05/09	Class 9	Dense direct solvers	Understand the principle of LU decomposition and the optimization and parallelization techniques that lead to the LINPACK benchmark.
05/12	Class 10	Dense eigensolvers	Determine eigenvalues and eigenvectors and understand the fast algorithms for diagonalization and orthonormalization.
05/16	Class 11	Sparse direct solvers	Understand reordering in AMD and nested dissection, and fast algorithms such as skyline and multifrontal methods.
05/19	Class 12	Sparse iterative solvers	Understand the notion of positive definiteness, condition number, and the difference between Jacobi, CG, and GMRES.
05/23	Class 13	Preconditioners	Understand how preconditioning affects the condition number and spectral radius, and how that affects the CG method.
05/26	Class 14	Multigrid methods	Understand the role of smoothers, restriction, and prolongation in the V-cycle.
05/30	Class 15	Fast multipole methods, H-matrices	Understand the concept of multipole expansion and low-rank approximation, and the role of the tree structure.

Iterative Methods

- **Initial error.**

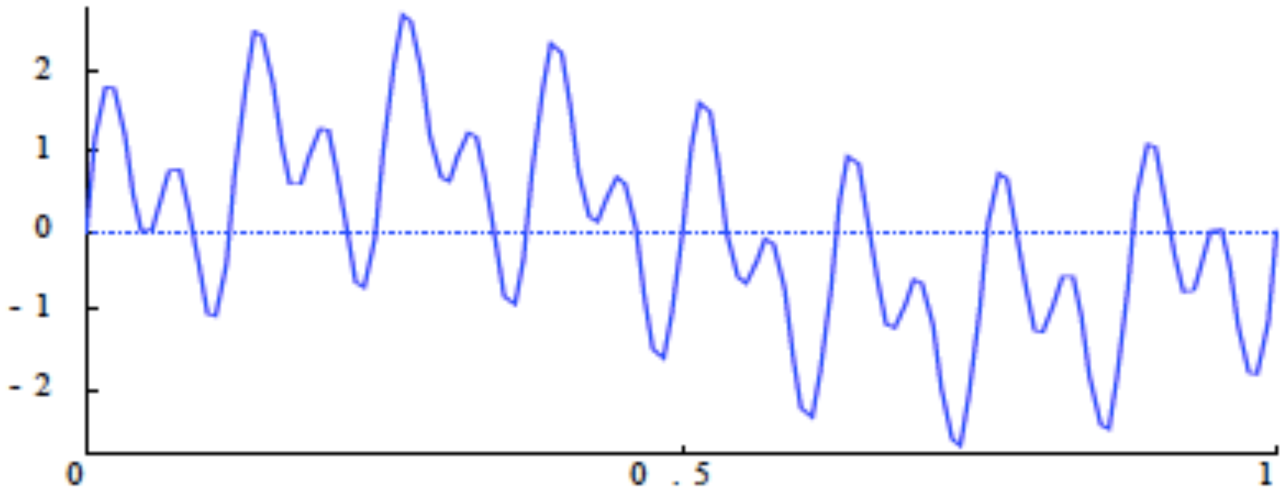

- **Error after several iteration sweeps:**

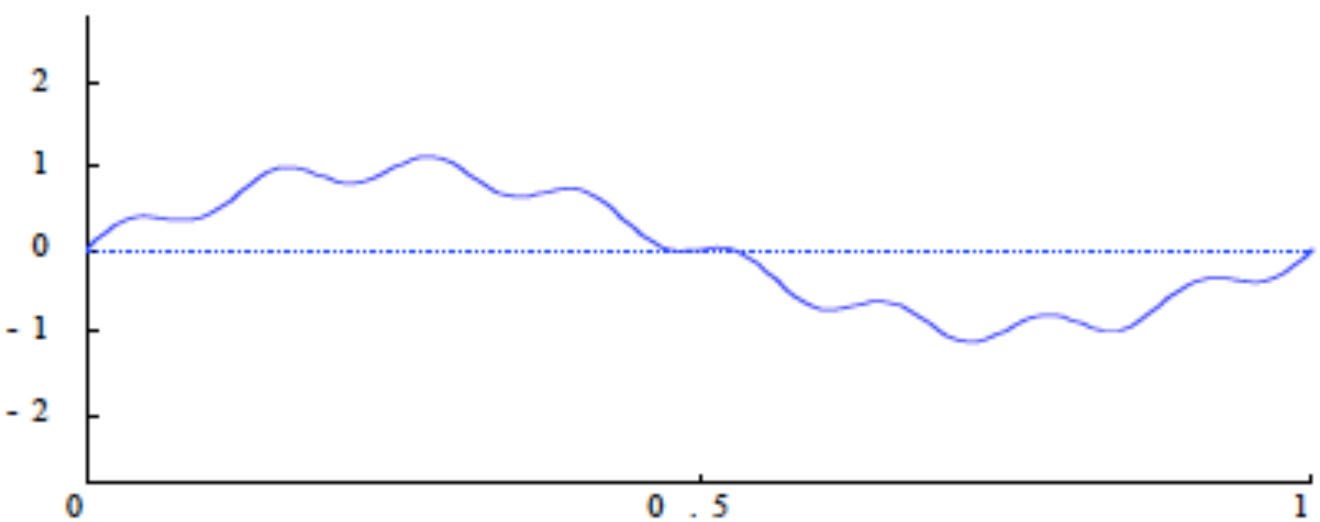

Many relaxation schemes have the smoothing property, where oscillatory modes of the error are eliminated effectively, but smooth modes are damped very slowly.

Multigrid Methods

- A smooth function:

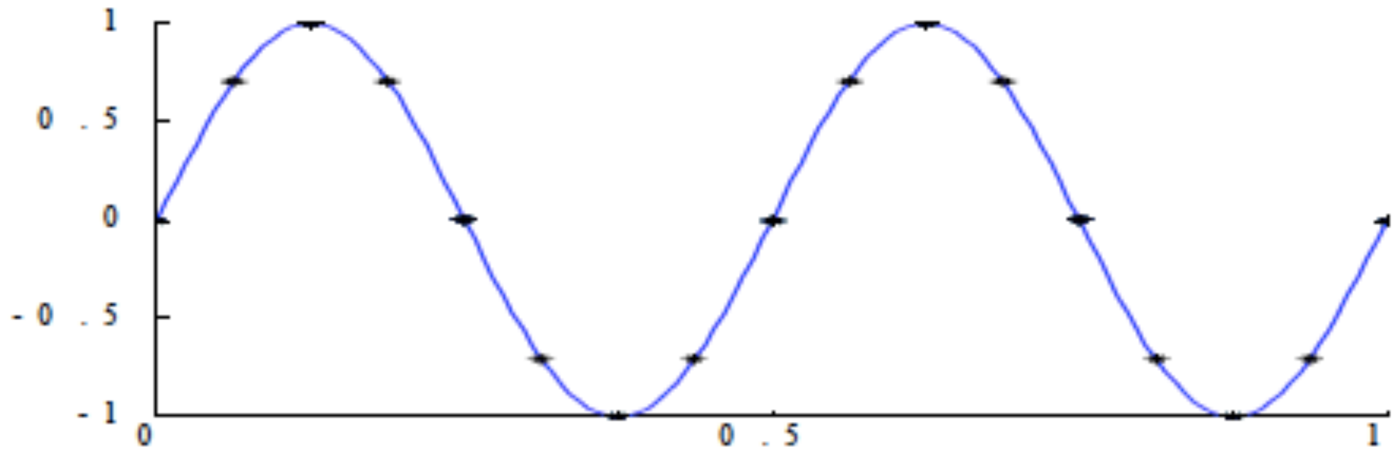

- Can be represented by linear interpolation from a coarser grid:

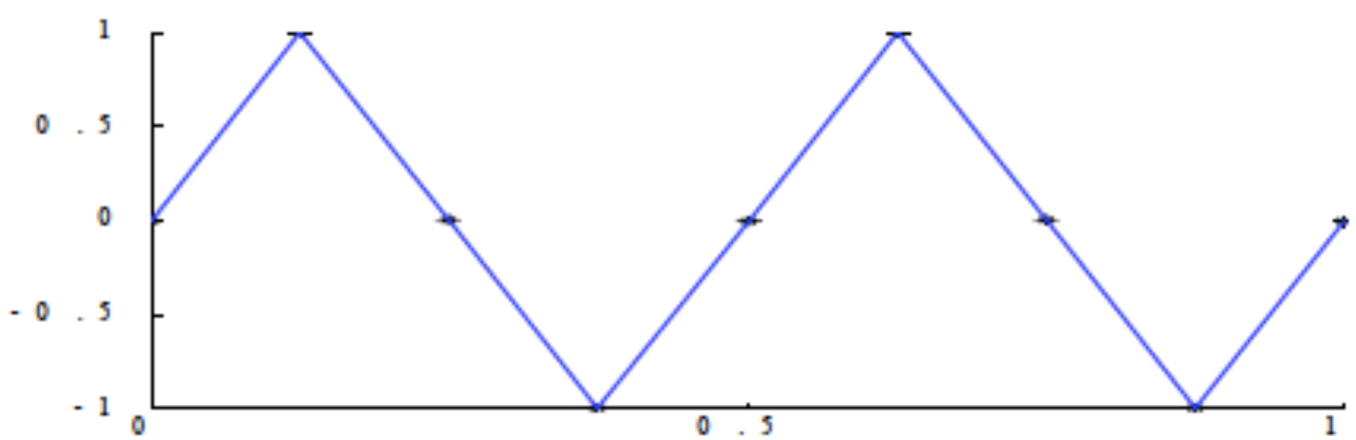

On the coarse grid, the smooth error appears to be relatively higher in frequency: in the example it is the 4-mode, out of a possible 16, on the fine grid, 1/4 the way up the spectrum. On the coarse grid, it is the 4-mode out of a possible 8, hence it is 1/2 the way up the spectrum.

Relaxation will be more effective on this mode if done on the coarser grid!!

Multigrid Method

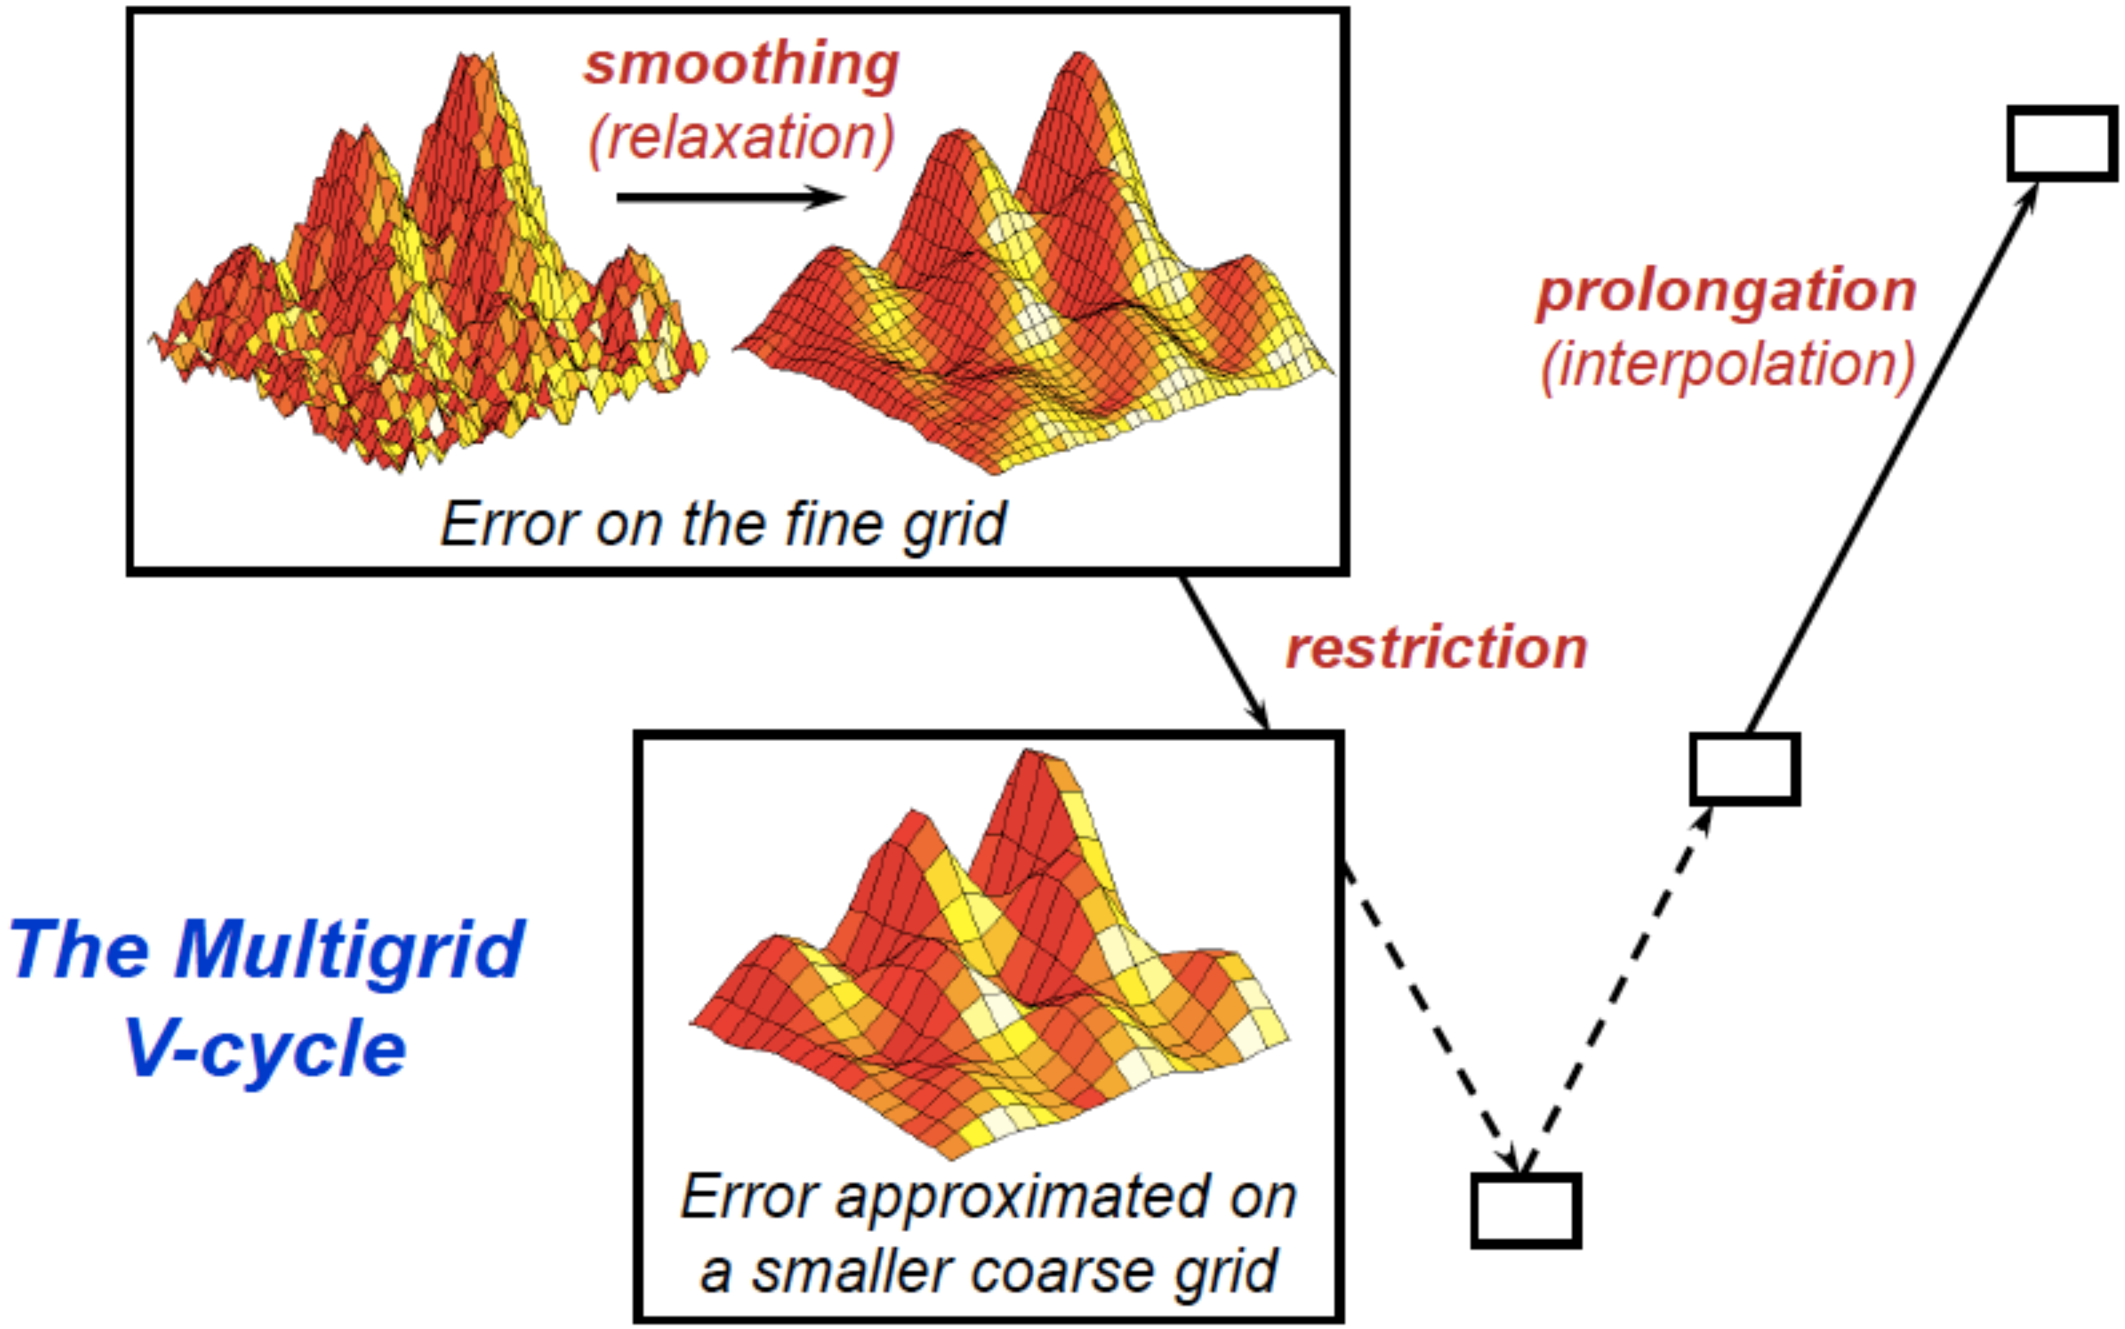

Multigrid Methods

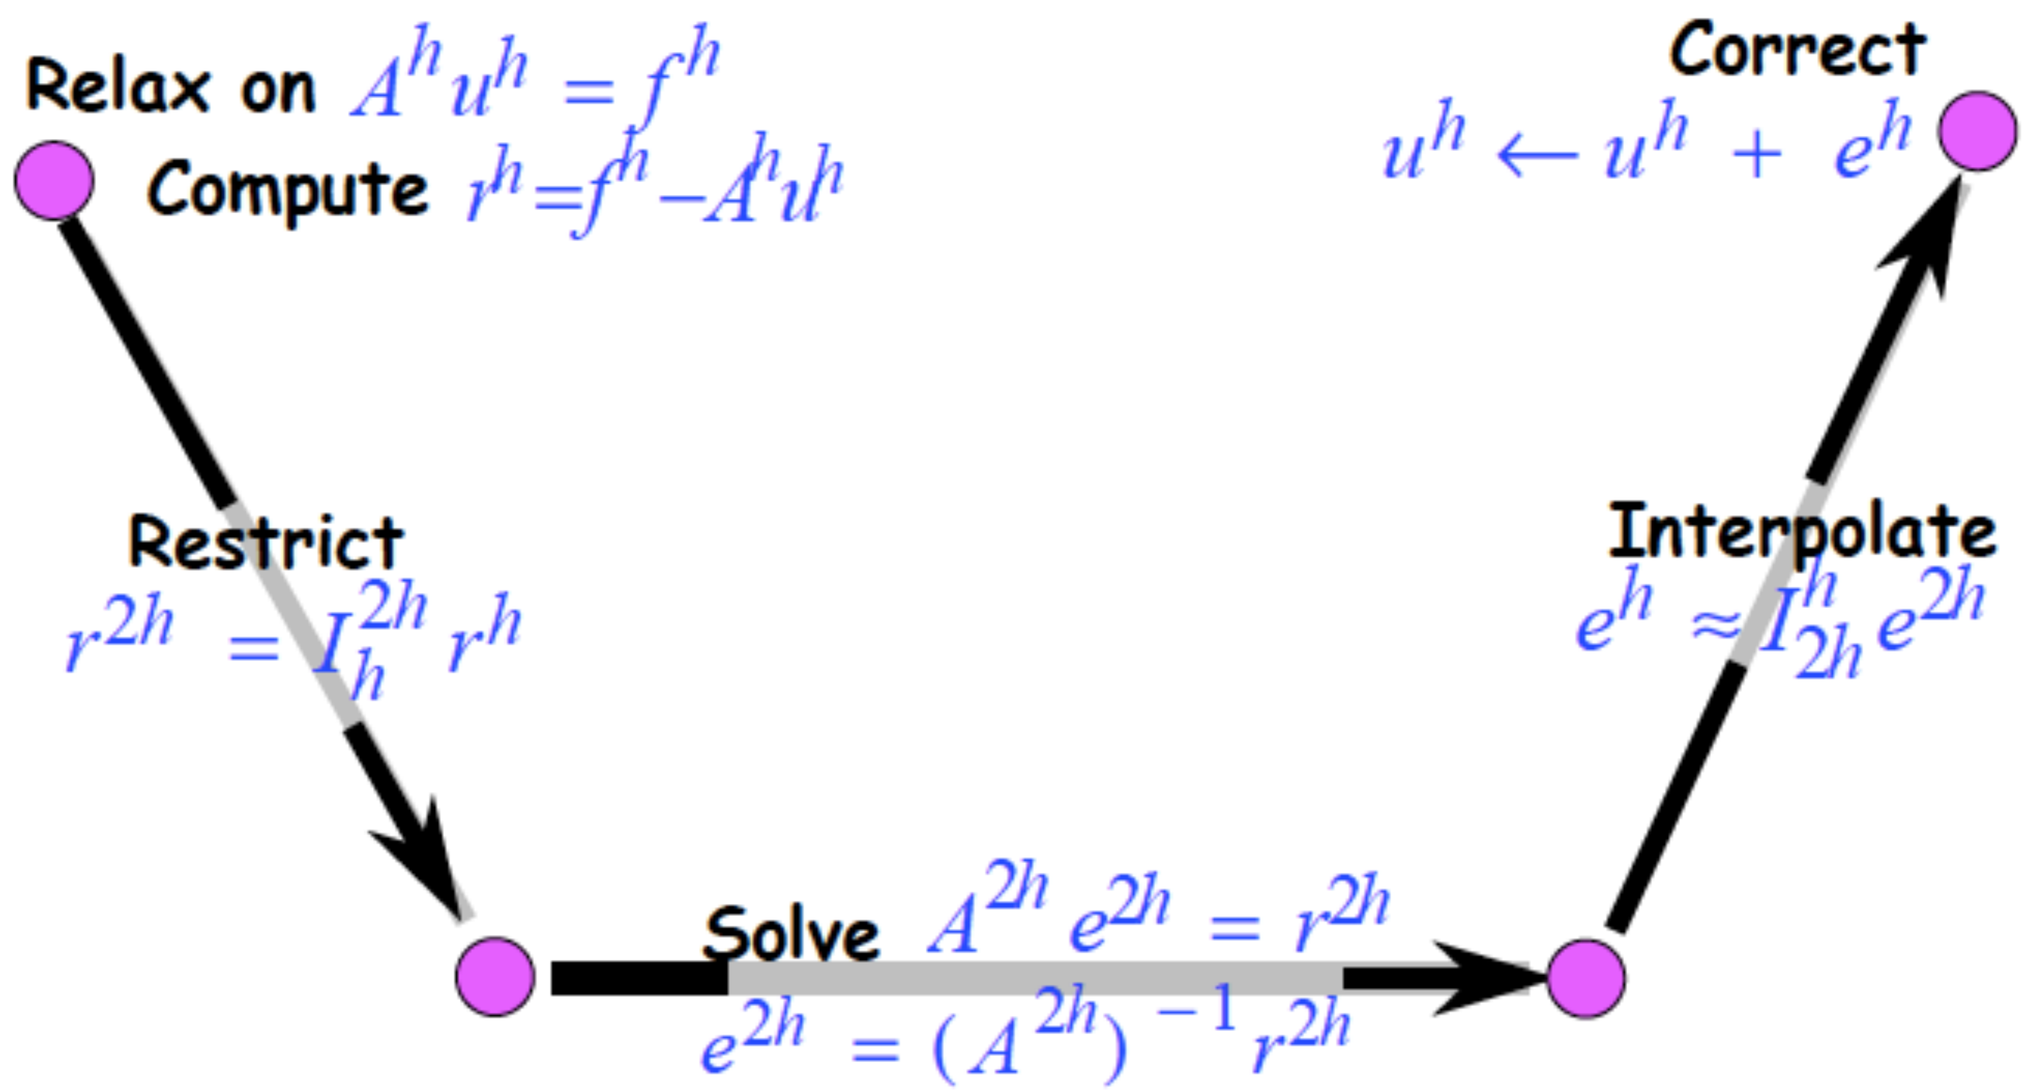

1-D Laplace Problem

- 1D Laplace on a uniform grid with spacing h

$$-u_{xx} = f \quad \text{on } \Omega = [0,1]$$
$$u = g \quad \text{on } \Gamma$$

