National Laboratory

California, USA

University of Waterloo

CASC

Outline

- introduction: AMG
- AMG coarsening: classical versus more aggressive
- convergence problems with more aggressively coarsened grids
- improved, long-range interpolation methods
- results
- conclusions and future work

Introduction

- *solve* $\mathbf{Au} = \mathbf{f}$
- $\mathbf{A}$ *from 3D PDE – sparse!*
- *large problems (10^9 dof) - parallel*
- *unstructured grid problems*

Algebraic Multigrid (AMG)

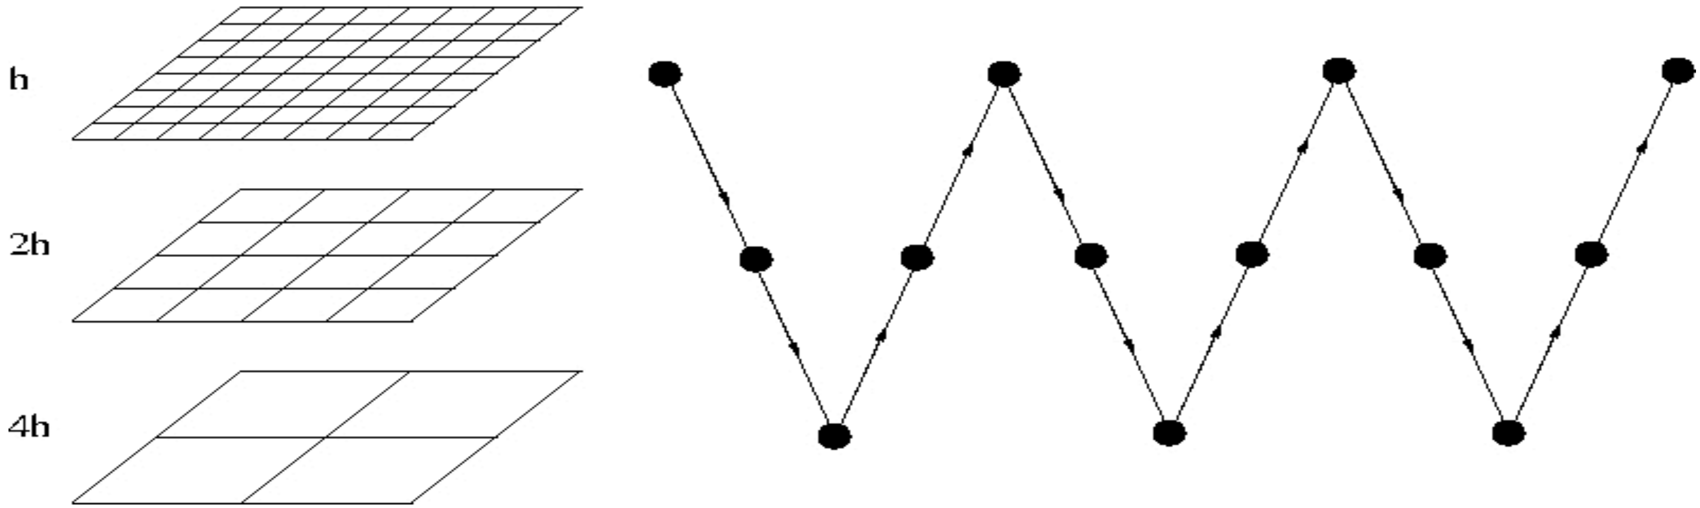

- *multi-level*
- *iterative*
- *algebraic: suitable for unstructured!*

AMG complexity - scalability

- *scalable algorithm:*

  O(n) operations per V-cycle (C_{op} bounded)

  AND

  number of V-cycles independent of n

  (ρ_{AMG} independent of n)