Continuous

$$-u_{i-1} + 2u_i - u_{i+1} = h^2 f_i$$
$$u_0 = u_{N+1} = g$$

Discrete

- Discrete problem is a linear system $Au = f$ with

$$A = \begin{pmatrix} 2 & -1 & & \\ -1 & 2 & -1 & \\ & & \ddots & \\ & & -1 & 2 \end{pmatrix}$$

Matrix

or

$$A = \begin{bmatrix} -1 & 2 & -1 \end{bmatrix}$$

Stencil

1-D Laplace Problem

- 1D Laplace on a uniform grid with spacing h

$$-u_{xx} = f \quad \text{on } \Omega = [0,1]$$
$$u = g \quad \text{on } \Gamma$$

Continuous

$$-u_{i-1} + 2u_i - u_{i+1} = h^2 f_i$$
$$u_0 = u_{N+1} = g$$

Discrete

- Discrete problem is a linear system $Au = f$ with

$$A = \begin{pmatrix} 2 & -1 & & \\ -1 & 2 & -1 & \\ & & \ddots & \\ & & -1 & 2 \end{pmatrix}$$

Matrix

or

$$A = \begin{bmatrix} -1 & 2 & -1 \end{bmatrix}$$

Stencil

Restriction

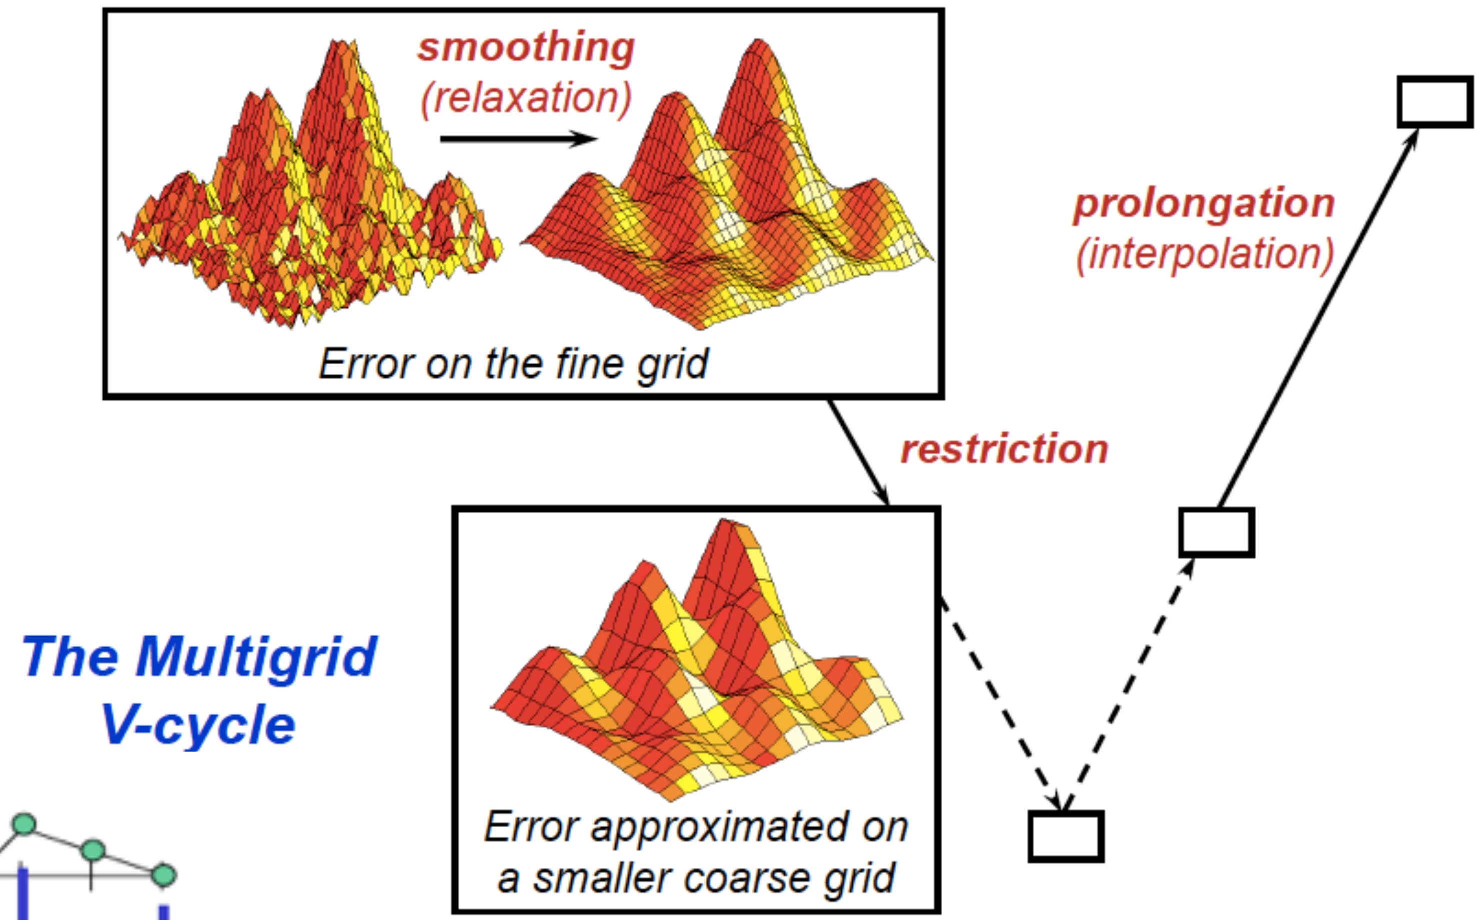

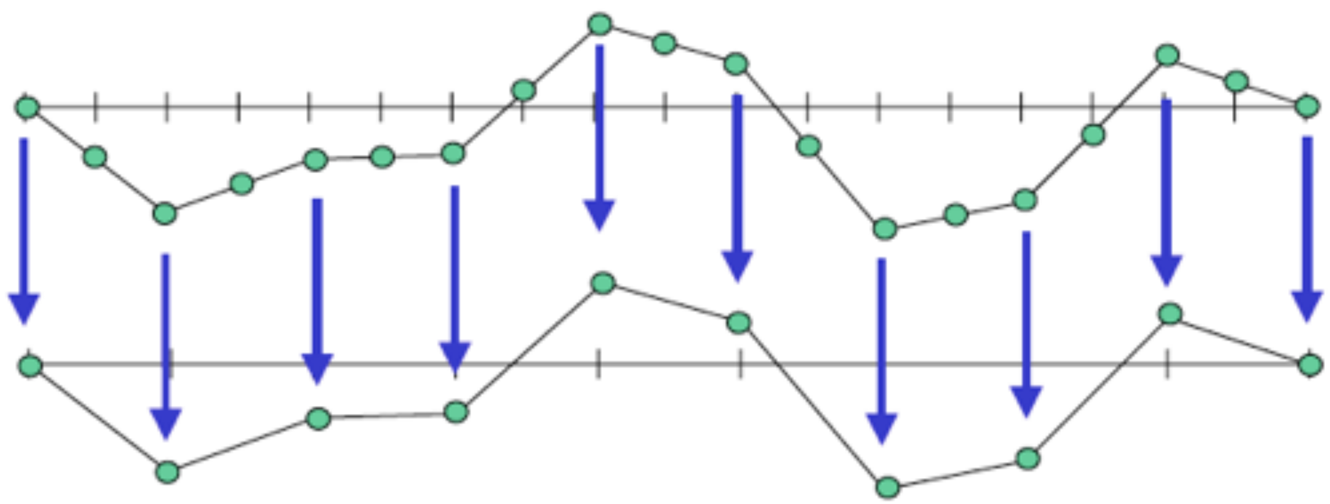

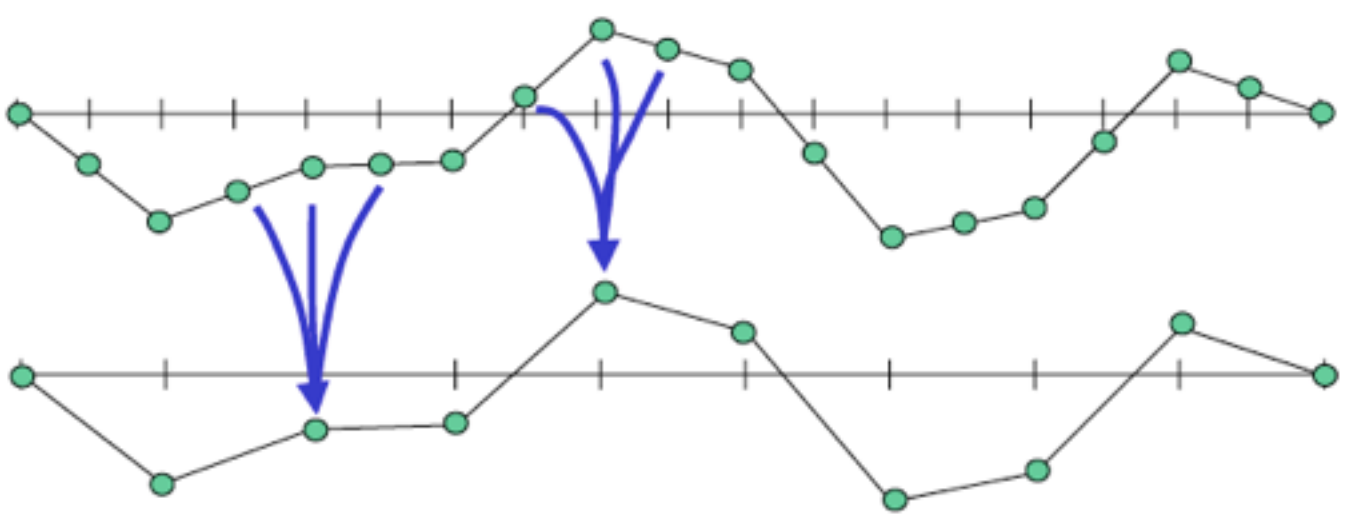

Interpolation (Prolongation)

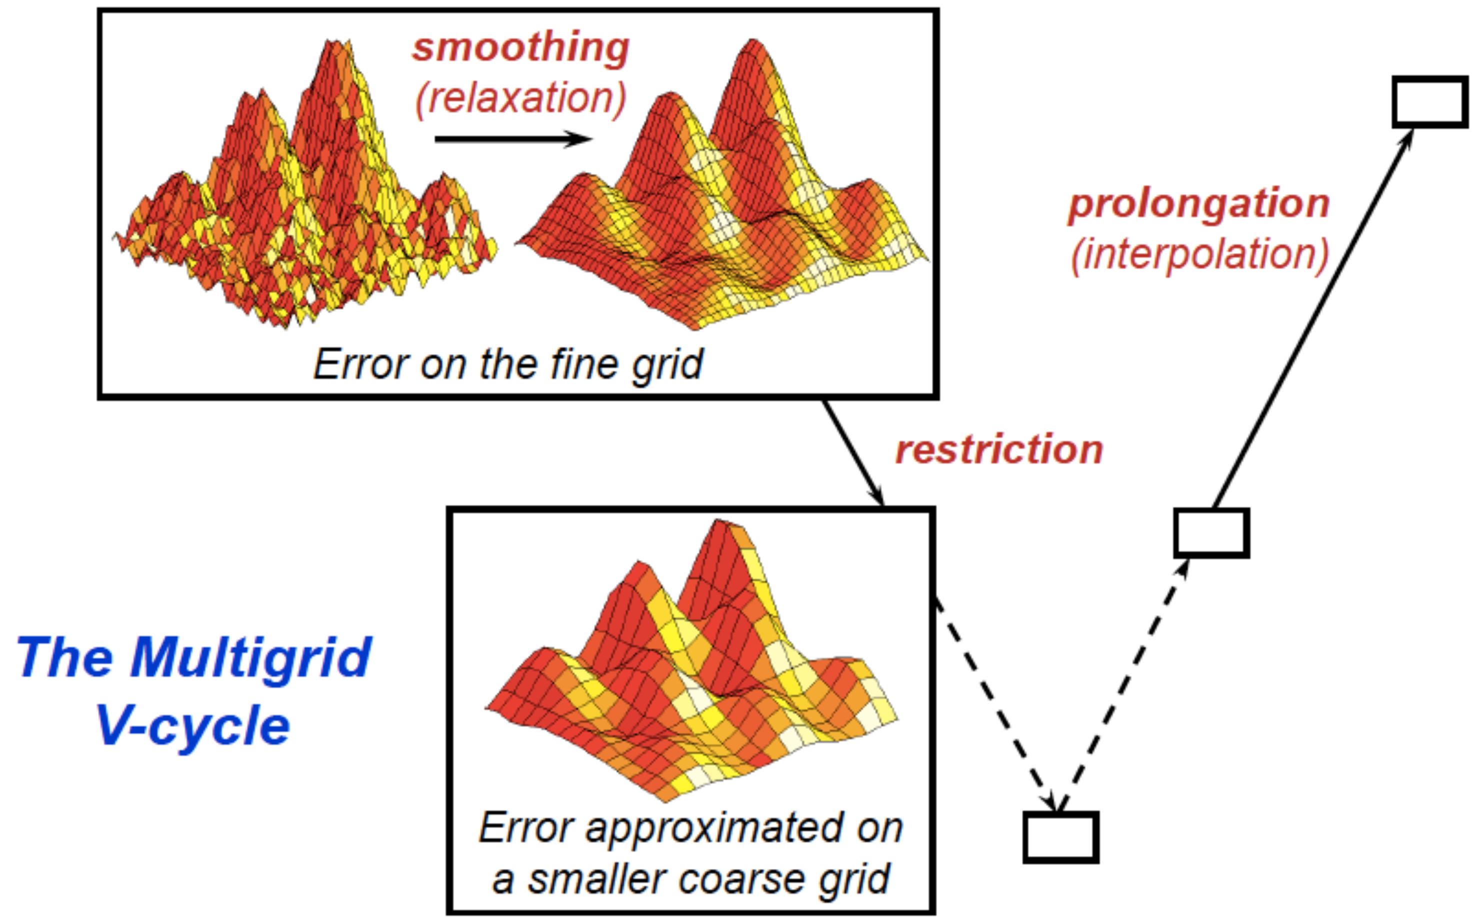

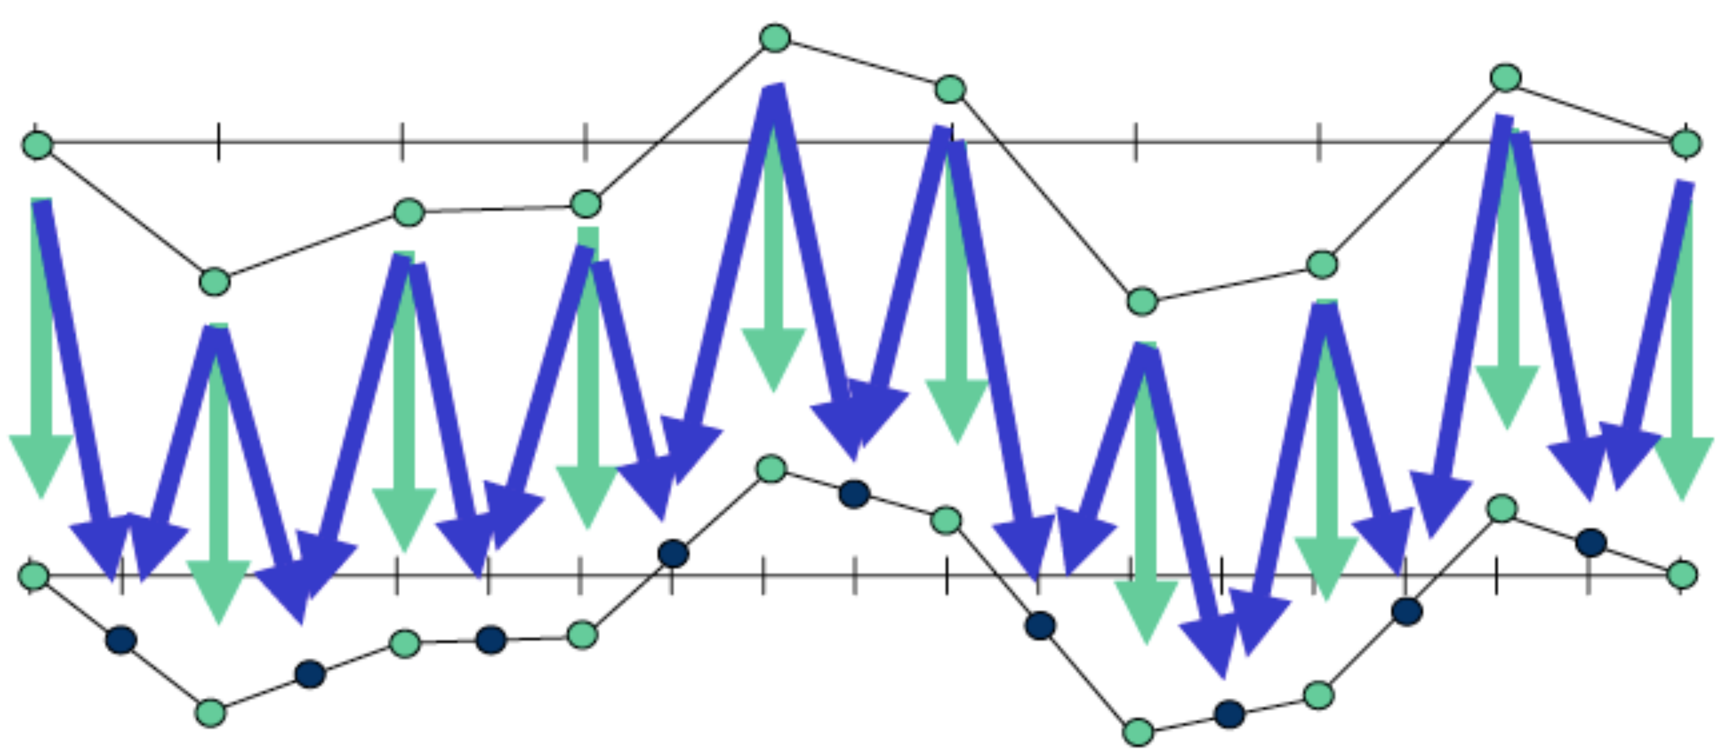

Interpolation (Prolongation)

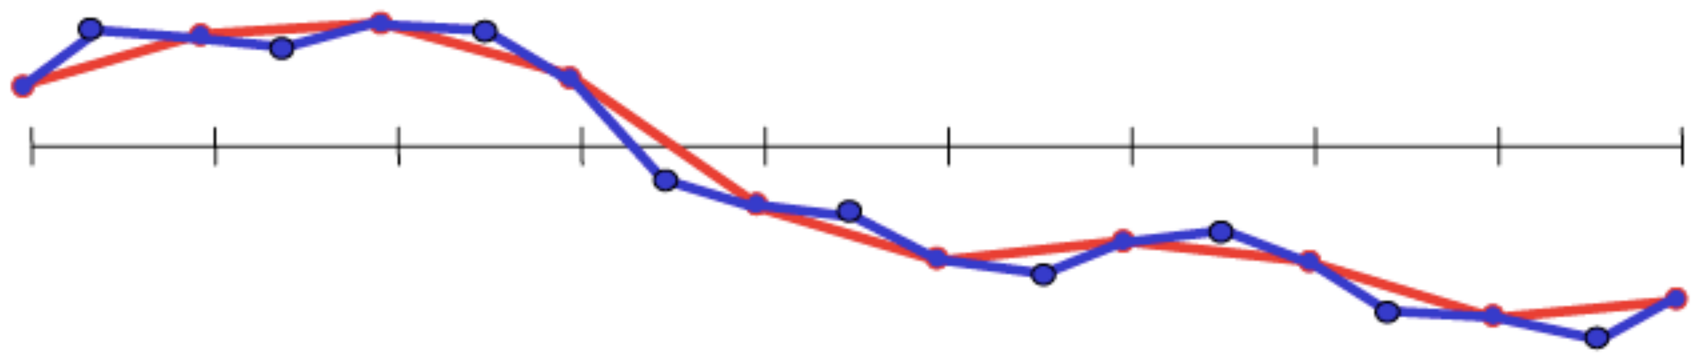

- If u is smooth, a coarse-grid interpolant of v^{2h} may do very well.

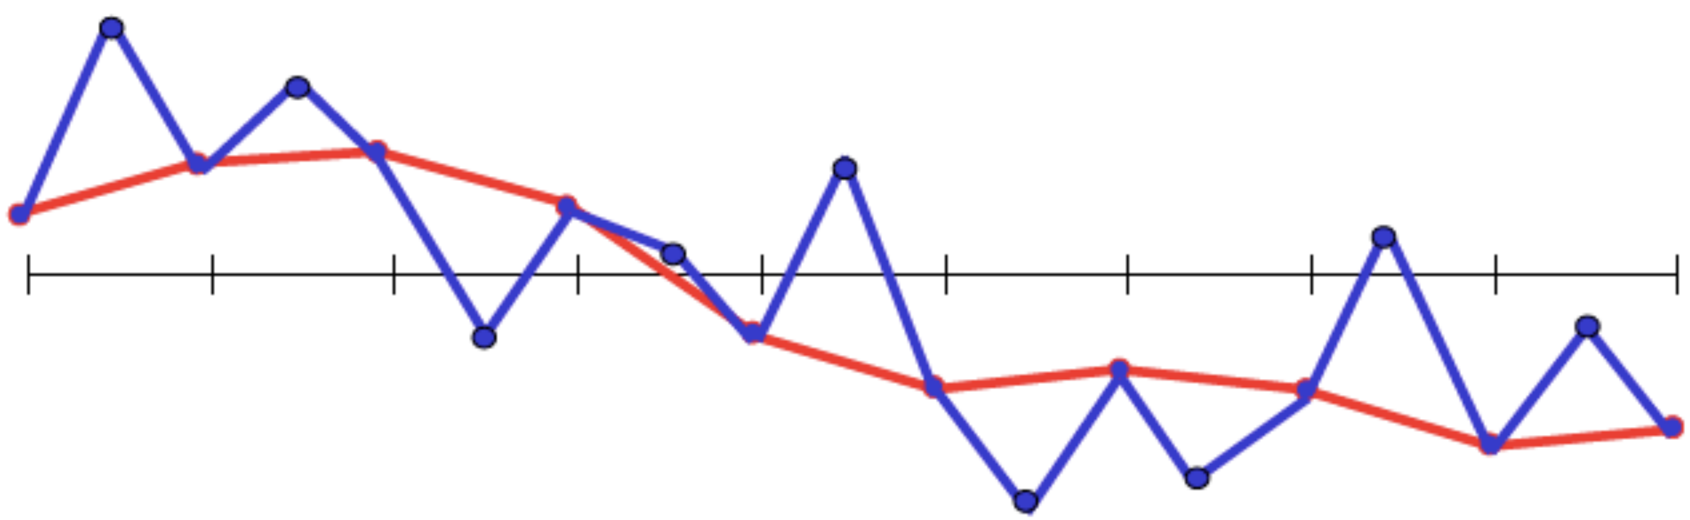

- If u is oscillatory, a coarse-grid interpolant of v^{2h} may ***not*** work well.

Smoothers

A classical linear iteration for a matrix $\mathtt{A}$ of the form

$$\mathtt{x}^{(\mathtt{i}+1)} = \mathtt{x}^{(\mathtt{i})} + \mathtt{R}(\mathtt{b} - \mathtt{A}\mathtt{x}^{(\mathtt{i})})$$

with some matrix $\mathtt{R}$, is a *smoothing iteration* if

$$(\mathtt{I} - \mathtt{RA})\mathtt{e} \quad \text{is smoother than e for any e.}$$

$$\begin{array}{rcccc} \mathtt{x}^{(\mathtt{i}+1)} & = & \mathtt{x}^{(\mathtt{i})} & + & \mathtt{R}(\mathtt{b} - \mathtt{A}\mathtt{x}^{(\mathtt{i})}) \\ \mathtt{x} & = & \mathtt{x} & + & \mathtt{R}(\mathtt{b} - \mathtt{A}\mathtt{x}) \\ \hline \mathtt{e}^{(\mathtt{i}+1)} & = & \mathtt{e}^{(\mathtt{i})} & - & \mathtt{R}\mathtt{A}\mathtt{e}^{(\mathtt{i})}. \end{array}$$

(Hence smoothing iterations smooth *errors*.)

Smoothers

A classical linear iteration for a matrix A of the form

$$x^{(i+1)} = x^{(i)} + R(b - Ax^{(i)})$$

with some matrix R, is a *smoothing iteration* if

$$(I - RA)e \quad \text{is smoother than e for any e.}$$

If D is the diagonal and L is the lower triangular part of A, then

Jacobi iteration: $R = D^{-1}$

Gauß-Seidel iteration: $R = (L + D)^{-1}$

Multilevel V-cycle

- **Major question: How do we "solve" the coarse-grid residual equation? Answer: *recursion*!**

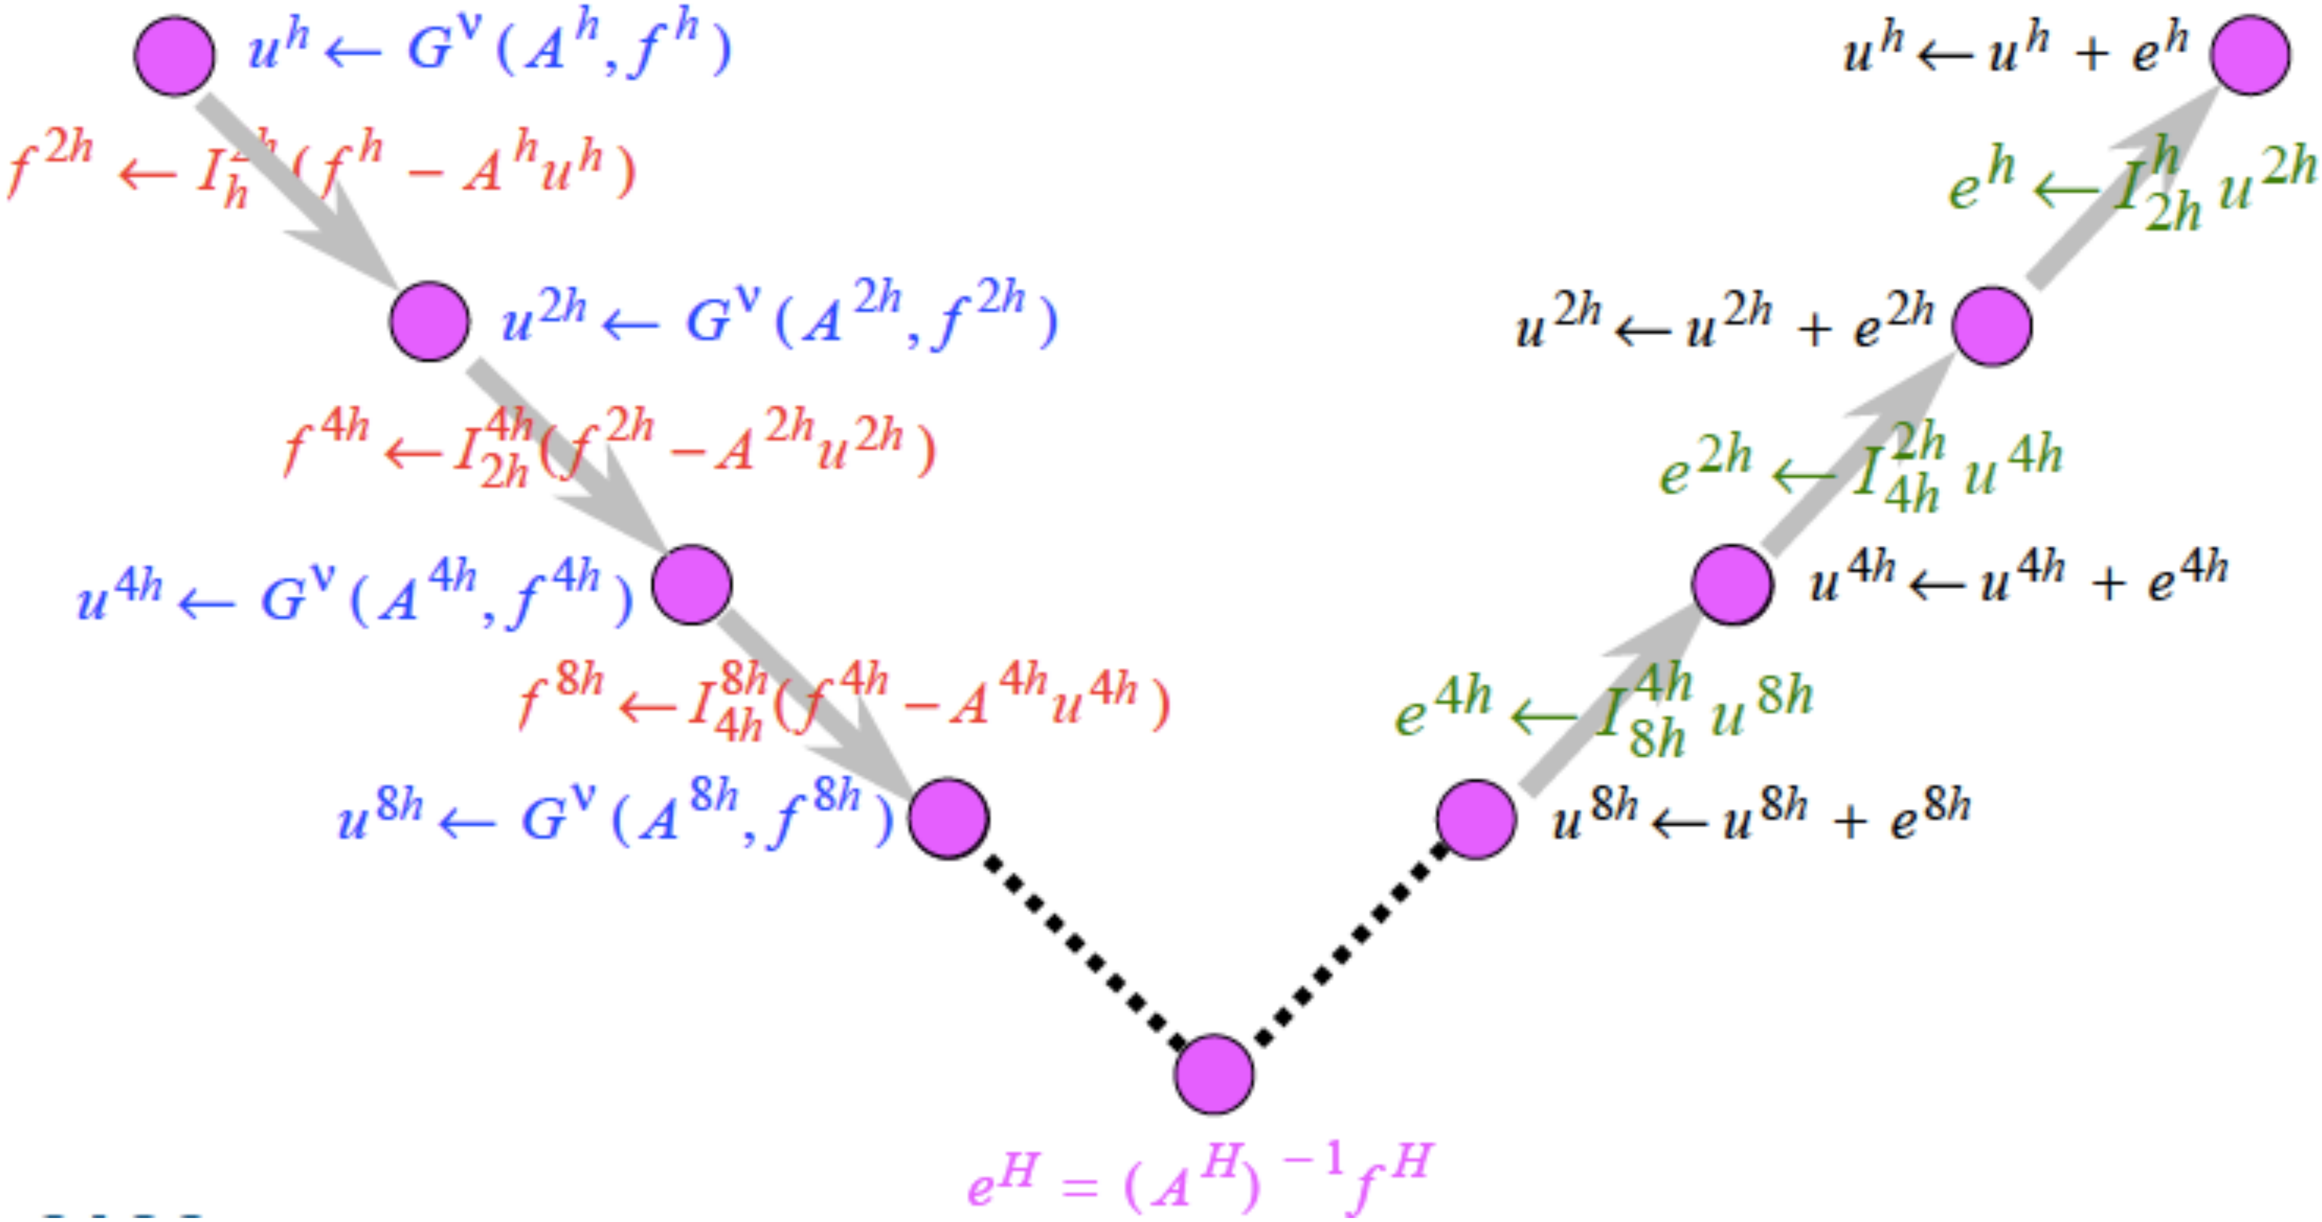

Multigrid cycles

Schedule of multilevel grids in standard multigrid algorithms:

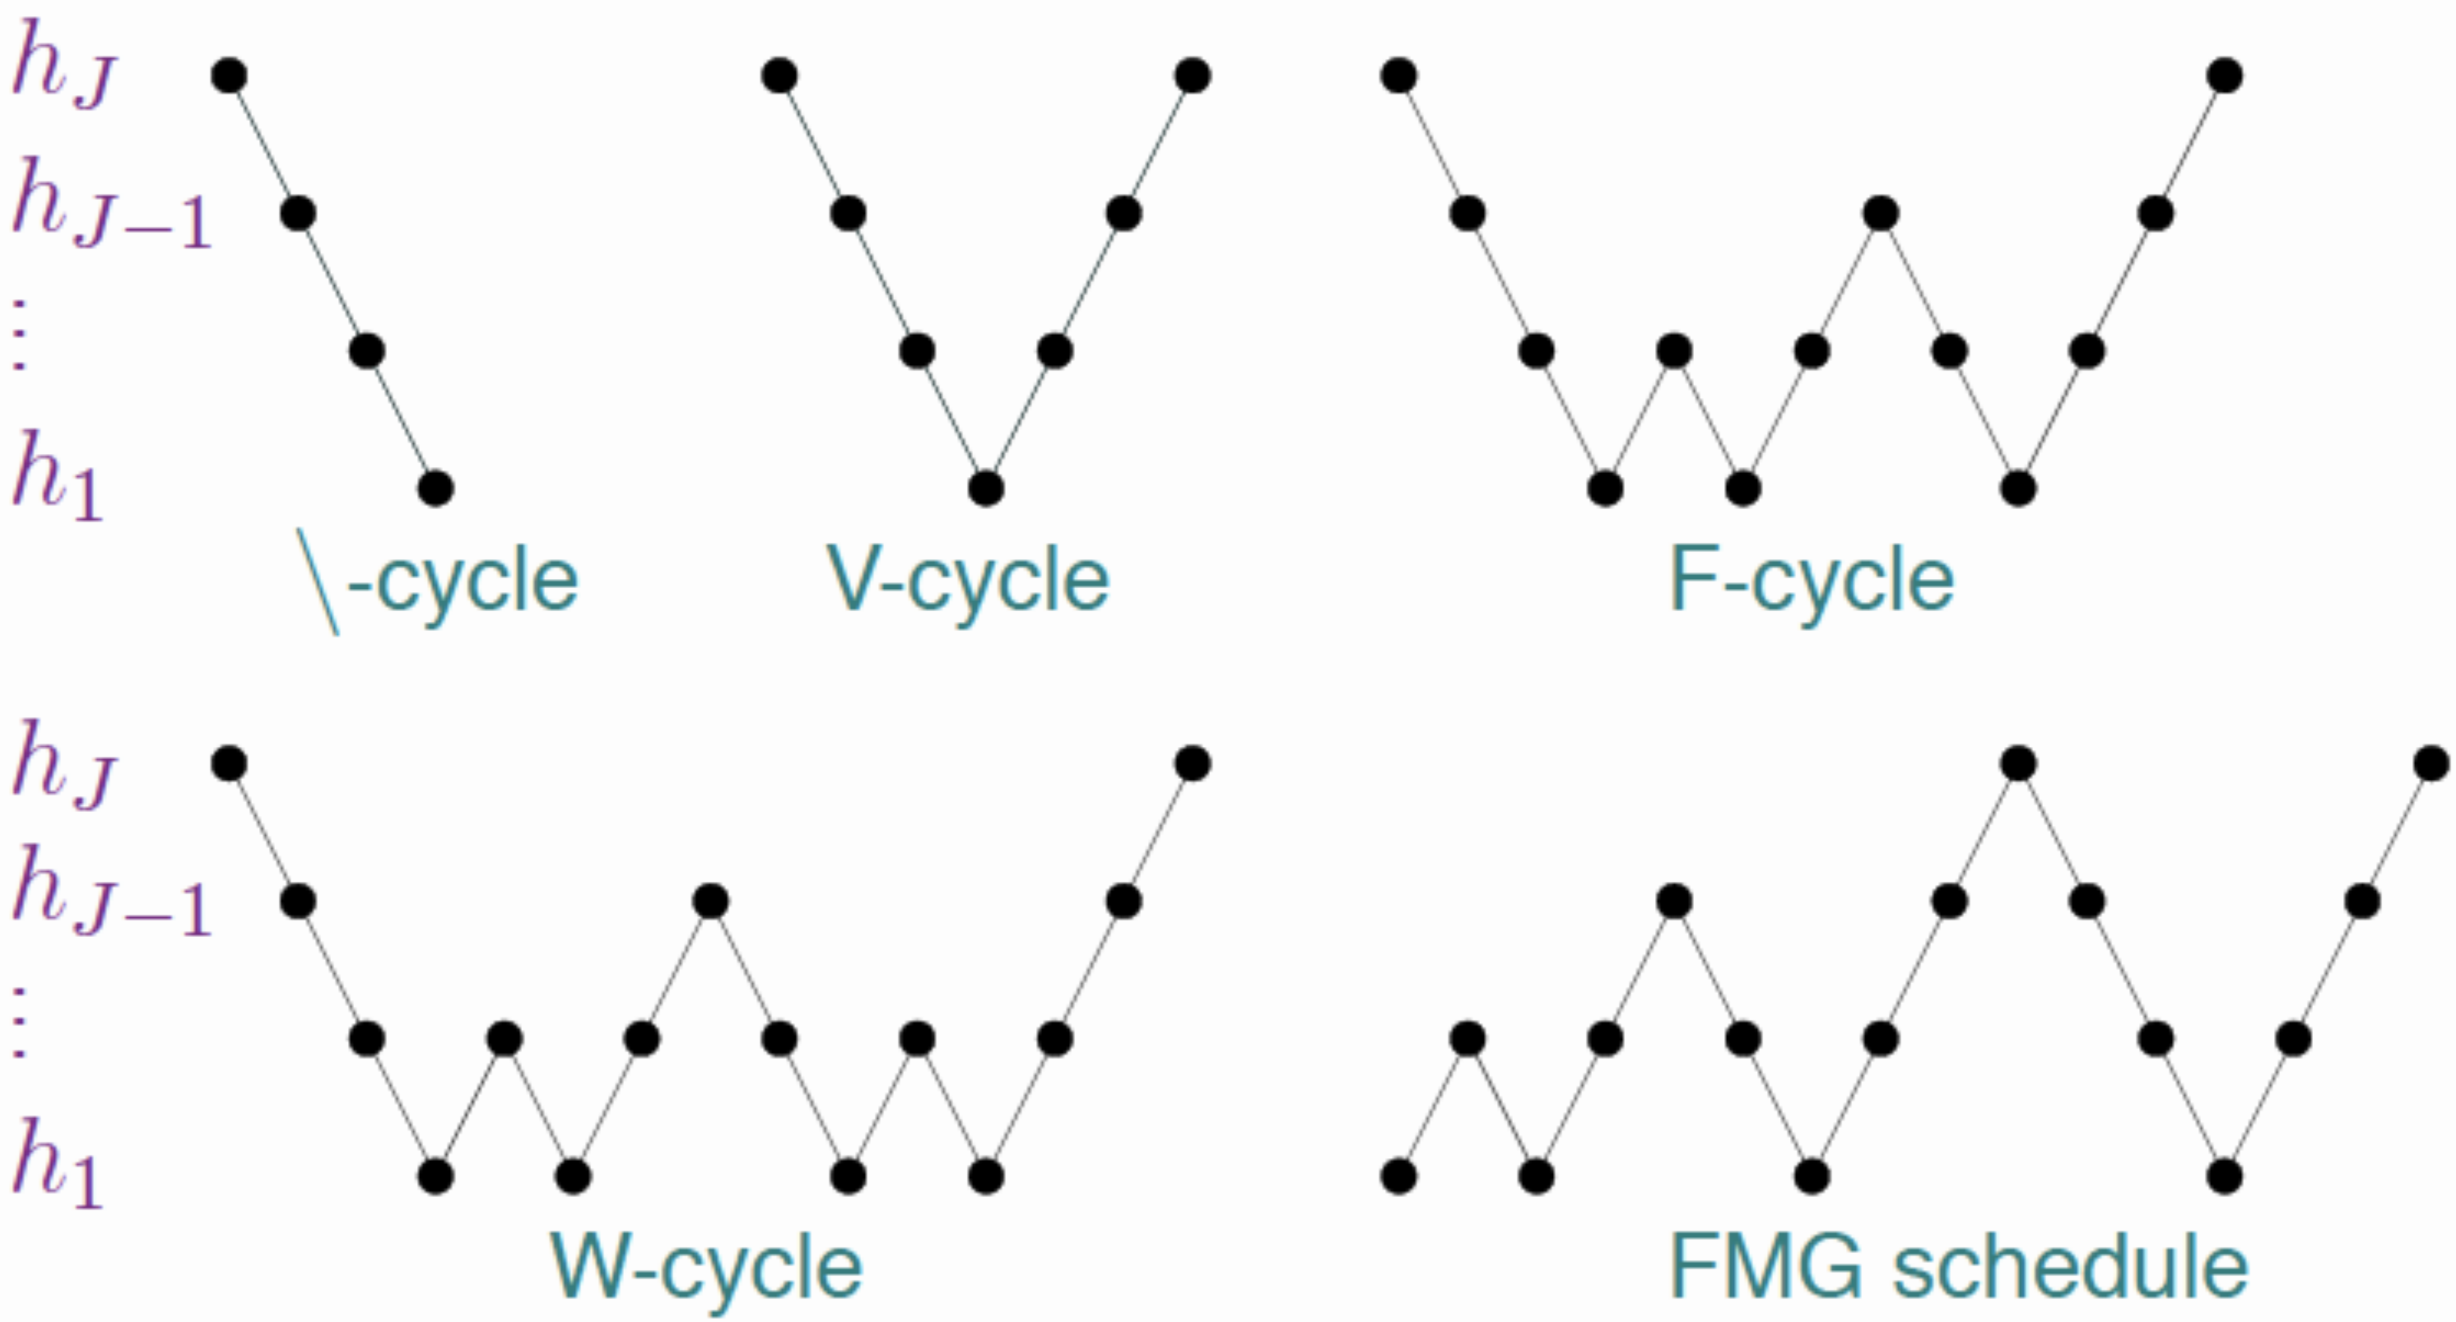

1-D Laplace Example

Testcase

Set $\Omega = (0, 1)$. Consider:

$$\begin{aligned} -\Delta u &= 0 \qquad \text{in } \Omega \\ u &= 0 \qquad \text{at } \partial\Omega = \{0, 1\} \end{aligned}$$

Set $N = 2^{10}$ and $u^0_{(\cdot)}$ according to:

$$u^0_i = \sin(2\pi x_i) + \sin(8\pi x_i) + \sin(32\pi x_i) + \sin(128\pi x_i)$$

Perform $\nu = 8$ GS relaxations followed by a coarse-grid correction.