- *Operator complexity* $C_{op} = \dfrac{\sum nonzeros(A_i)}{nonzeros(A_0)}$

e.g., 3D: $C_{op} = 1 + 1/8 + 1/64 + \ldots < 8/7$

measure of memory use, and work in solve phase

Classical AMG coarsening

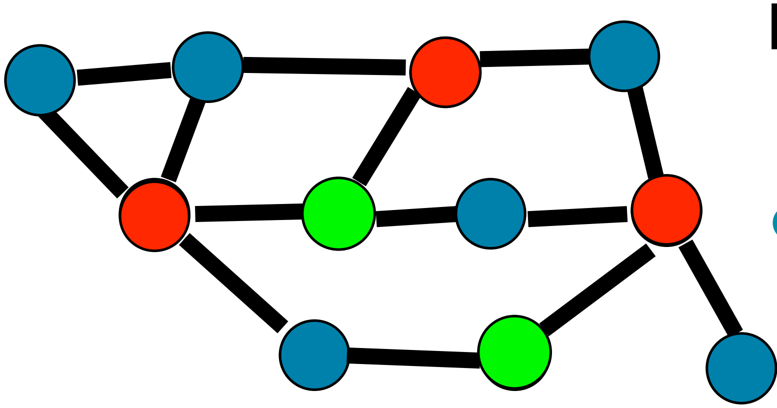

- (C1) Maximal Independent Set:

Independent: no two C-points are connected

Maximal: if one more C-point is added, the independence is lost

- (C2) All F-F connections require connections to a common C-point (for good nearest-neighbor interpolation)
- *F*-points have to be changed into *C*-points, to ensure (C2); (C1) is violated

more C-points, higher complexity

Classical coarsening: scalability results

- *example: finite difference Laplacian, parallel CLJP coarsening algorithm*

- *2D (5-point): near-optimal scalability* **(250^2 dof/proc)**

Procs	C_{op}	t_{tot}	Iter
16	4.48	2.89	9
64	4.50	3.85	9
256	4.50	5.01	9

Classical coarsening: complexity growth in some cases

- *3D (7-point): complexity growth*

dof	C_{op}	Iter
32^3	16.17	8
64^3	22.51	11

- *increased memory use, long solution times, long setup times, loss of scalability*

Classical coarsening: complexity growth in some cases

- *4D (9-point), 5D (11-point): complexity growth!!*

	dof	C_{op}	Iter
4D	20^4	127.5	8
5D	9^5	256.9	5

- *excessive memory use*

our approach to reduce complexity: PMIS (parallel modified independent set)

- *do not add C points for strong F-F connections that do not have a common C point*

- *less C points, reduced complexity, but worse convergence factors expected*

- *combine with GMRES acceleration*

- *in many cases (3D…), large gains*

PMIS coarsening: reduce complexity

- *finite difference Laplacian (CLJP-PMIS+GMRES)*

	dof	C_{op}	Iter	t_{tot}
2D	120^2	4.16	12	0.22
	120^2	1.90	24	0.24
3D	100^3	25.94	12	129.42
	100^3	2.36	20	27.68
4D	20^4	127.5	8	88.39
	20^4	2.95	11	4.31
5D	9^5	256.9	5	73.92
	8^5	3.14	8	0.91
	20^5	4.02	12	181.93

Convergence problems on PMIS-coarsened grids

- *PMIS coarsening works well for many problems*
- *for some problems, too many iterations are necessary because interpolation is not accurate enough ("not enough C-points")*
- *one solution: add C-points (CLJP…)*
- *other solution: use distance-two C-points for interpolation = long-range interpolation*
    - *Stuebe's multipass interpolation*
    - *F-F interpolation*

Convergence problems

- *3D elliptic PDE with jumps in coefficient a*

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

- *1000 processors, 40^3 dof/proc*

	t_{tot}	C_{op}	Iter
CLJP	52.48	17.00	17
PMIS	211.79	2.40	686

- *remedy: improve interpolation used with PMIS*

classical AMG Interpolation

- *after relaxation:*

$$Ae \approx 0 \text{ (relative to e)}$$

- *heuristic: error after interpolation should also satisfy this relation approximately*

- *derive interpolation from:*

$$a_{ii}e_i + \sum_{j \in C} a_{ij}e_j + \sum_{j \in F} a_{ij}e_j = 0 \qquad \forall i \in F$$

classical AMG interpolation

$$a_{ii}e_i + \sum_{j\in C} a_{ij}e_j + \sum_{j\in F} a_{ij}e_j = 0 \qquad \forall i \in F$$

- *"large" a_{ij} should be taken into account accurately*
- *"strong connections": i strongly depends on j (and j strongly influences i) if*

$$-a_{ij} \geq \theta \max_{k \neq i} \{-a_{ik}\}, \qquad 0 < \theta \leq 1$$

with strong threshold θ

classical AMG interpolation

- *strong F-F connections interpolated from common C-point*
- *interpolation only from nearest-neighbor C-points*