1-D Laplace Example

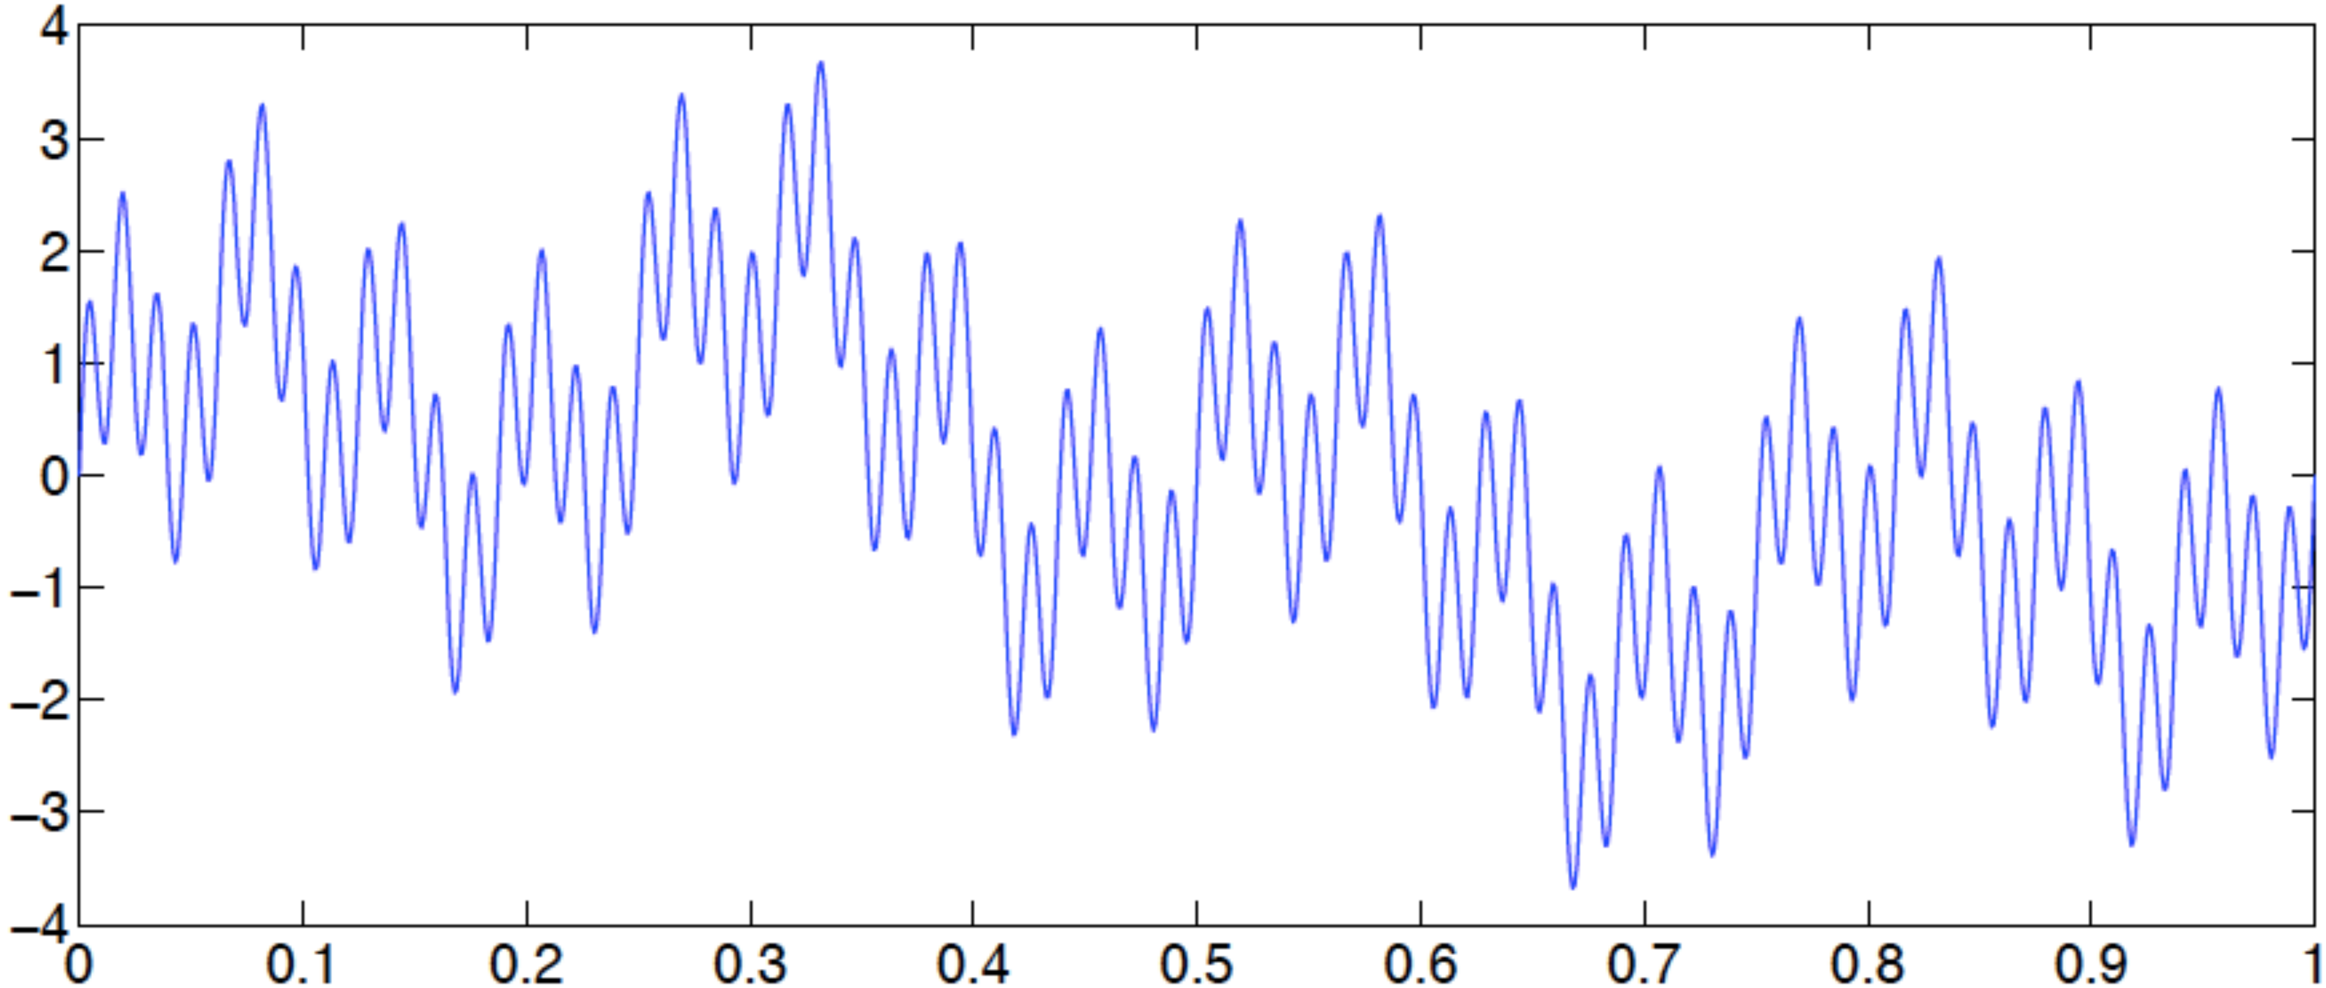

0 GS^8 relaxation   0 CG correction

1-D Laplace Example

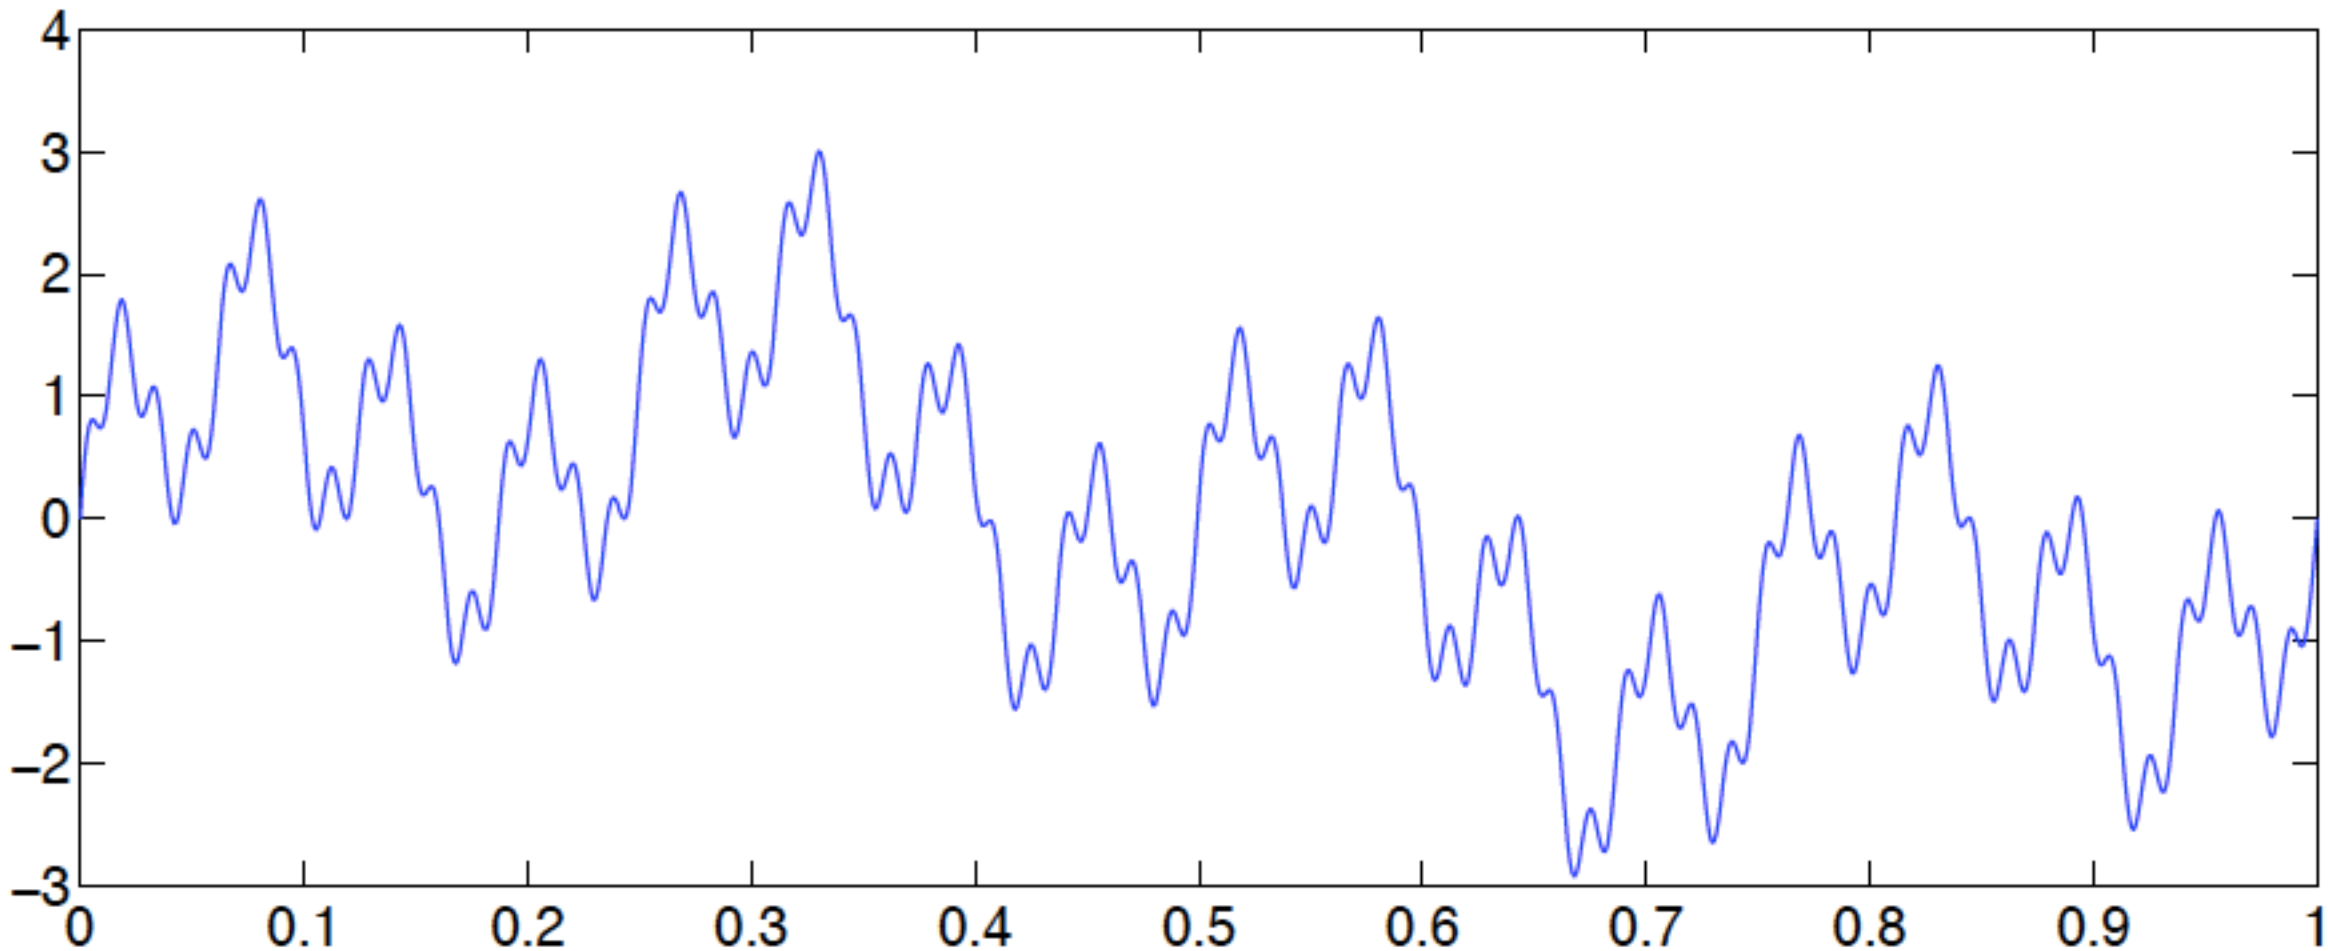

1 GS^8 relaxation   0 CG correction

1-D Laplace Example

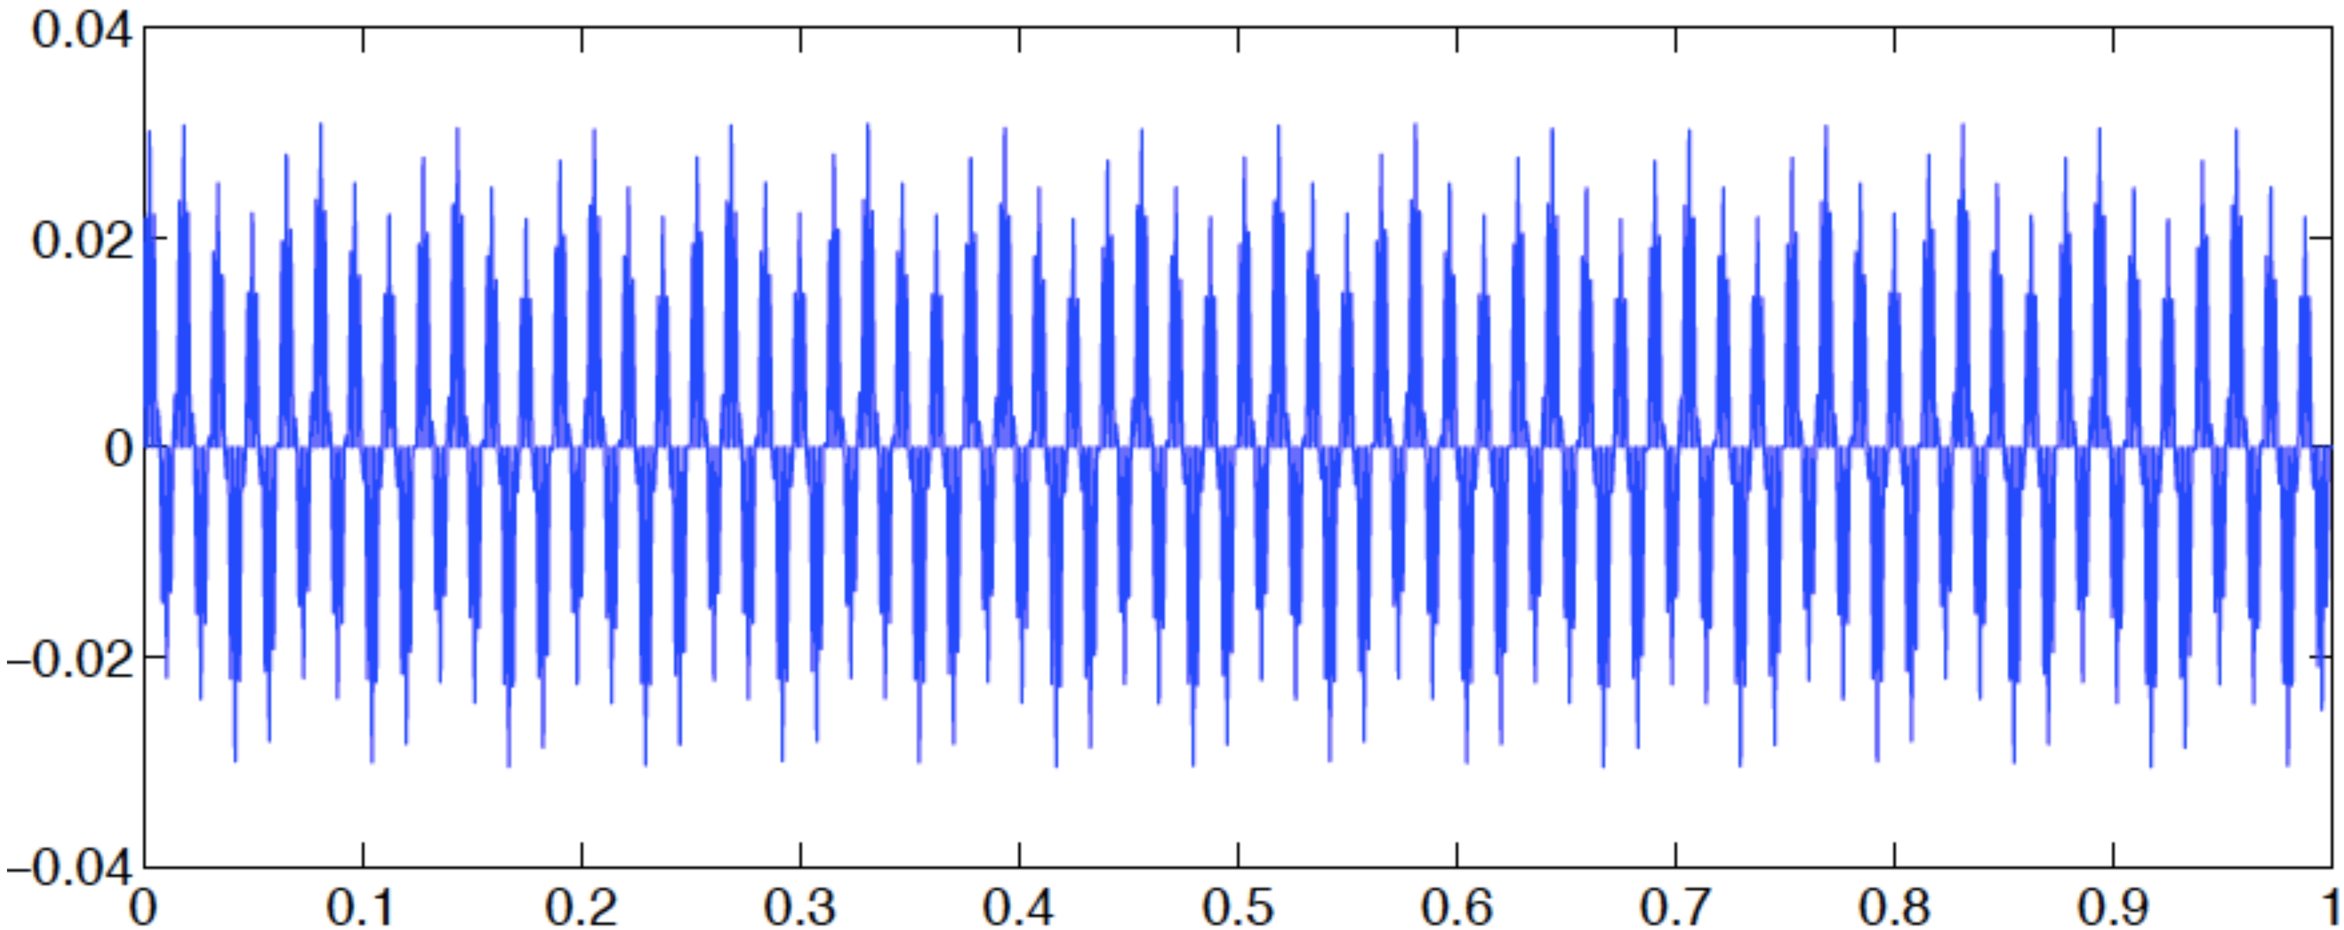

1 GS^8 relaxation 1 CG correction

1-D Laplace Example

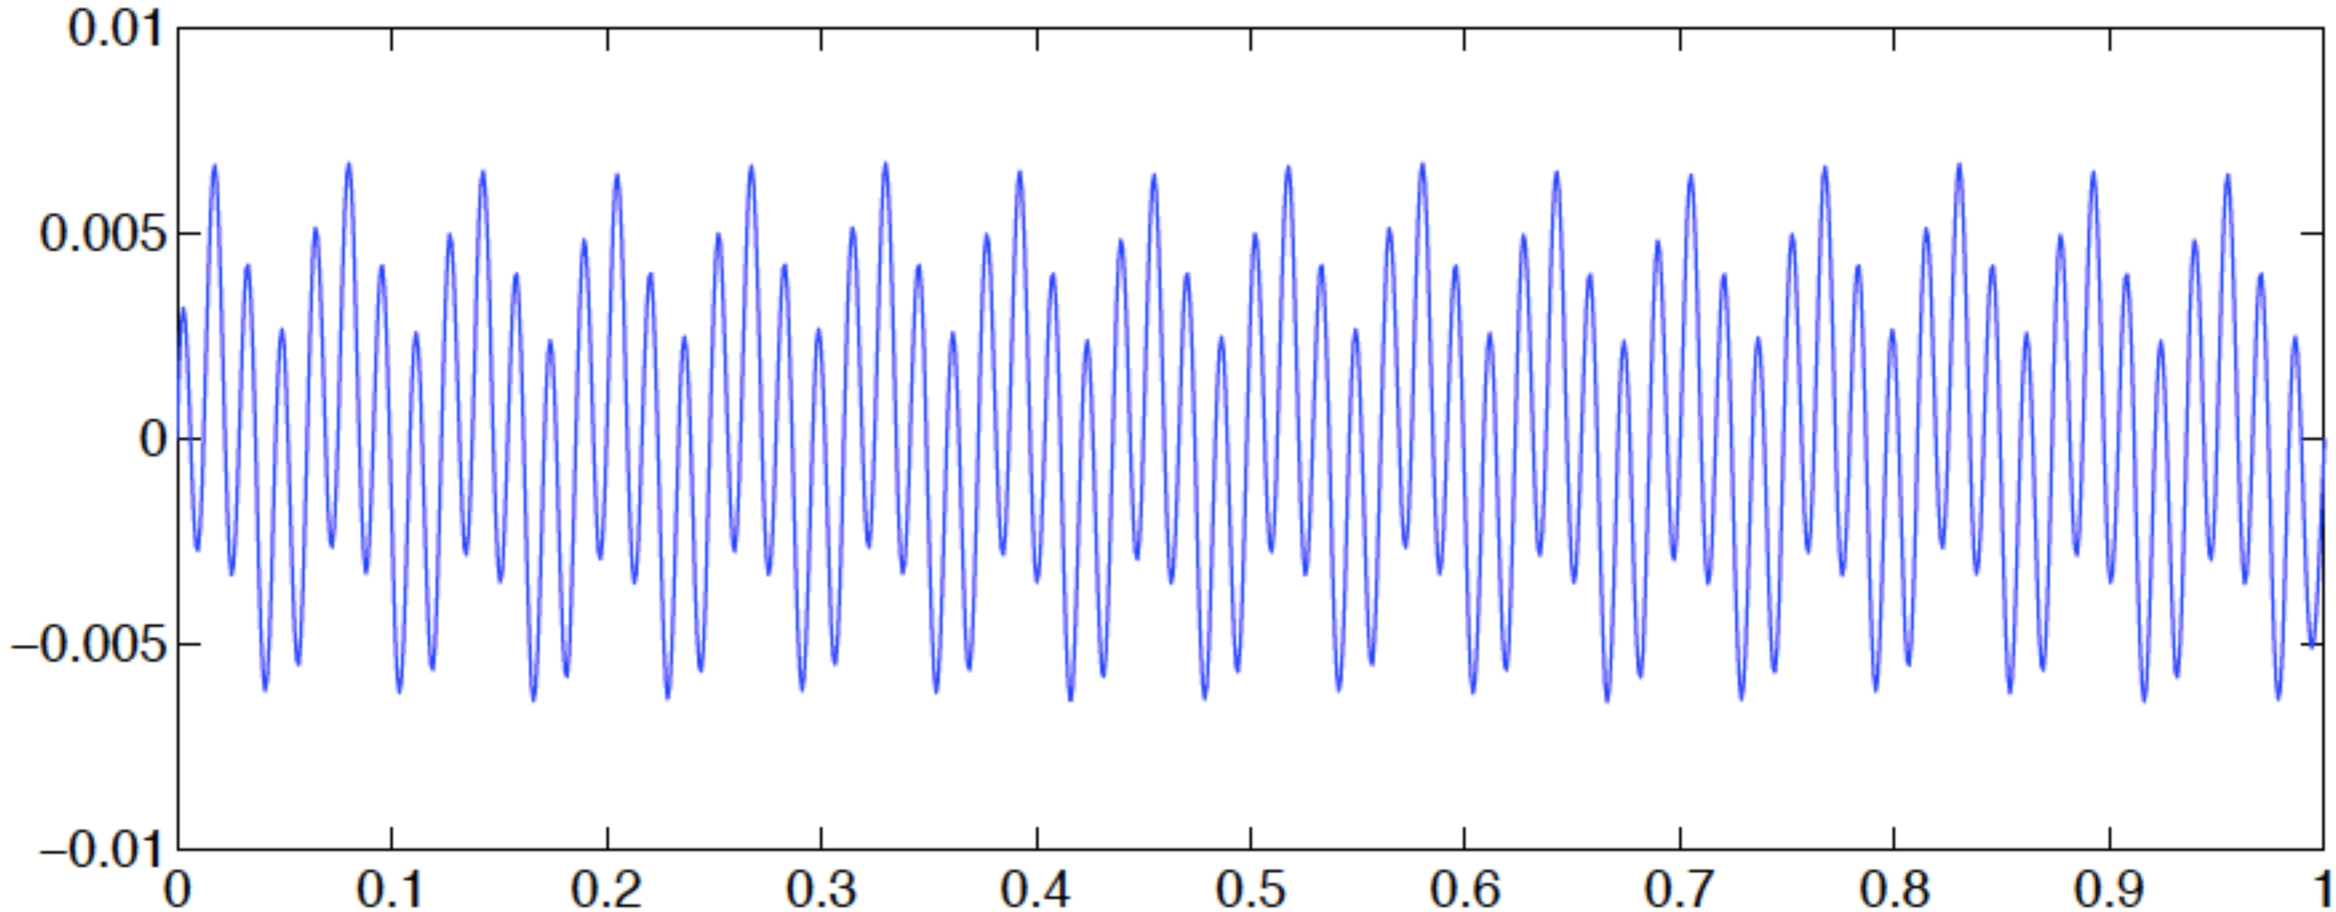

2 GS^8 relaxation 1 CG correction

1-D Laplace Example

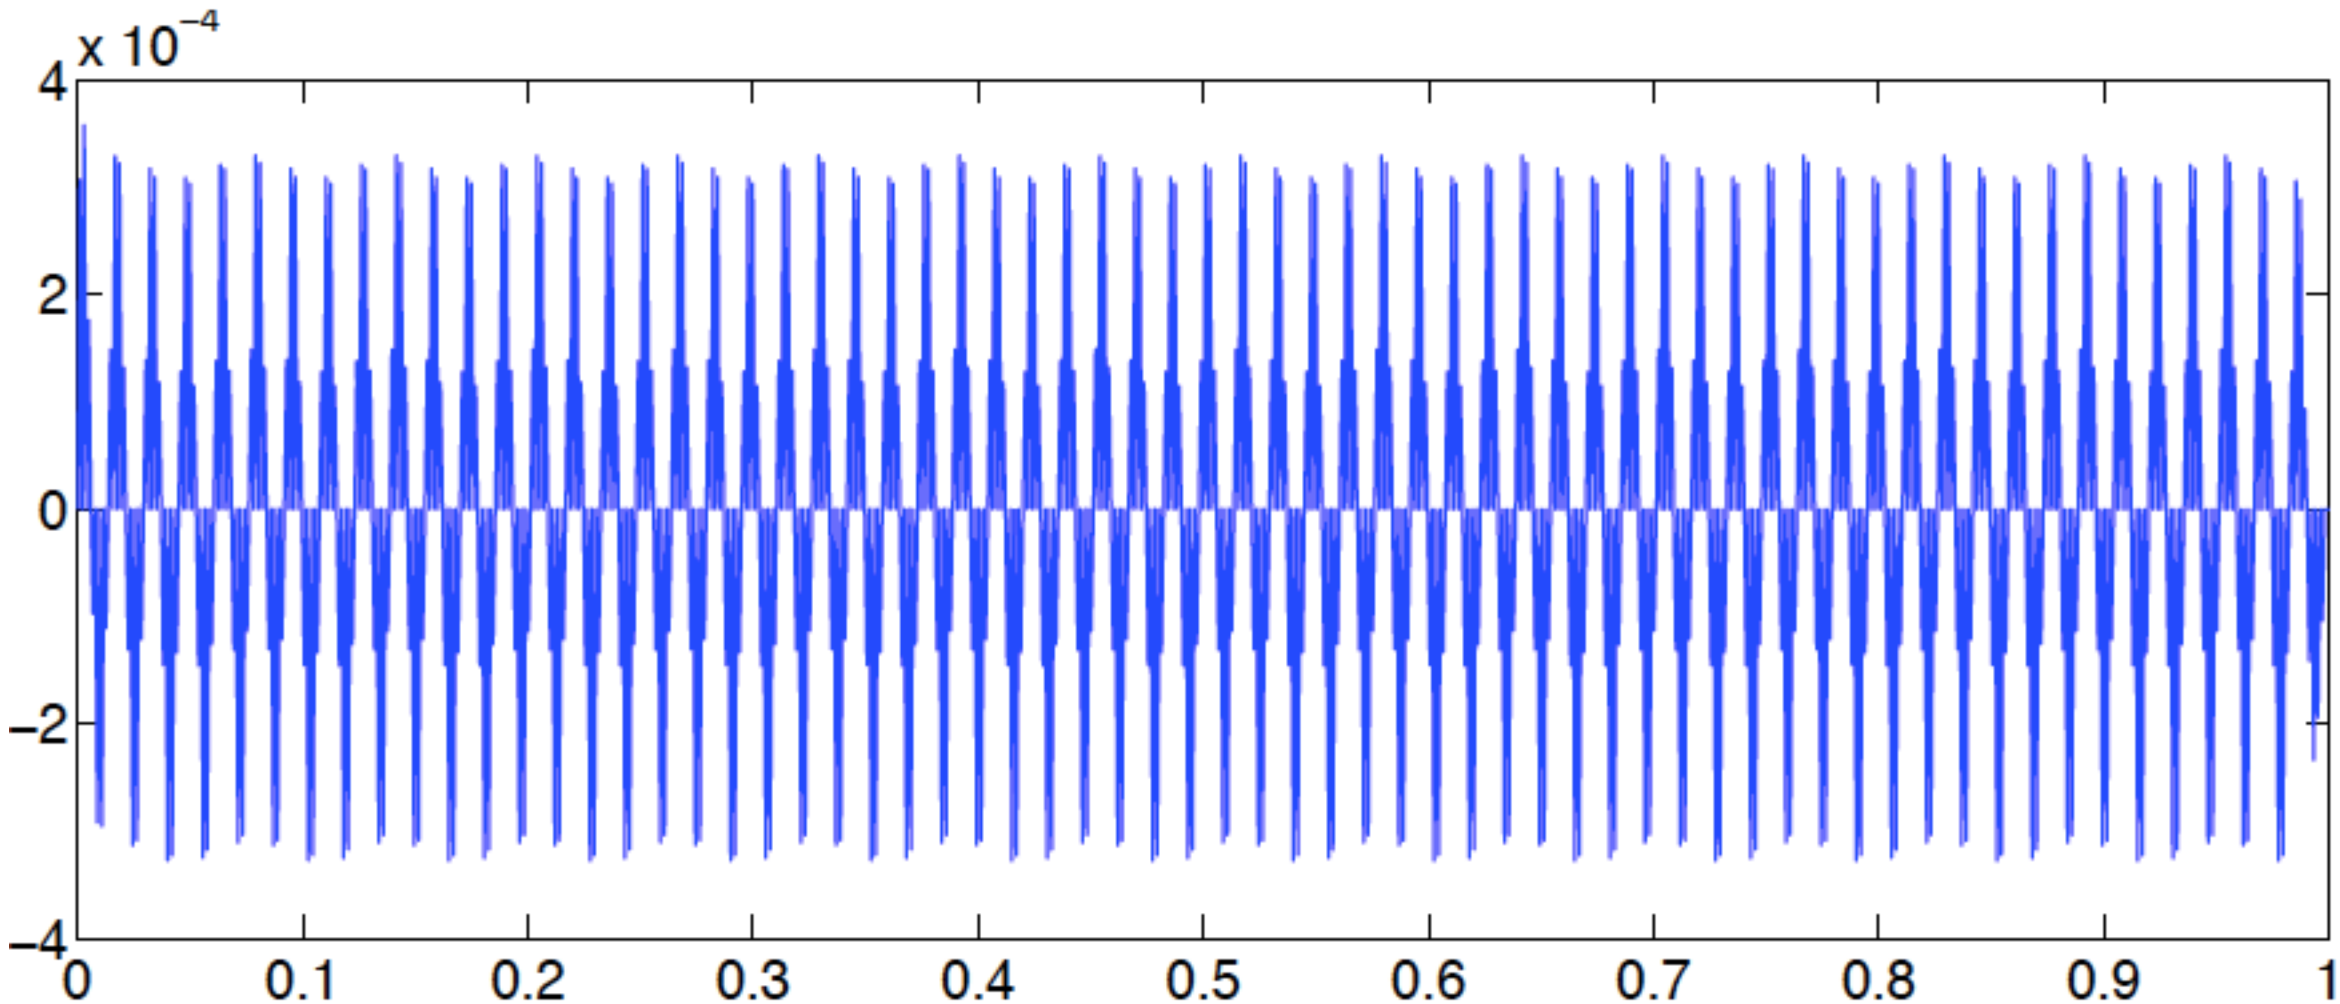

2 GS^8 relaxation   2 CG correction

1-D Laplace Example

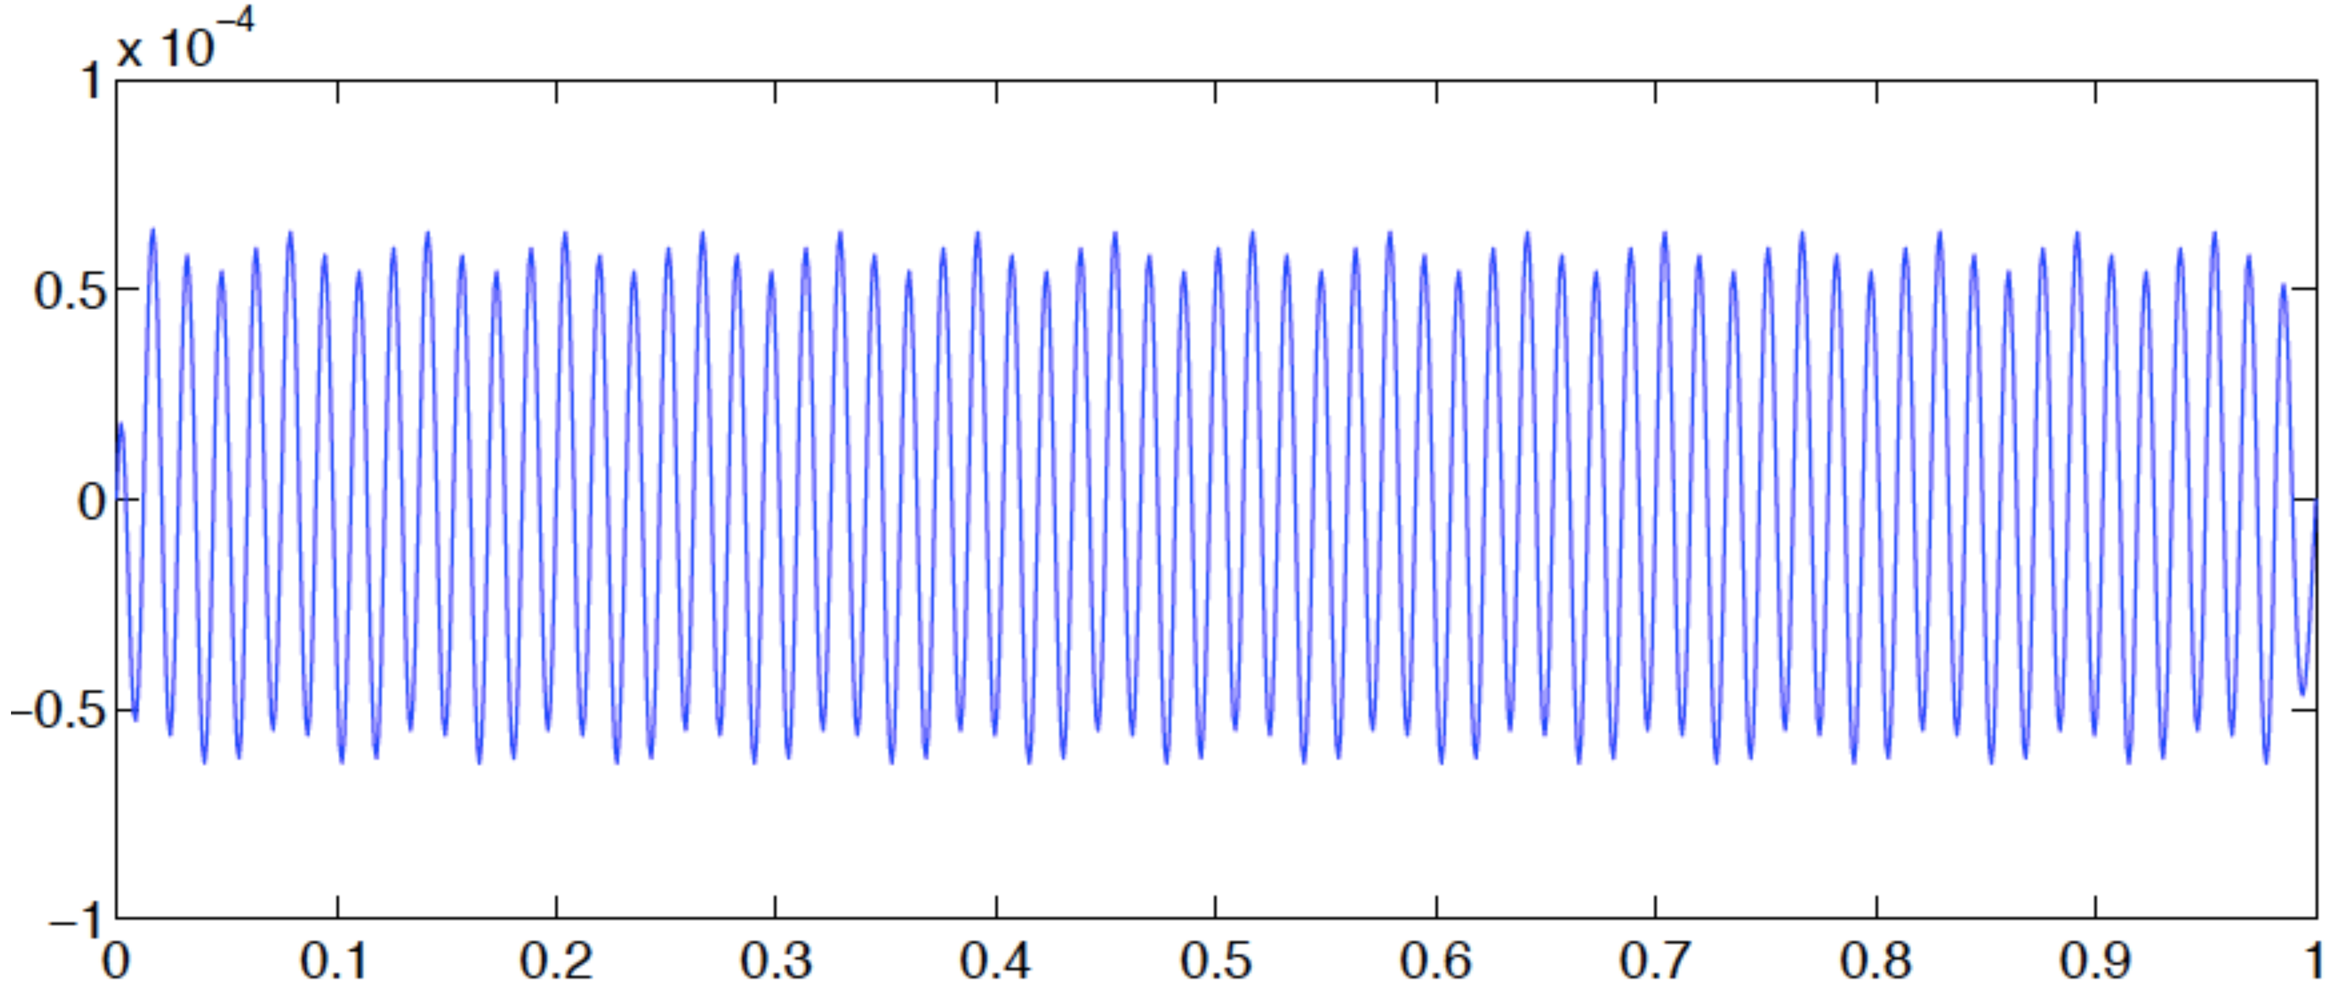

3 GS^8 relaxation   1 CG correction

1-D Laplace Example

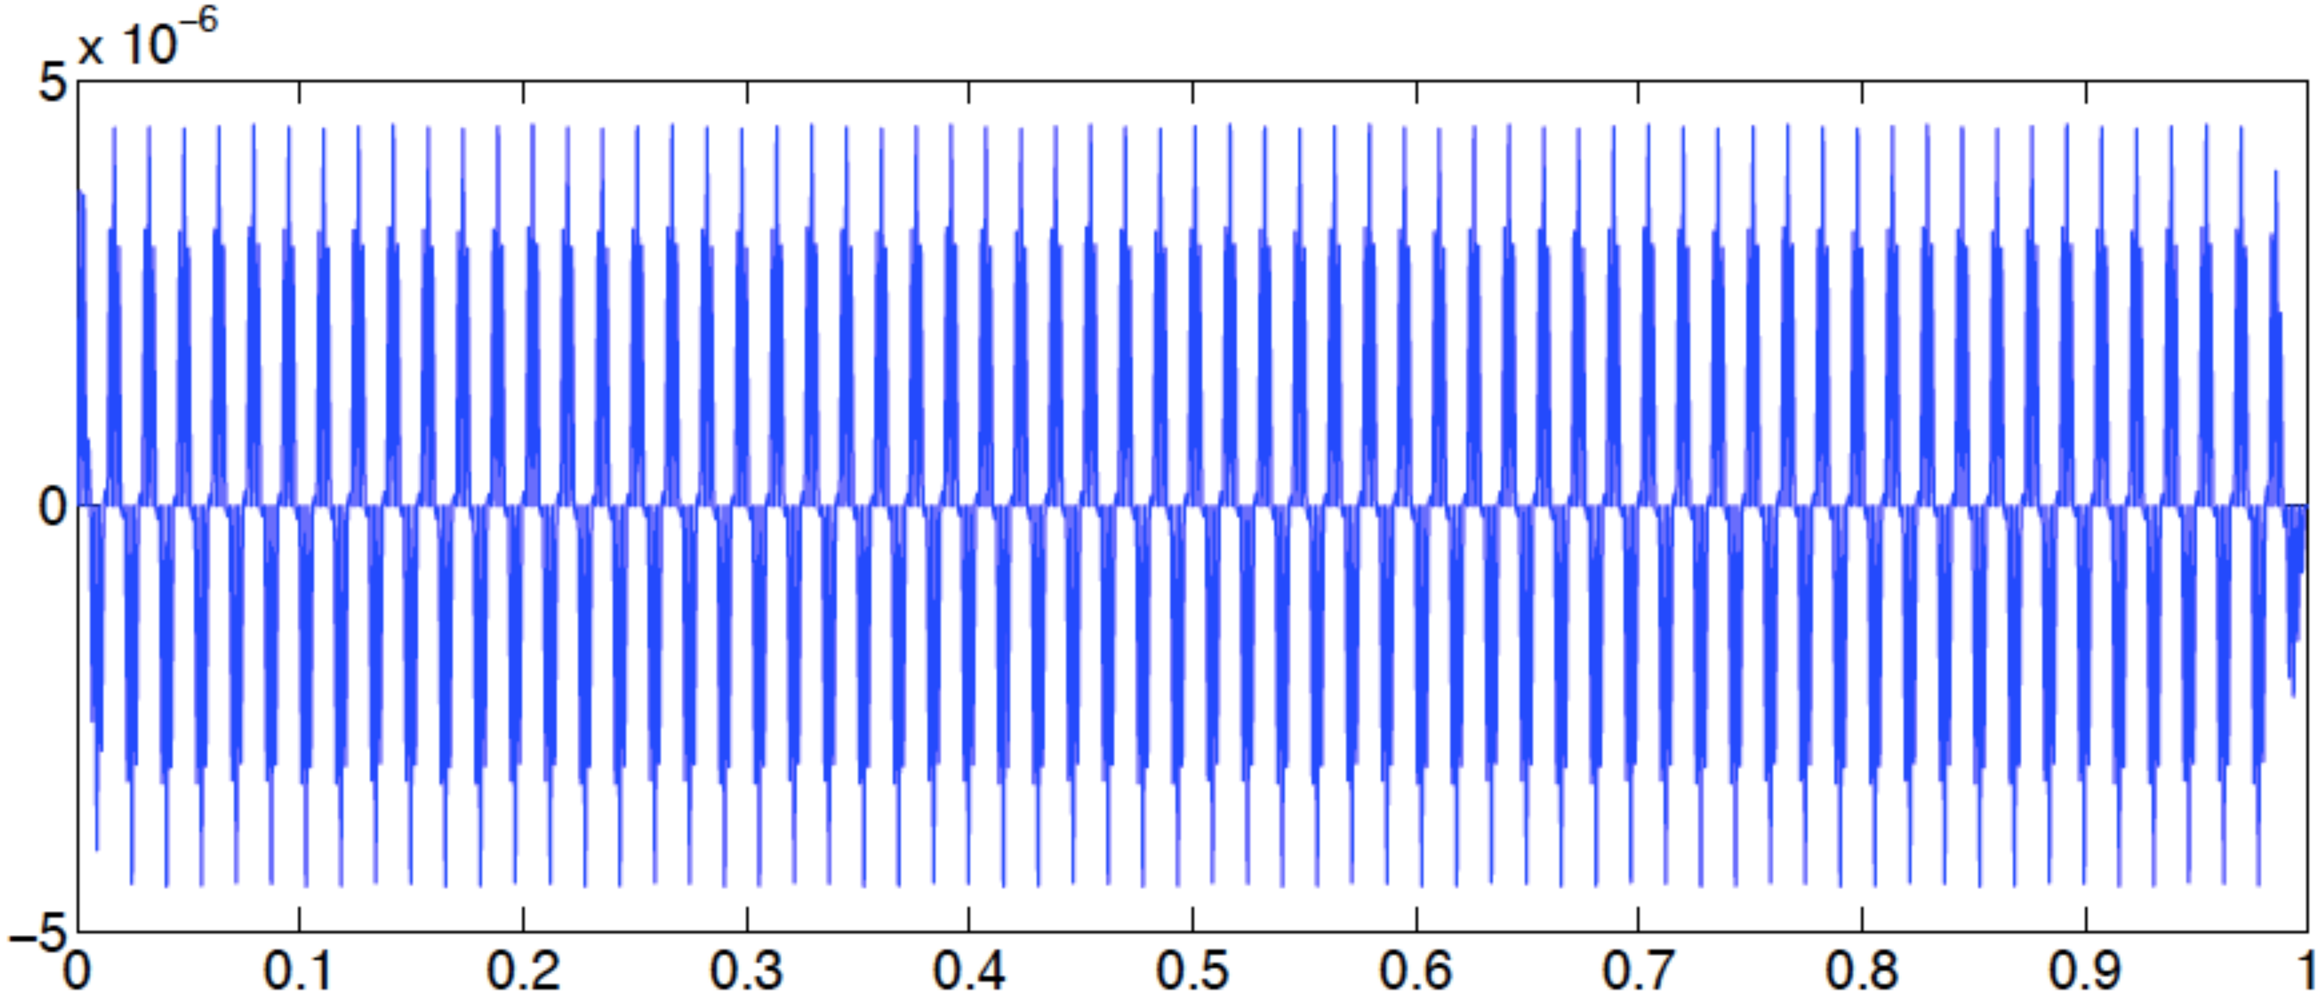

3 GS^8 relaxation   3 CG correction

1-D Laplace Example

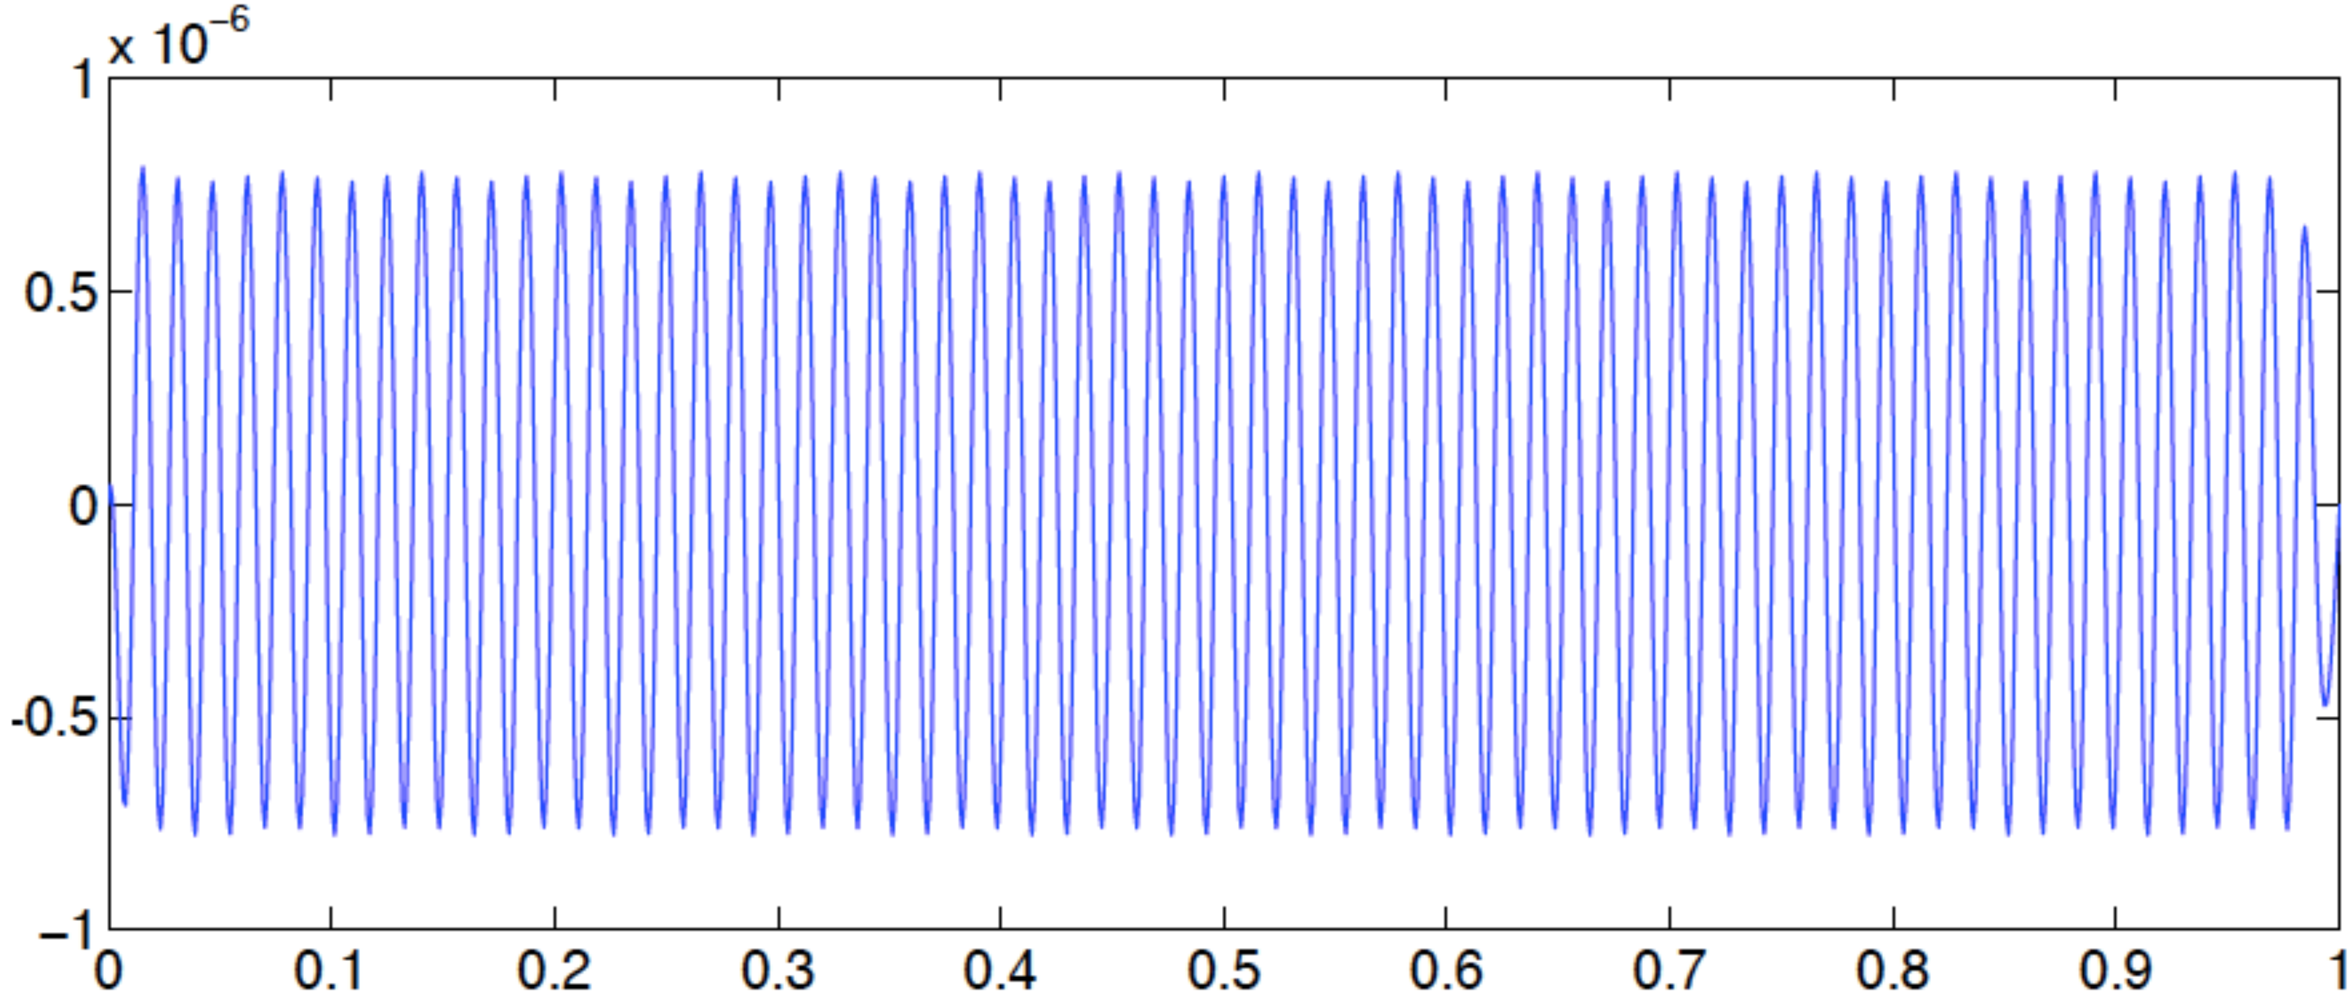

4 GS^8 relaxation 3 CG correction

I-D Laplace Example

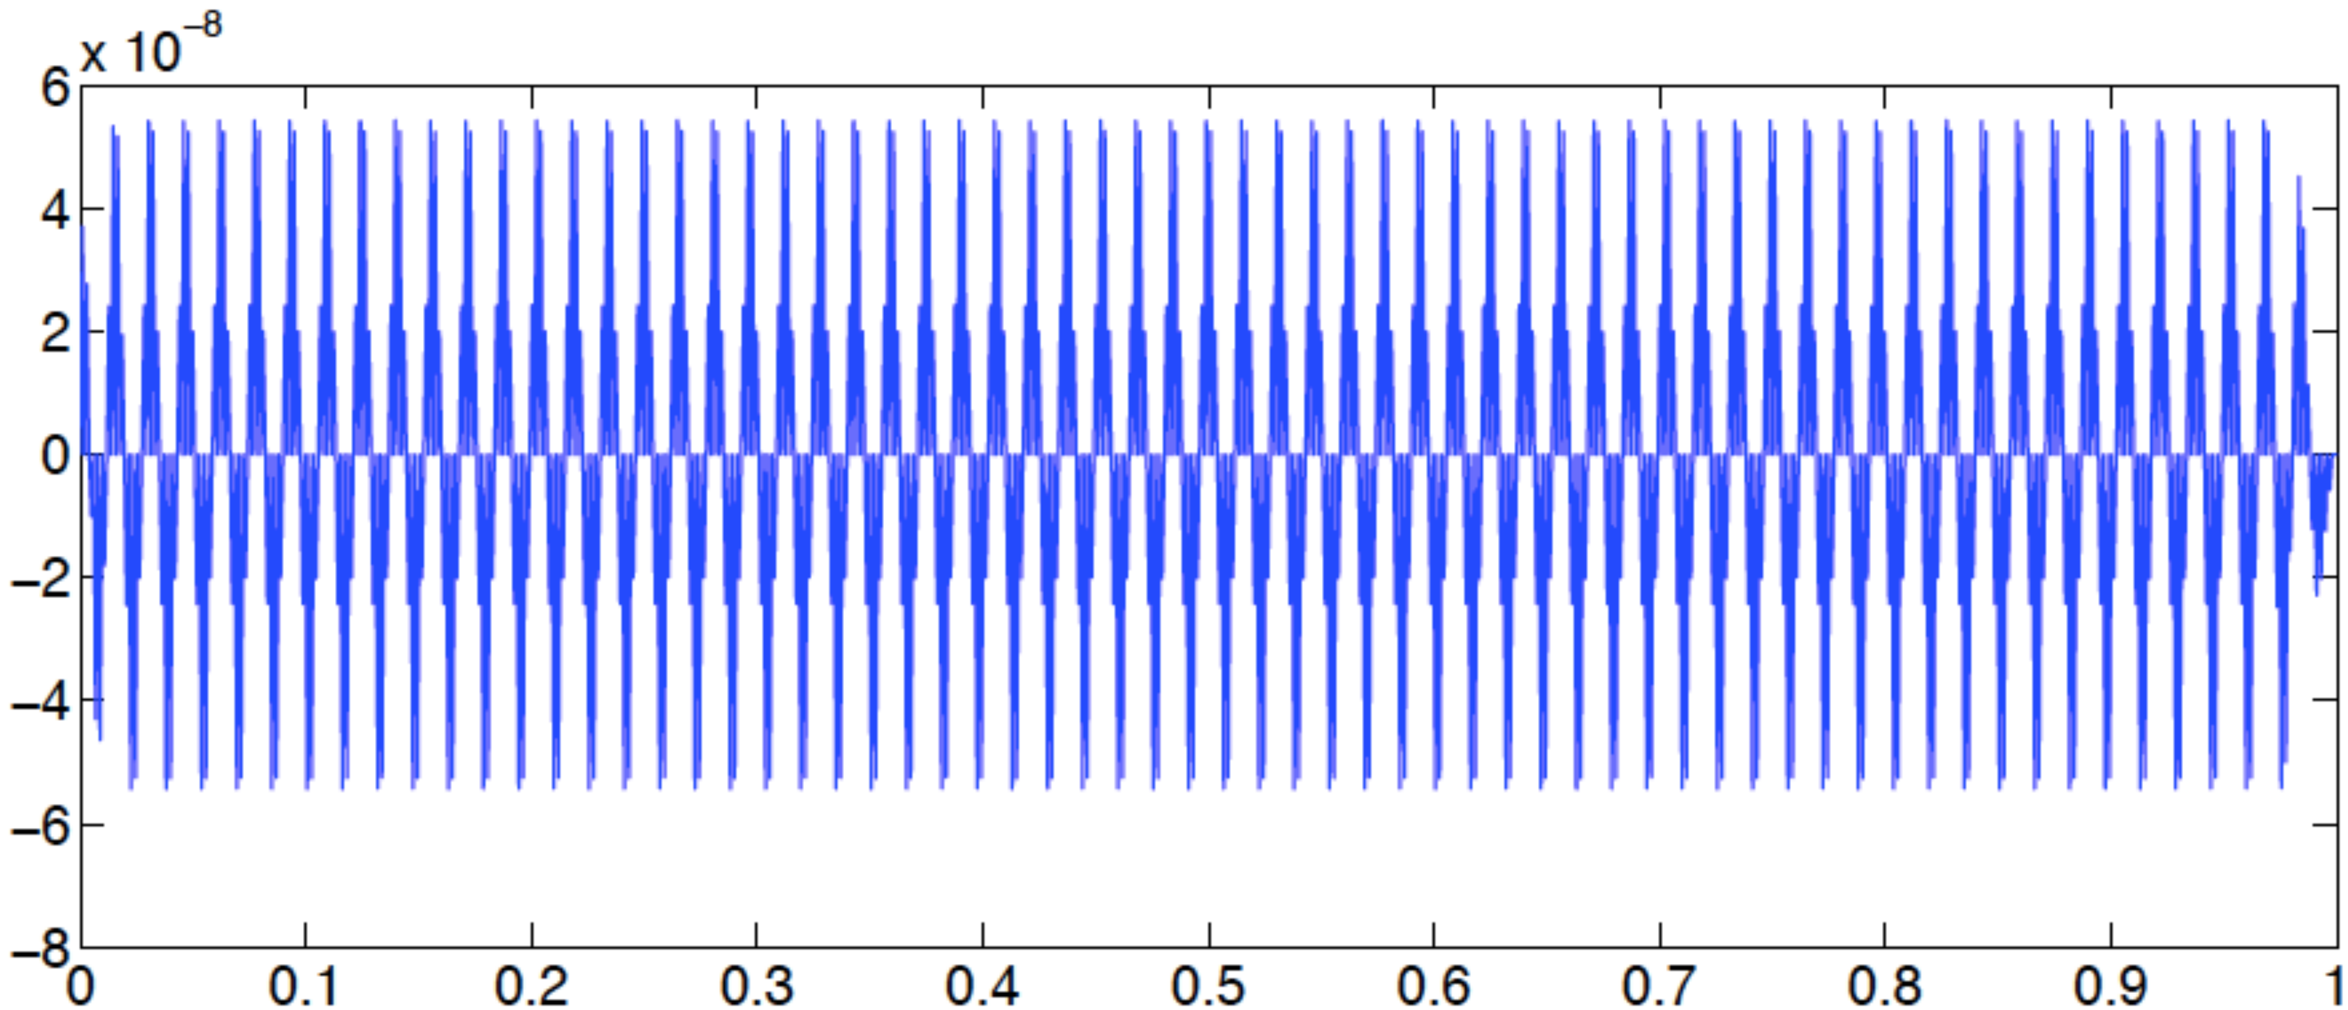

4 GS^8 relaxation 4 CG correction

2-D Laplace Problem

- Five-point stencil discretization on a uniform grid

$$-\nabla^2 u = f$$

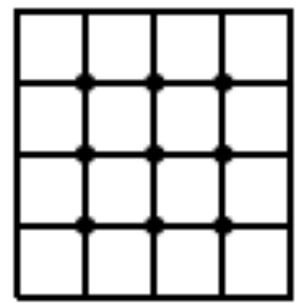

$$A = \begin{bmatrix} & -1 & \\ -1 & 4 & -1 \\ & -1 & \end{bmatrix}$$

- Smoothers: weighted Jacobi or GS (lexicographical or red/black)
- Full coarsening, bilinear interpolation

$$P = \frac{1}{4} \left] \begin{matrix} 1 & 2 & 1 \\ 2 & 4 & 2 \\ 1 & 2 & 1 \end{matrix} \right[$$