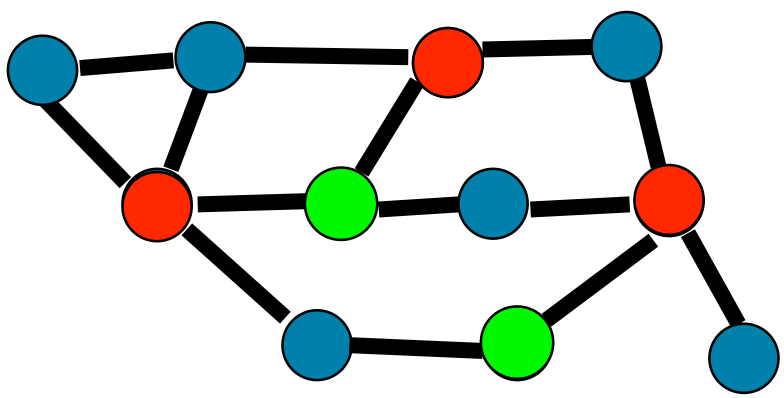

(1) Stueben's multipass interpolation

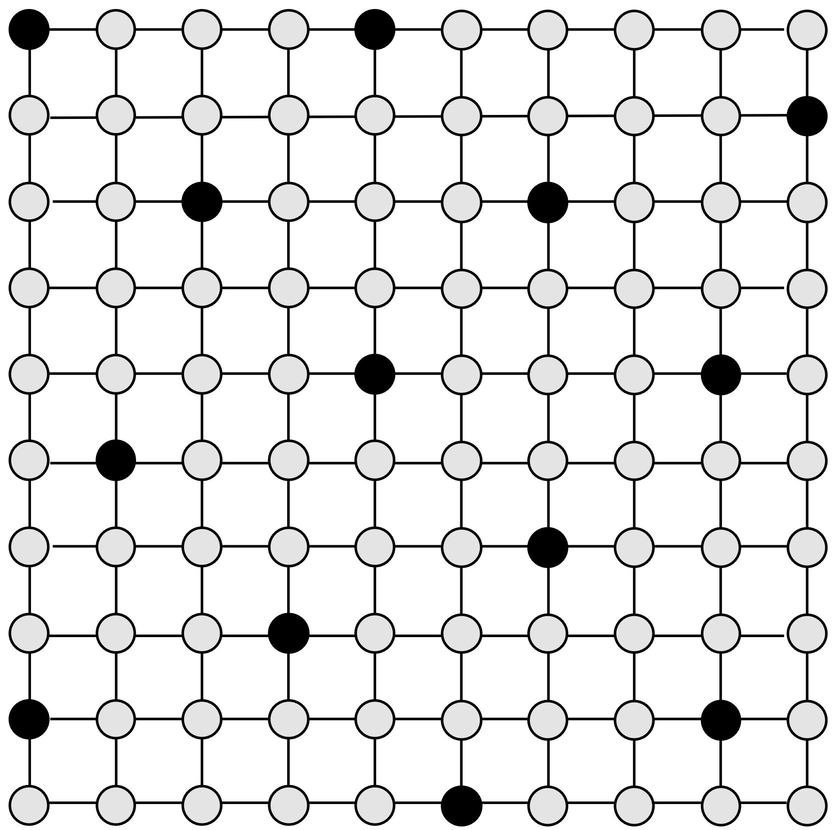

1st pass:

Coarse points

Multipass interpolation

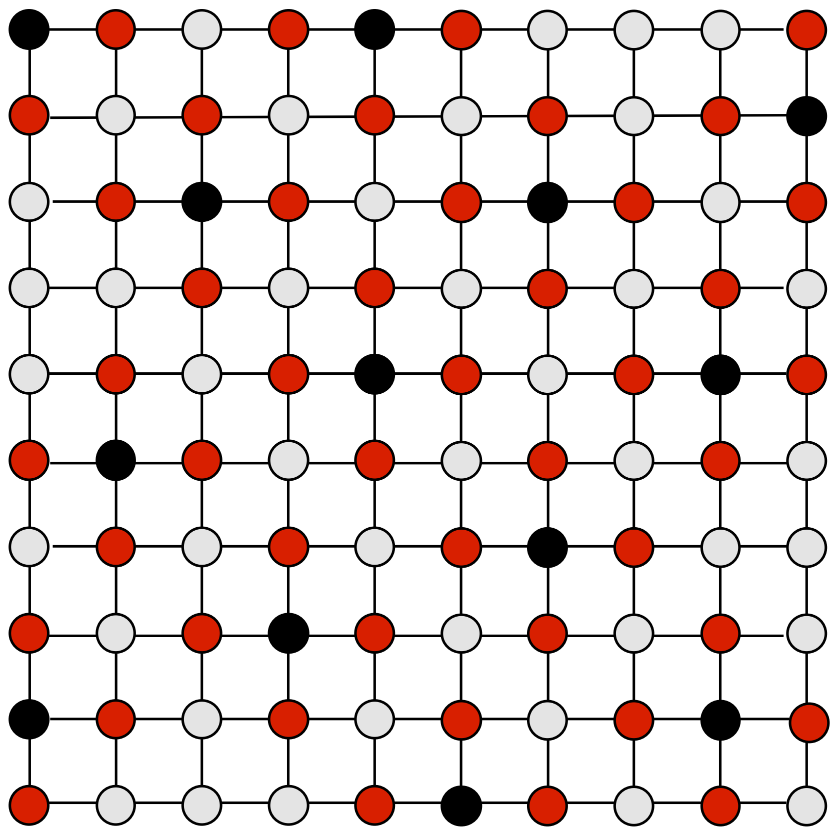

2nd pass:

direct interpolation from coarse C-neighbor

Multipass interpolation

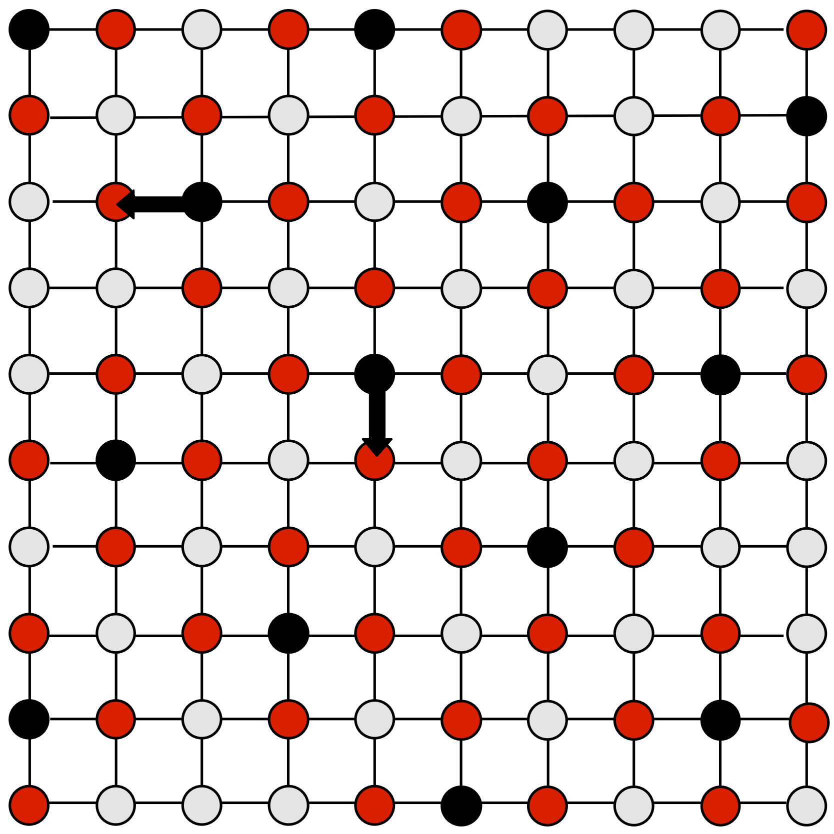

2^{nd} pass:

direct interpolation from coarse C-neighbor

Multipass interpolation

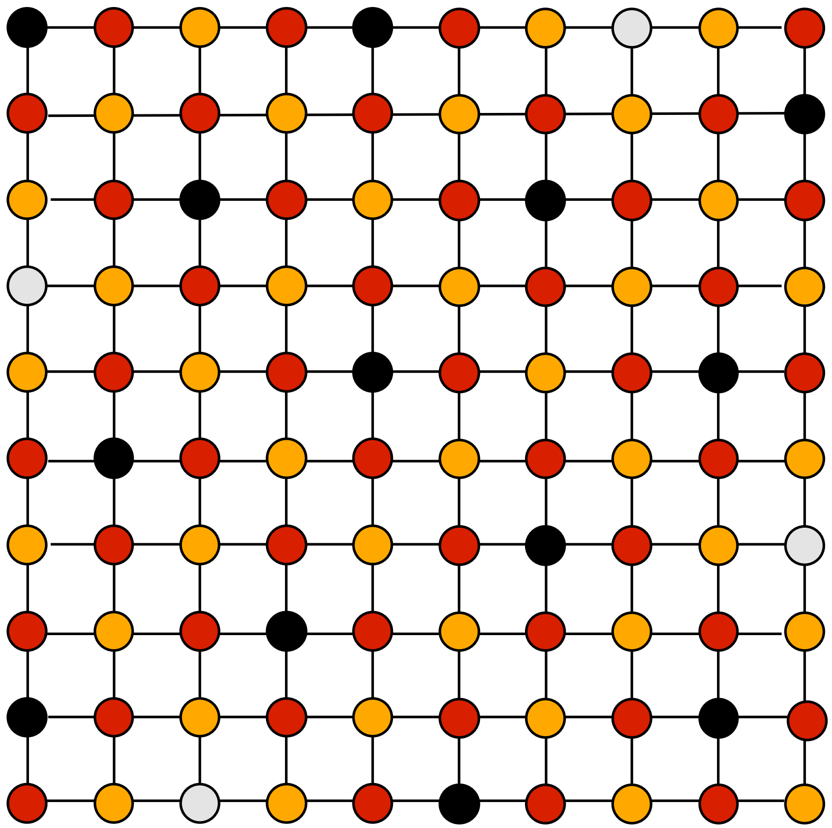

3rd pass:

direct interpolation from coarse F-neighbor (indirectly from distance-2 C-point)

Multipass interpolation

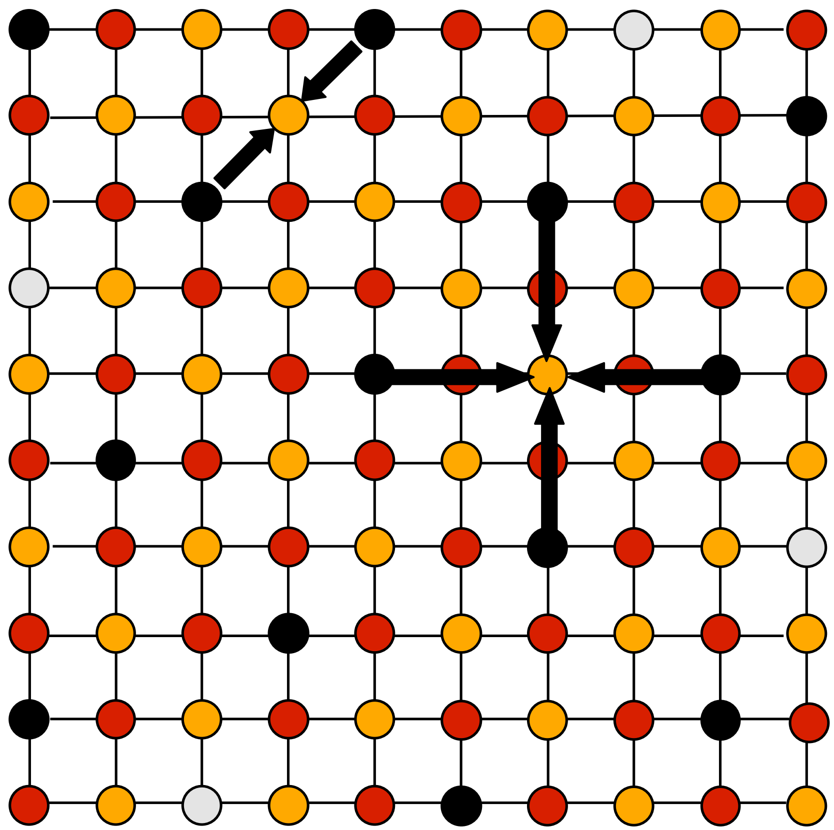

3rd pass:

direct interpolation from coarse F-neighbor (indirectly from distance-2 C-point)