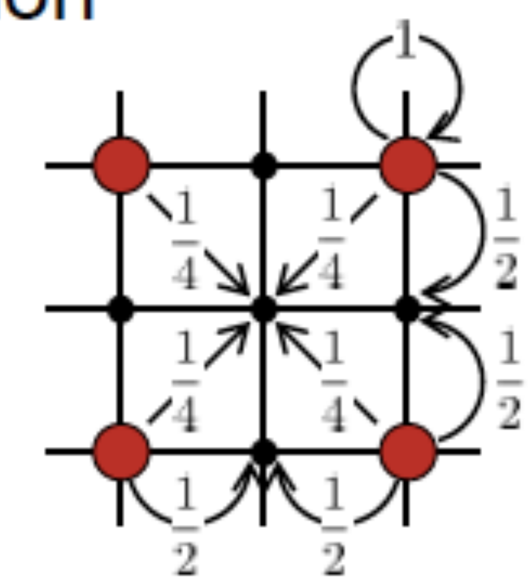

- Coarse discretization (scaled appropriately) for A_c

Full Multigrid (FMG)

Coarse-grid prediction

The coarse grid can also be used to construct an initial approximation for the fine grid.

1. Solve the H-grid problem:

$$A^H u^H = f^H$$

2. Construct an initial approximation for h-grid by prolongation:

$$u^{h,0} = Pu^H$$

Remarks

1. Of course, the coarse-grid prediction can again be applied recursively
2. Since Pu^H is only an initial approximation, it is not necessary to fully resolve u^H: a suitable approximation will do

Full Multigrid (FMG)

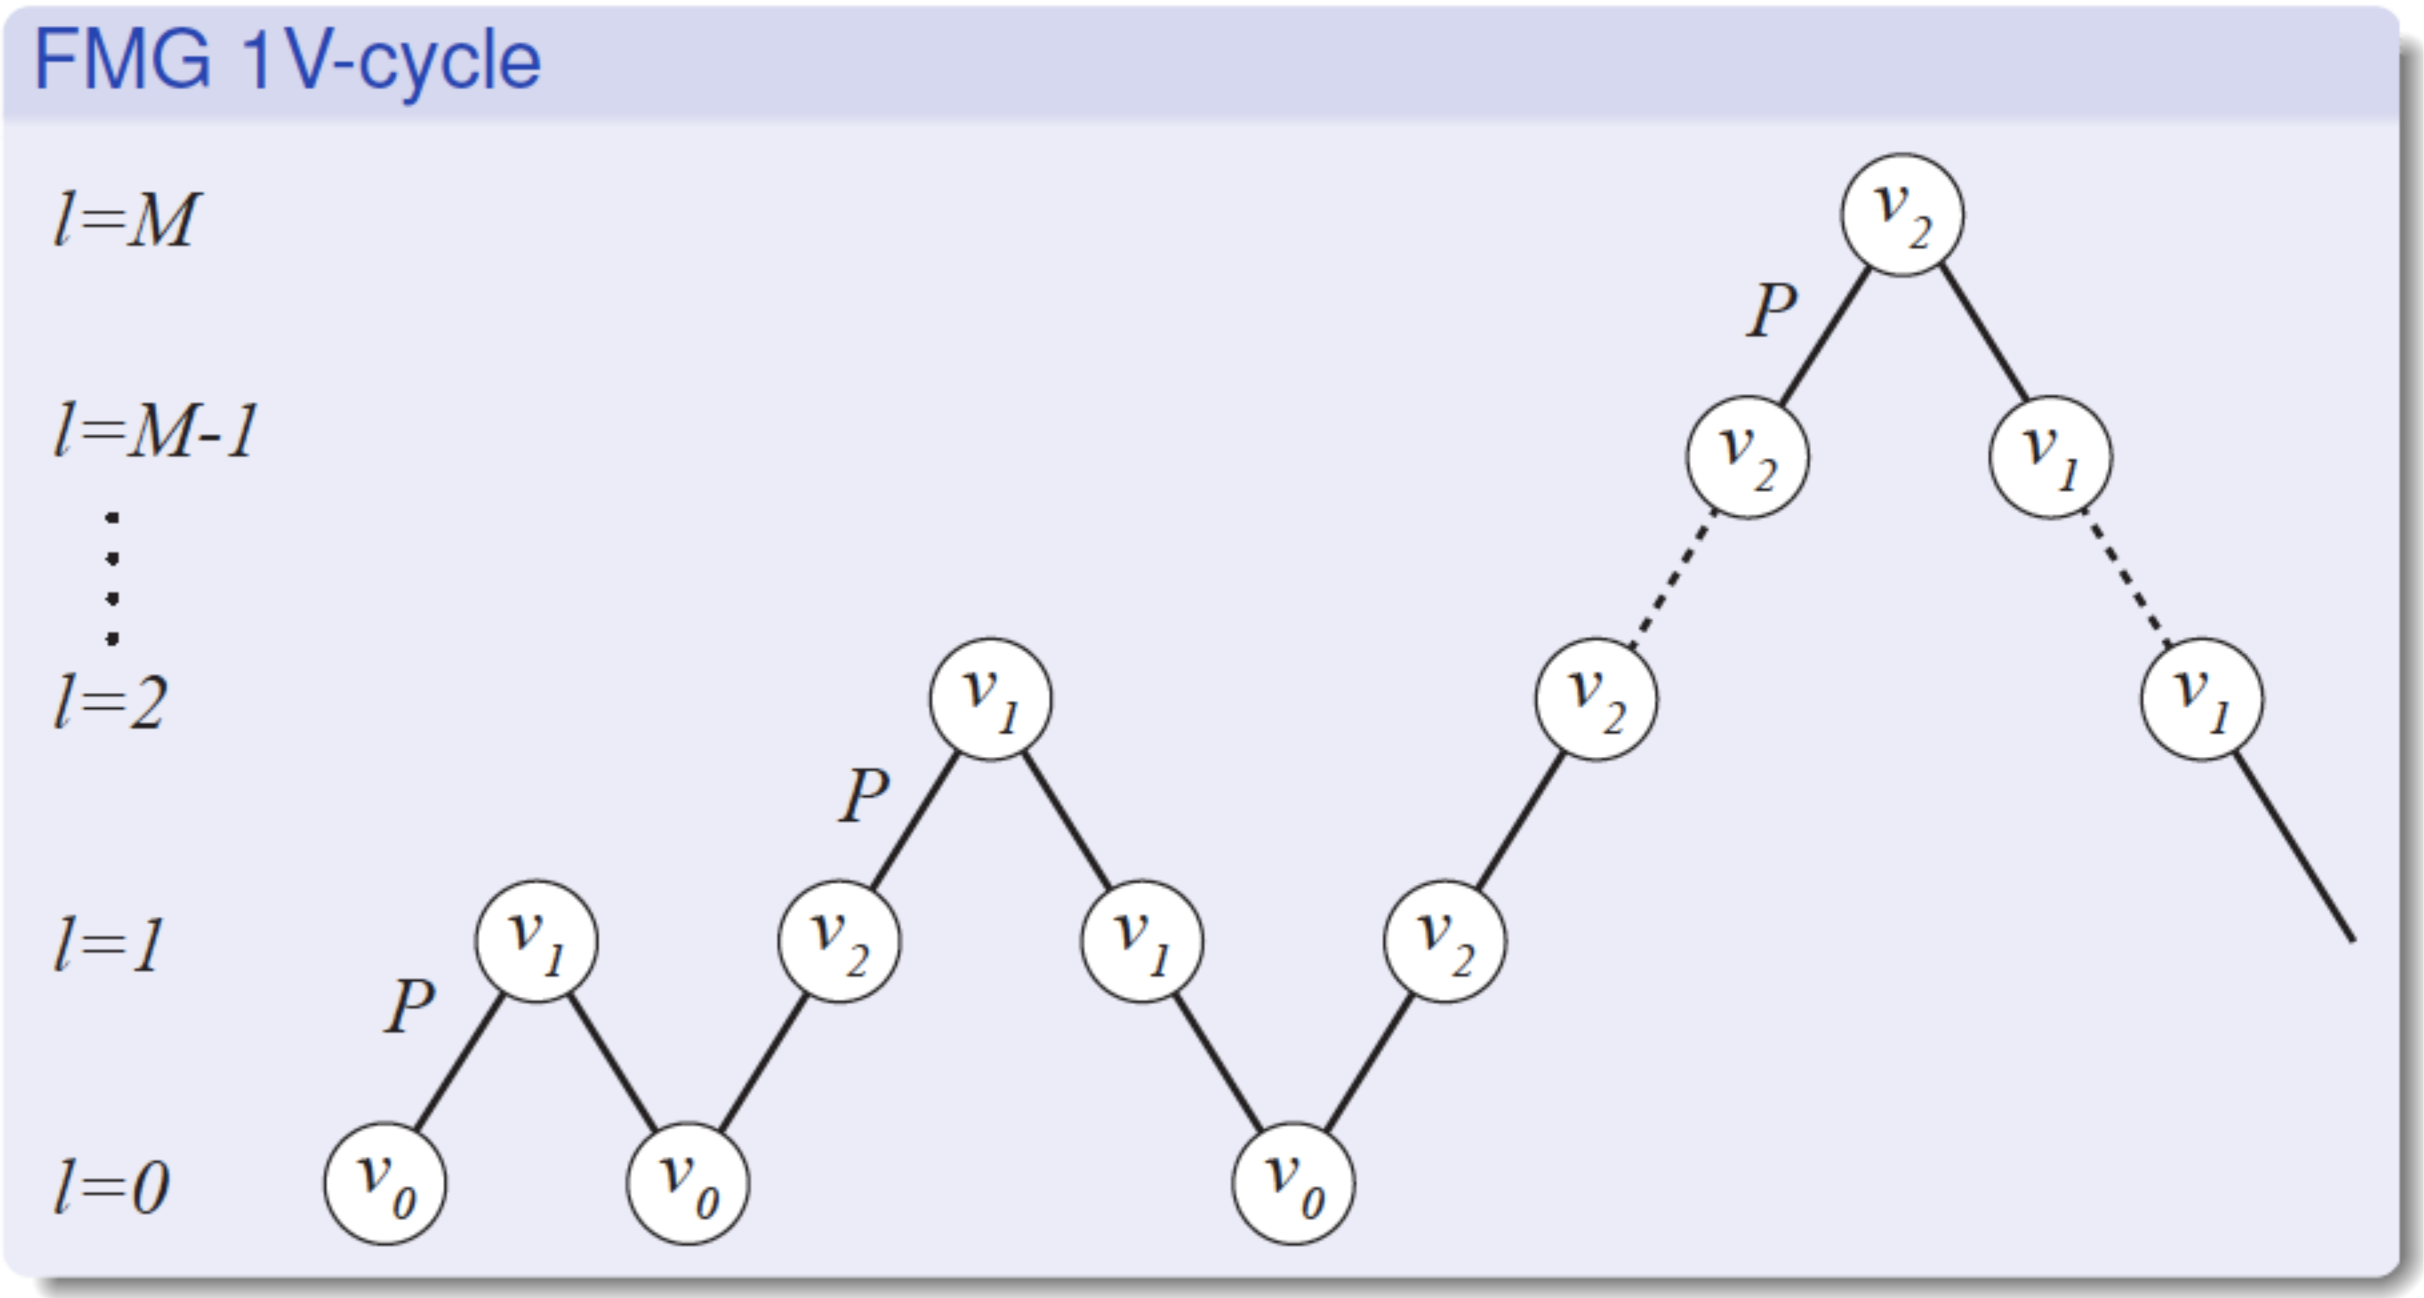

Parallel Multigrid

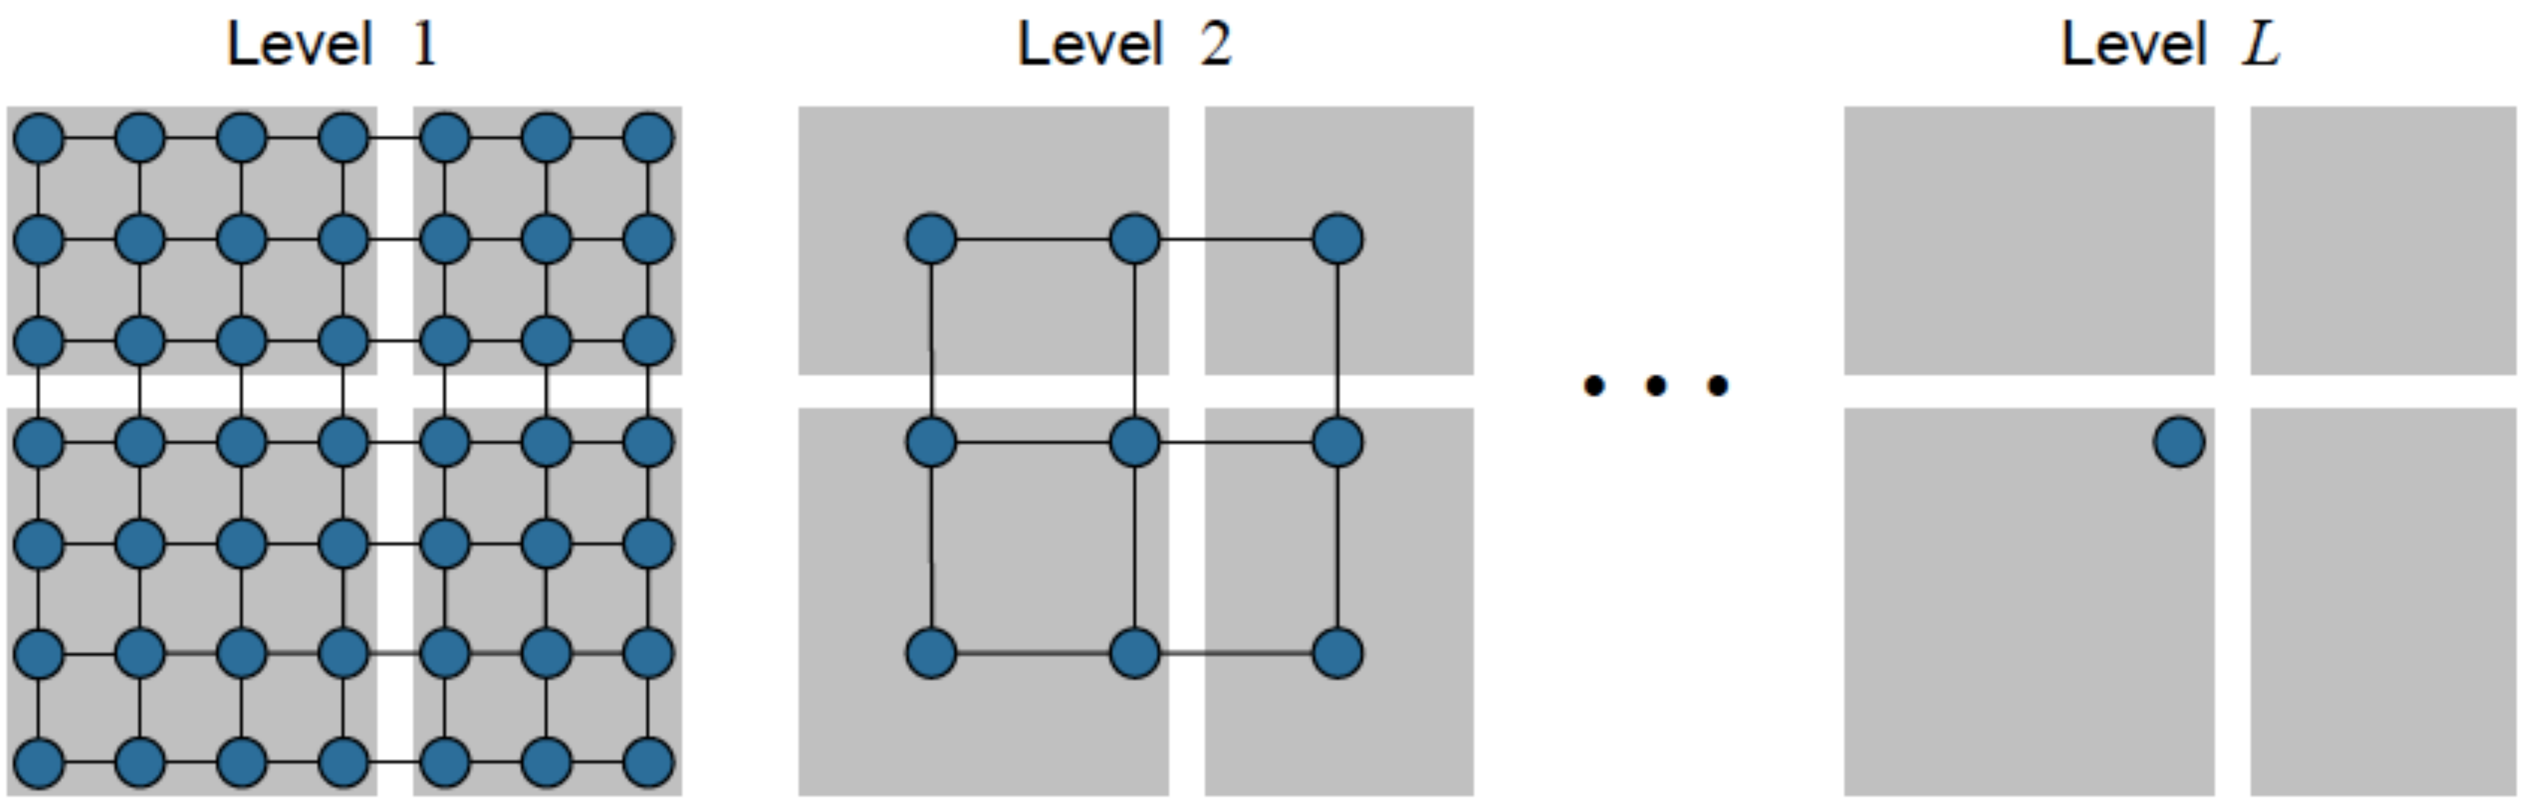

- Basic communication pattern is "nearest neighbor"
  - Relaxation, interpolation, & Galerkin not hard to implement
- Different neighbor processors on coarse grids
- Many idle processors on coarse grids
  - Algorithms to take advantage have had limited success

Parallel Multigrid

- Standard communication / computation models

$$T_{comm} = \alpha + m\beta \quad \text{(communicate m doubles)}$$
$$T_{comp} = m\gamma \quad \text{(compute m flops)}$$

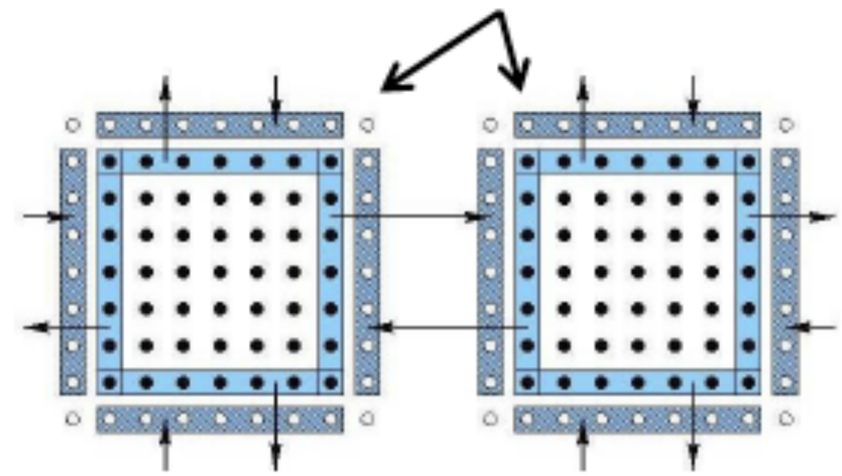

- Time to do relaxation

$$T \approx 4\alpha + 4n\beta + 5n^2\gamma$$

- Time to do relaxation in a V(1,0) multigrid cycle

$$T_V \approx (1 + 1 + \cdots)4\alpha + (1 + 1/2 + \ldots)4n\beta + (1 + 1/4 + \ldots)5n^2\gamma$$
$$\approx (\log N)4\alpha + (2)4n\beta + (4/3)5n^2\gamma$$

- For achieving optimality in general, the *log* term is unavoidable!

- More precise: $T_{V,better} \approx T_V + (\log P)(4\beta + 5\gamma)$

Algebraic Multigrid

- For best results, geometry alone is not enough

Linear Interpolation

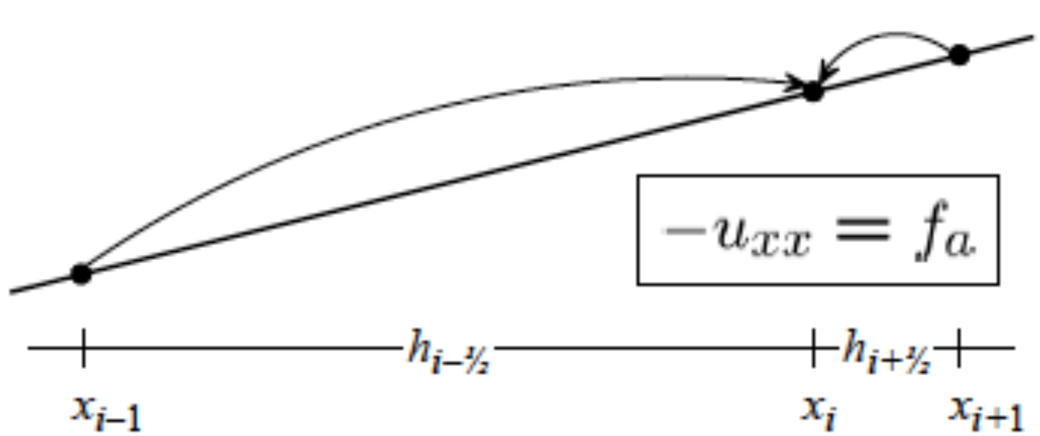

Operator-Dependent Interpolation

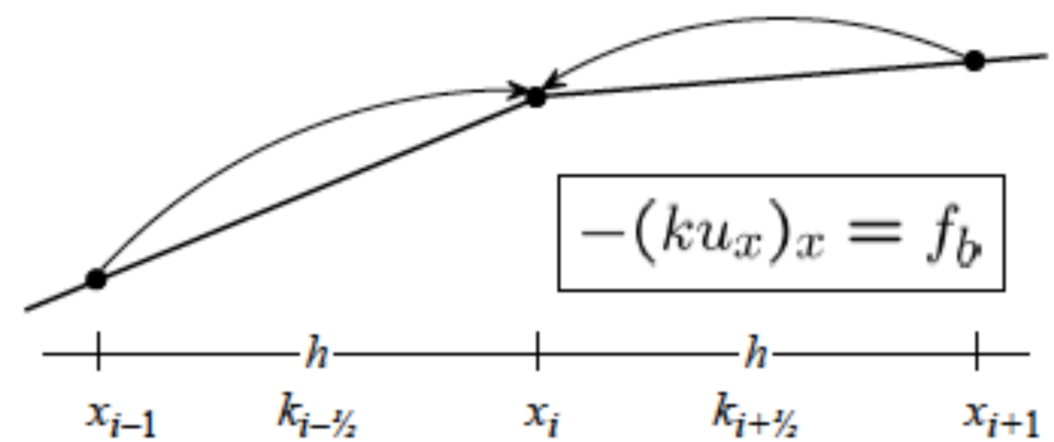

- AMG ignores geometric information altogether, but captures both linear & operator-dep interpolation

$$(A\mathbf{u})_i = a_{i,i-1}u_{i-1} + a_{i,i}u_i + a_{i,i+1}u_{i+1}$$

$$u_i = \left(-\frac{a_{i,i-1}}{a_{i,i}}\right)u_{i-1} + \left(-\frac{a_{i,i+1}}{a_{i,i}}\right)u_{i+1}$$

Algebraic Multigrid

- 7 GS sweeps on

$$-au_{xx} - bu_{yy} = f$$

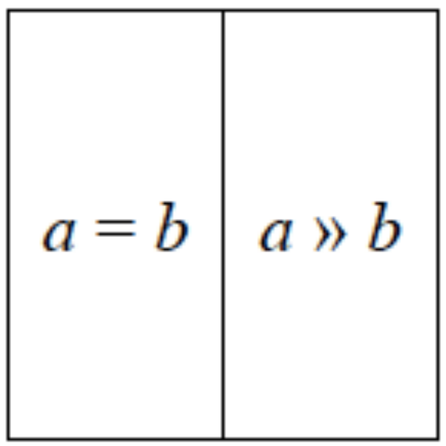

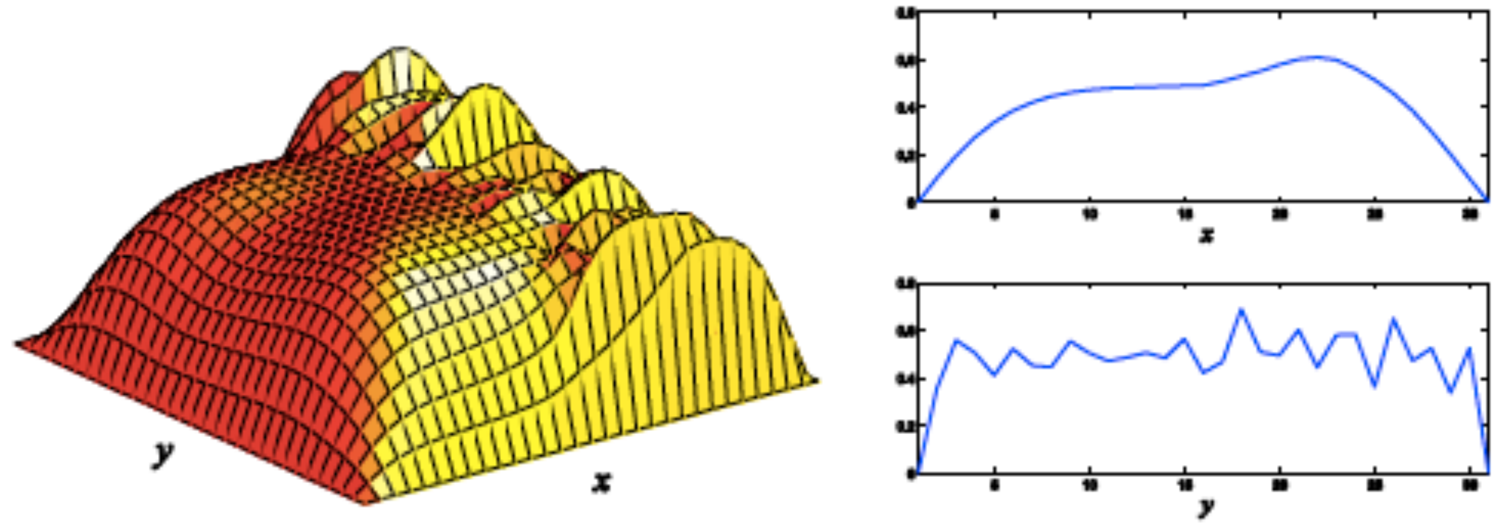

- This example…
  - targets geometric smoothness
  - uses pointwise smoothers
- Not sufficient for some problems!