Multipass interpolation

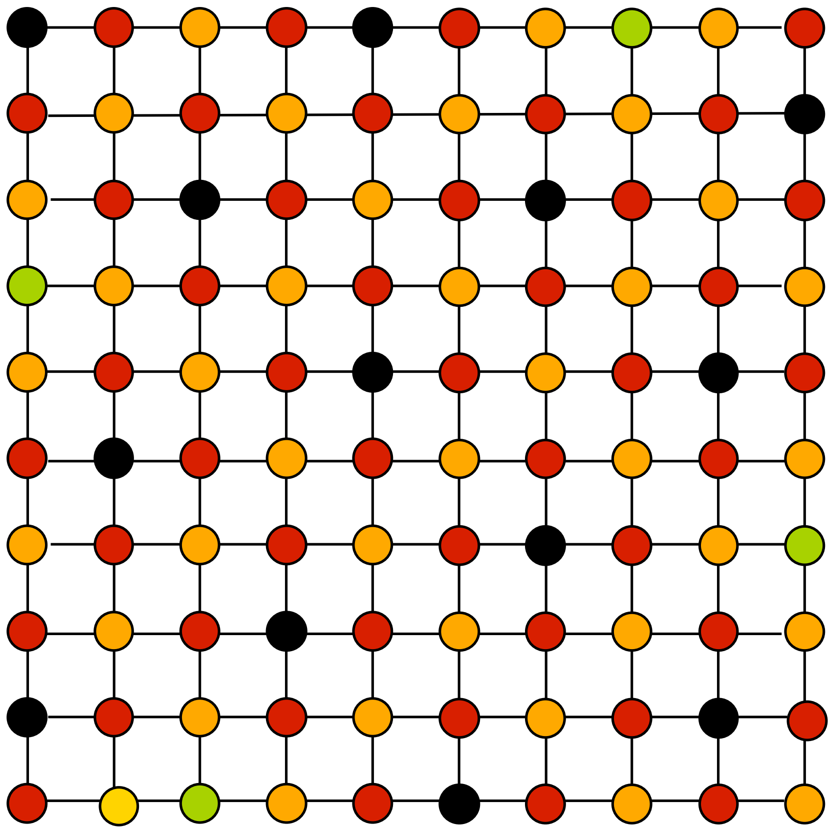

Final pass

Multipass interpolation

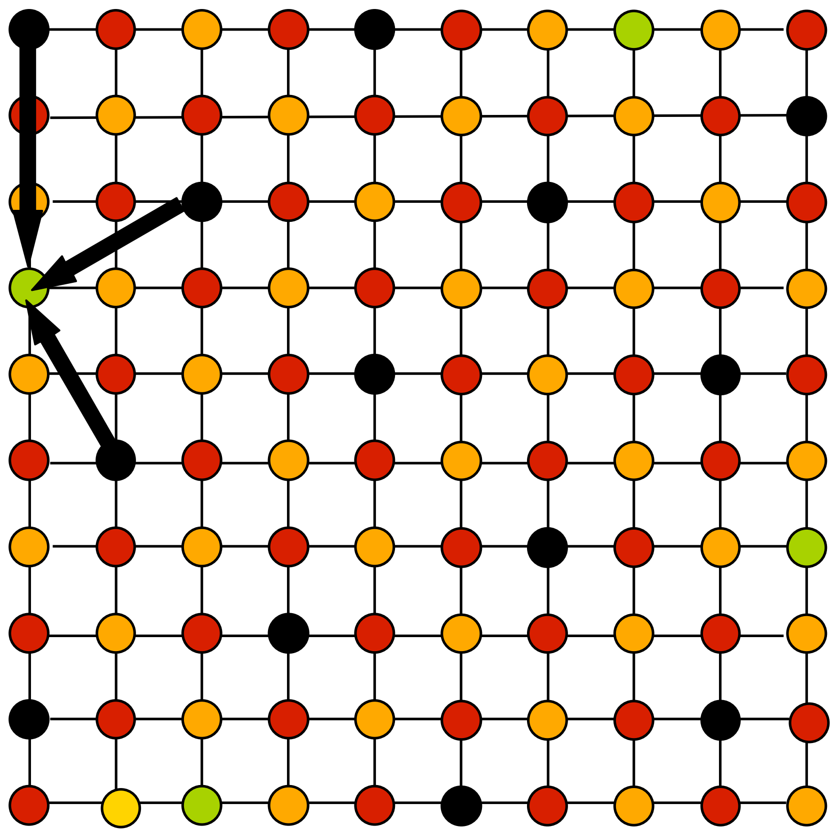

Final pass

(2) F-F interpolation

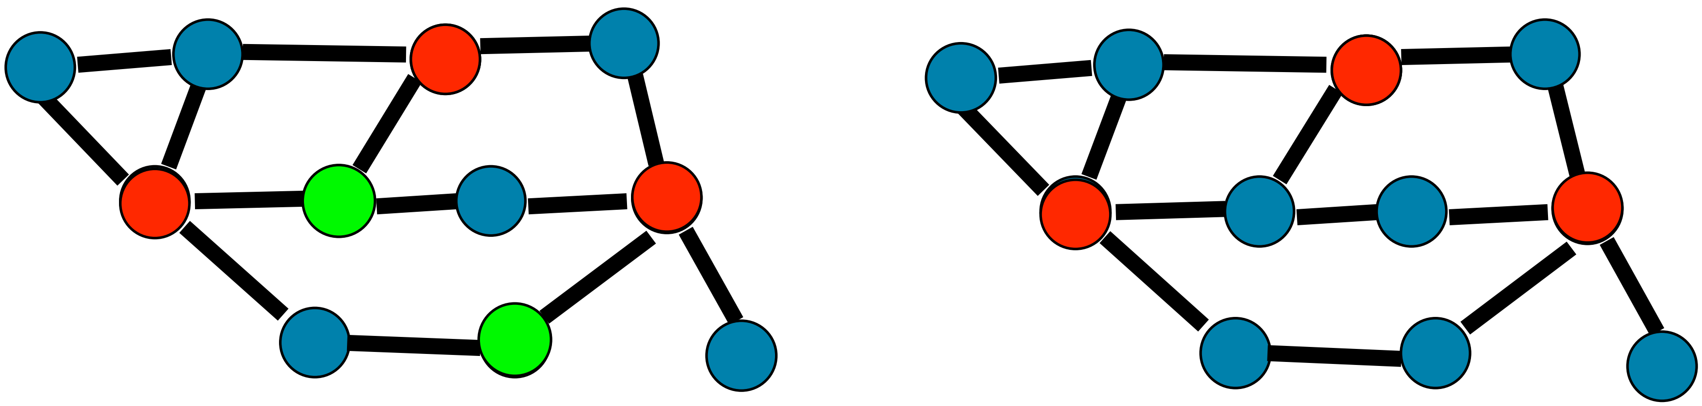

- *when strong F-F connection without a common C-point is detected, do not add C-point, but extend interpolation stencil to distance-two C-points*
- *no C-points added, but larger interpolation stencils*

results using long-range interpolation

- *3D elliptic PDE with jumps in coefficient a*

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

- *1 processor, AMG+GMRES, 80^3 dof*

	t_{tot}	C_{op}	s_{avg} (level)	Iter
CLJP	48.0	21.54	1007 (9)	7
PMIS	94.6	2.46	54 (3)	188
PMIS + mp	13.7	2.47	56 (3)	21
PMIS + F-F	21.4	4.90	204 (3)	9

results using long-range interpolation

- *3D elliptic PDE with jumps in coefficient a*
- *1 processor, AMG+GMRES*

	dof	C_{op}	s_{avg}	Iter	t_{setup}	t_{solve}	t_{tot}
PMIS + mp	***40***	***2.53***	***44***	***17***	***0.33***	***0.98***	***1.31***
	80	***2.47***	***56***	***21***	***3.11***	***10.55***	***13.66***
	120	***2.44***	***59***	***26***	***10.98***	***46.84***	***57.82***
PMIS + F-F	**40**	**4.64**	**114**	**9**	**1.31**	**0.70**	**2.01**
	80	**4.90**	**204**	**9**	**15.06**	**6.38**	**21.44**
	120	**4.94**	**248**	**9**	**55.47**	**22.94**	**78.41**

- *mp uses less memory, is faster than F-F*

Conclusions

- *PMIS leads to reduced, scalable complexities for large problems on parallel computers*

- *for difficult problems, nearest-neighbor interpolation is not sufficient on PMIS grids*

- *long-range interpolation improves convergence*

- *multipass appears superior to F-F*

Future work

- *parallel implementation of multipass interpolation*
- *investigate scalability of parallel AMG algorithms on Blue Gene/L-class machines*