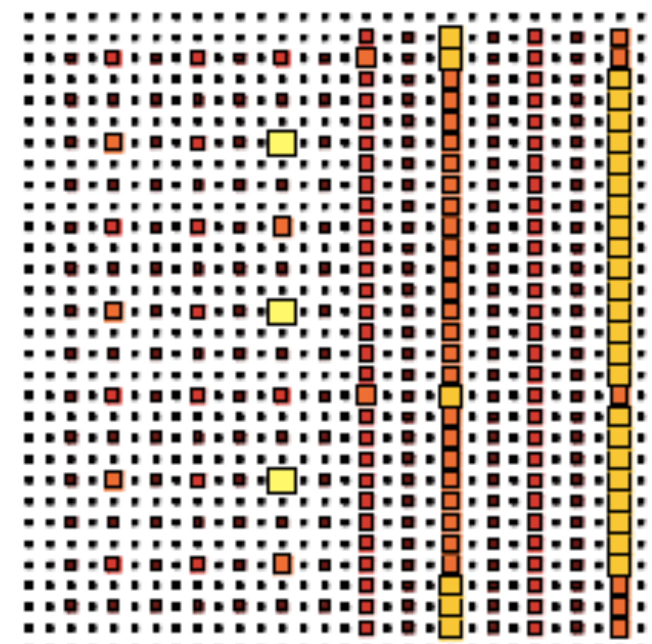

AMG coarsens grids in the direction of geometric smoothness

AMG Coarsening

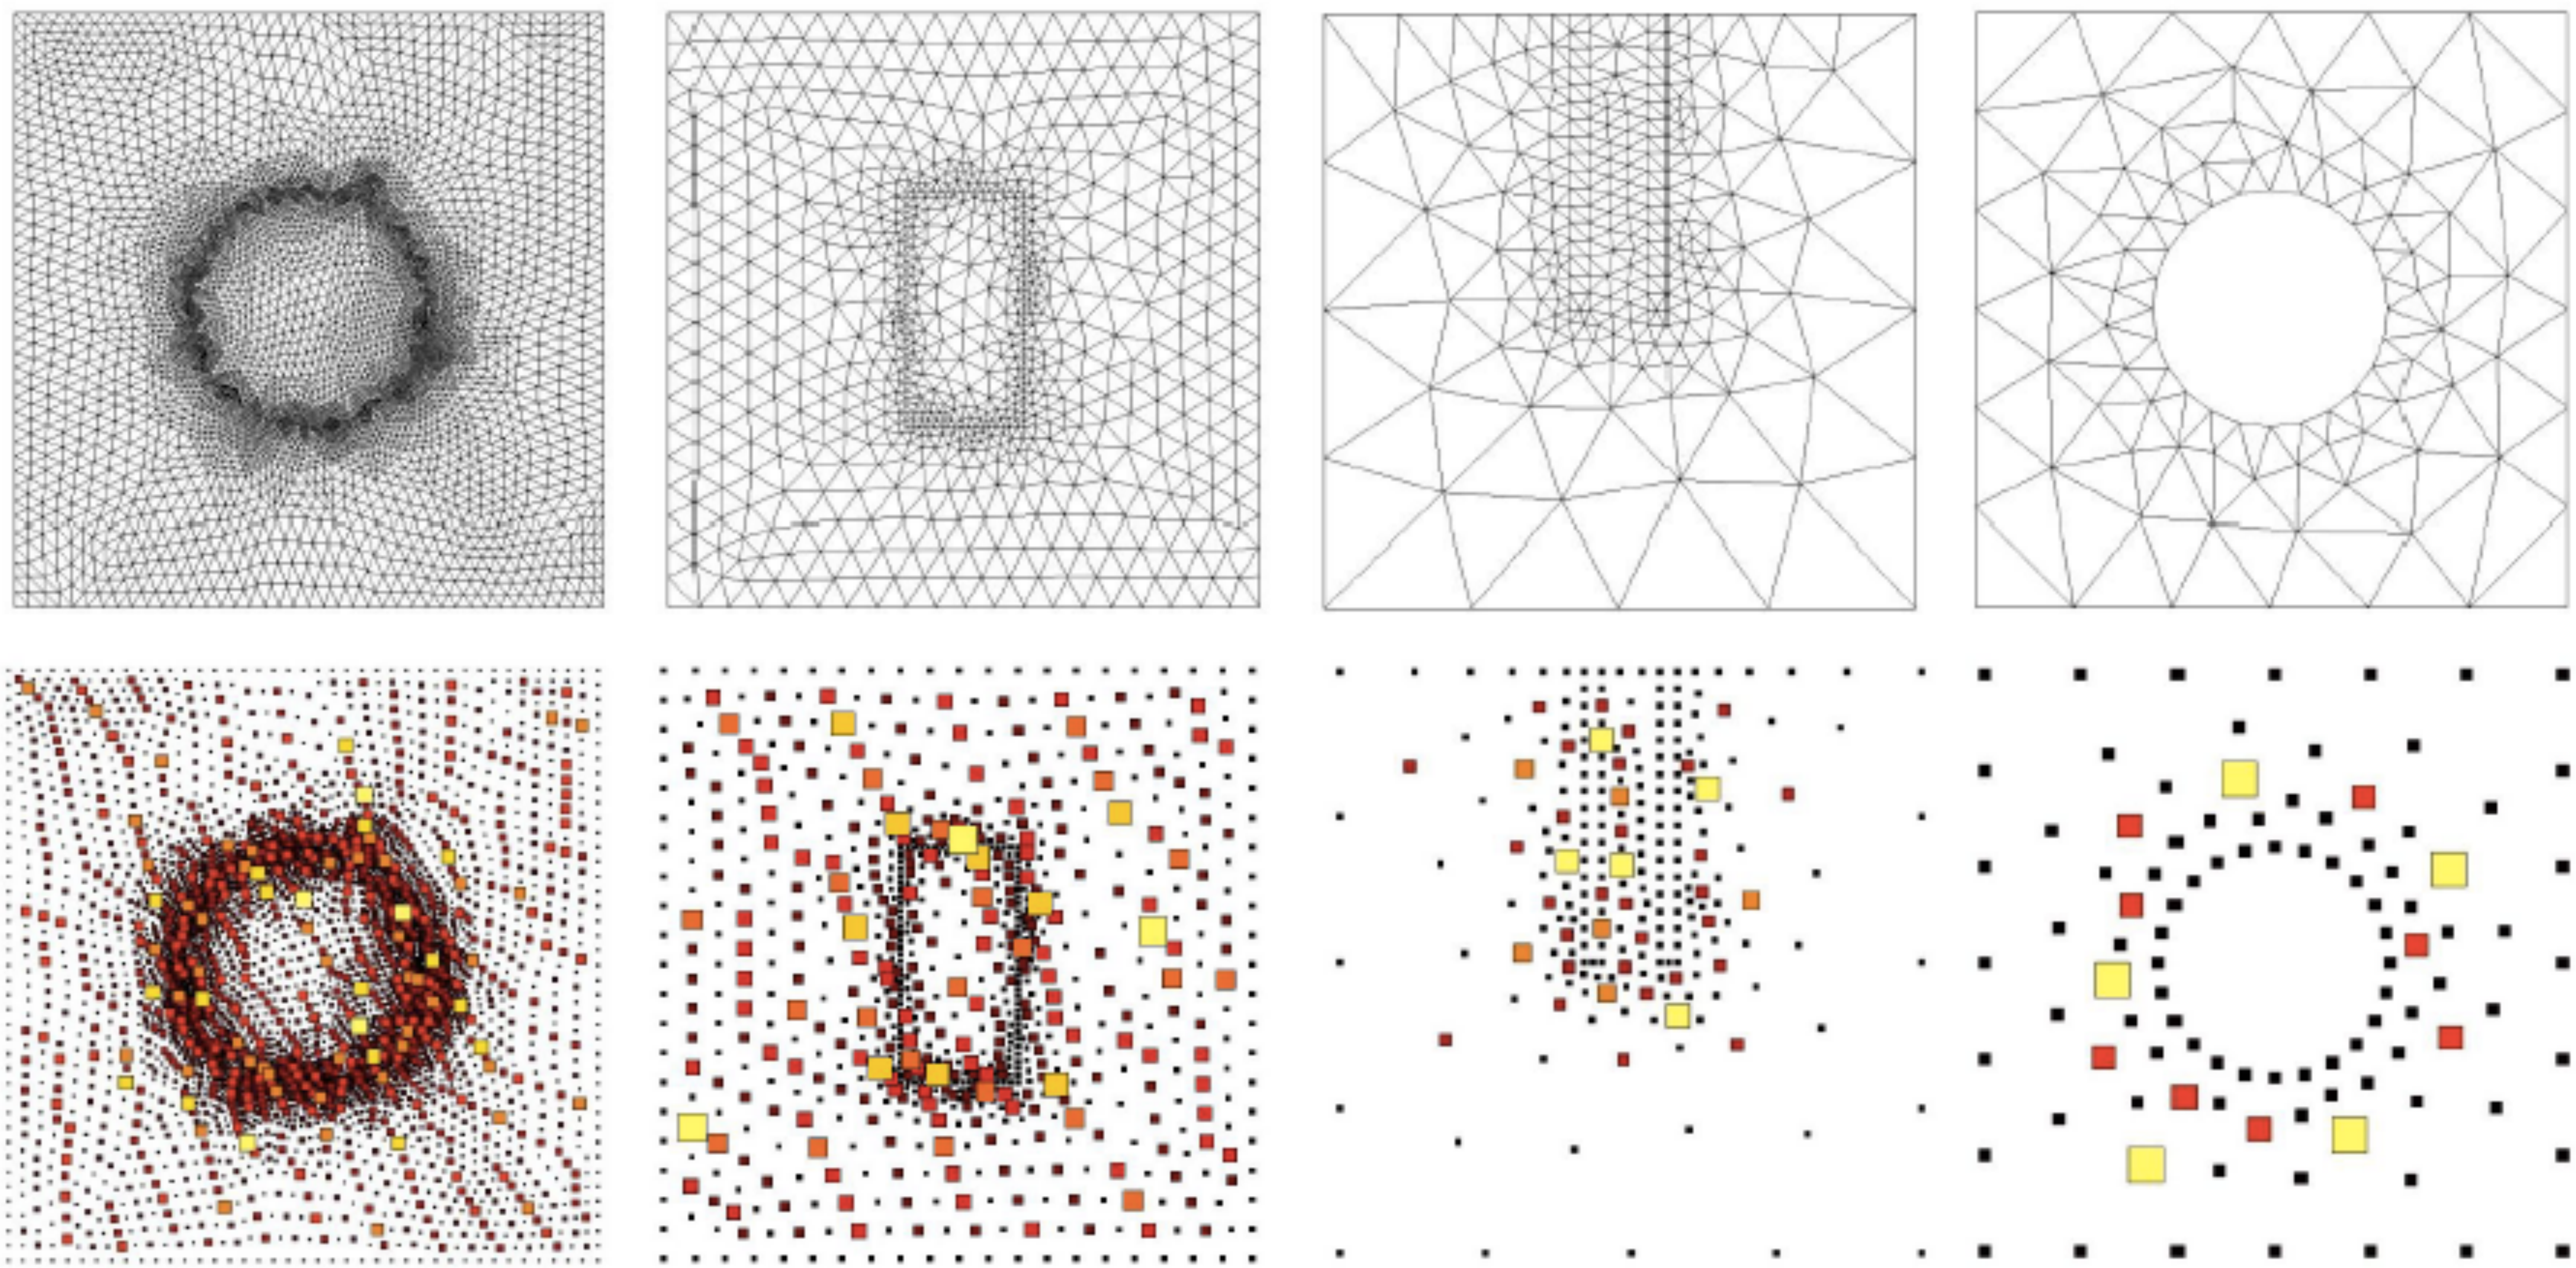

AMG Coarsening

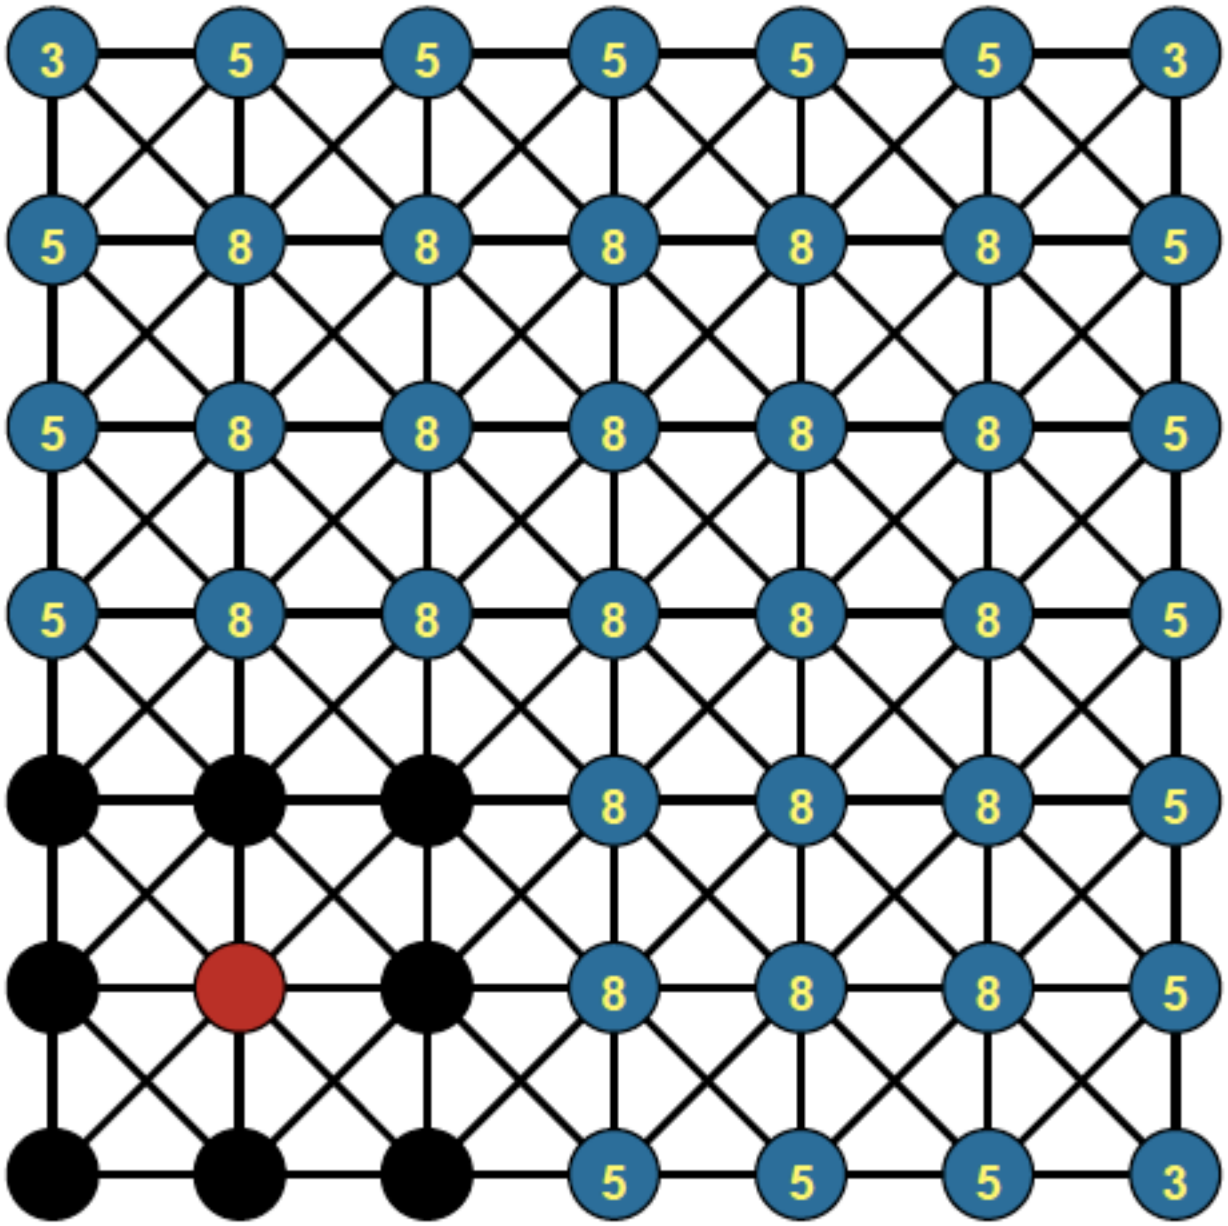

➔ select C-pt with maximal measure

➔ select neighbors as F-pts

➔ update measures of F-pt neighbors

AMG Coarsening

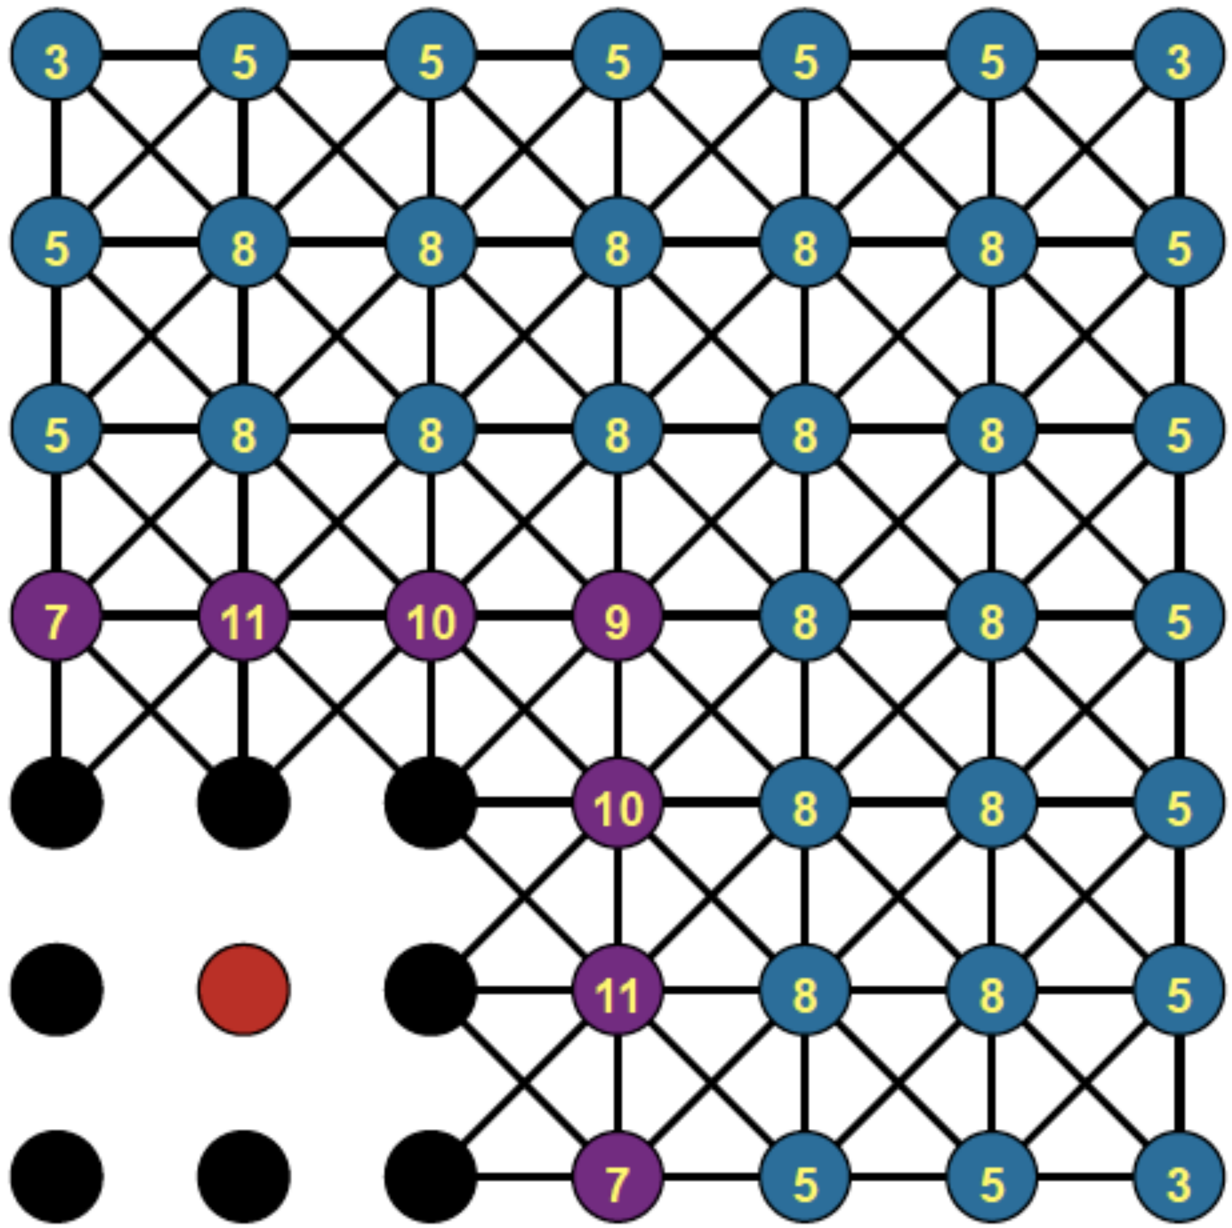

- ➔ **select C-pt with maximal measure**
- ➔ **select neighbors as F-pts**
- ➔ **update measures of F-pt neighbors**

AMG Coarsening

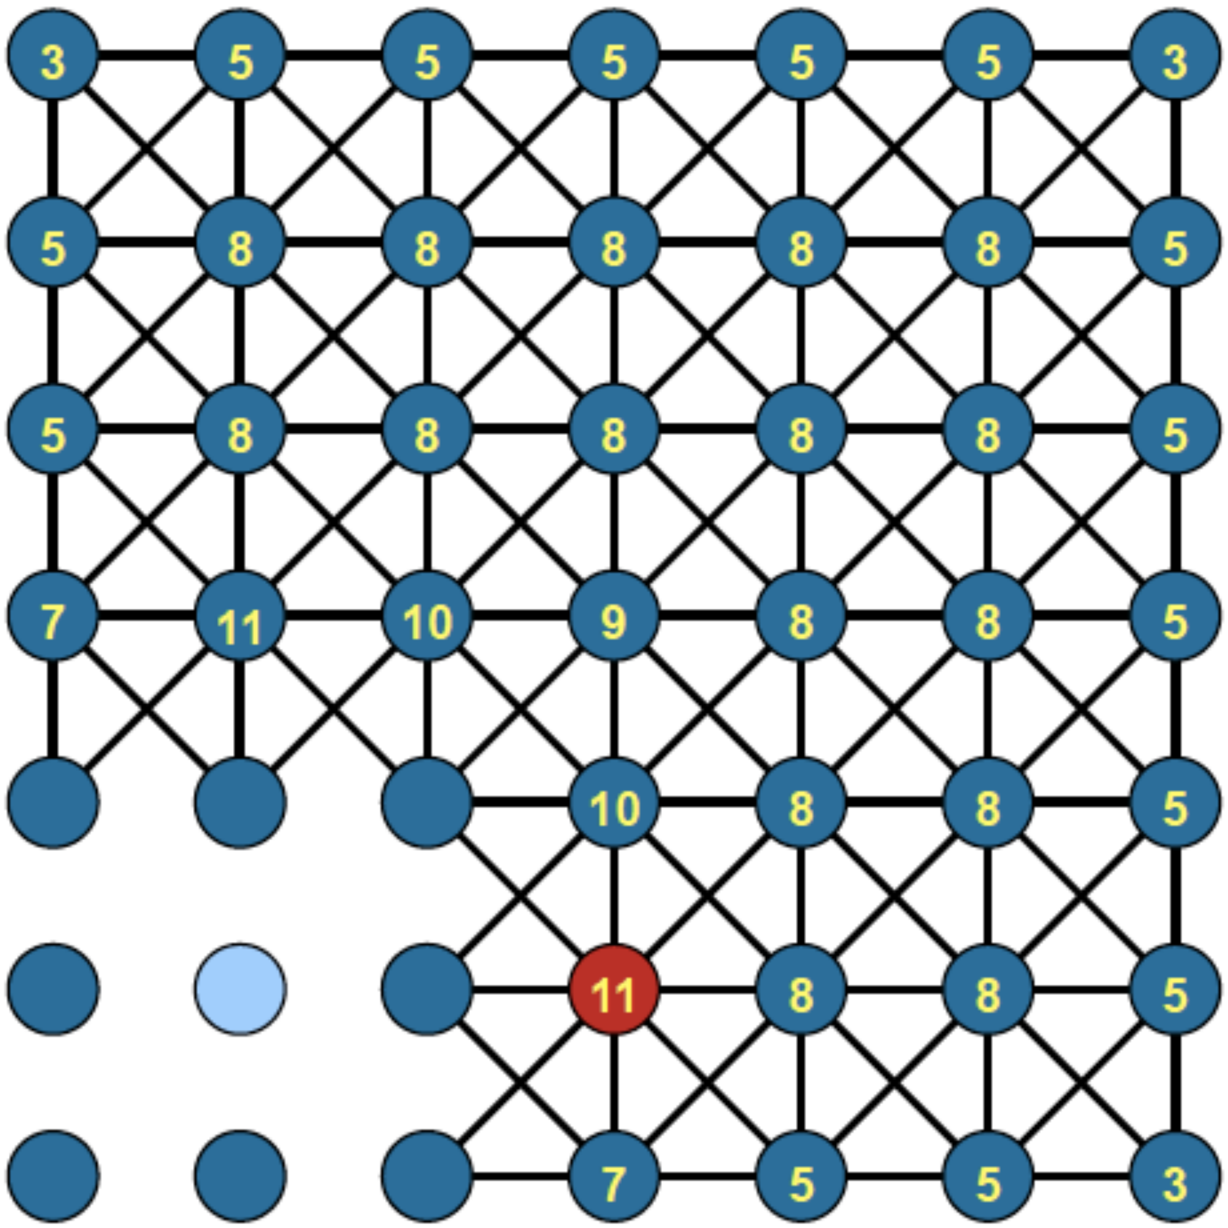

- ➔ select C-pt with maximal measure
- ➔ select neighbors as F-pts
- ➔ update measures of F-pt neighbors

AMG Coarsening

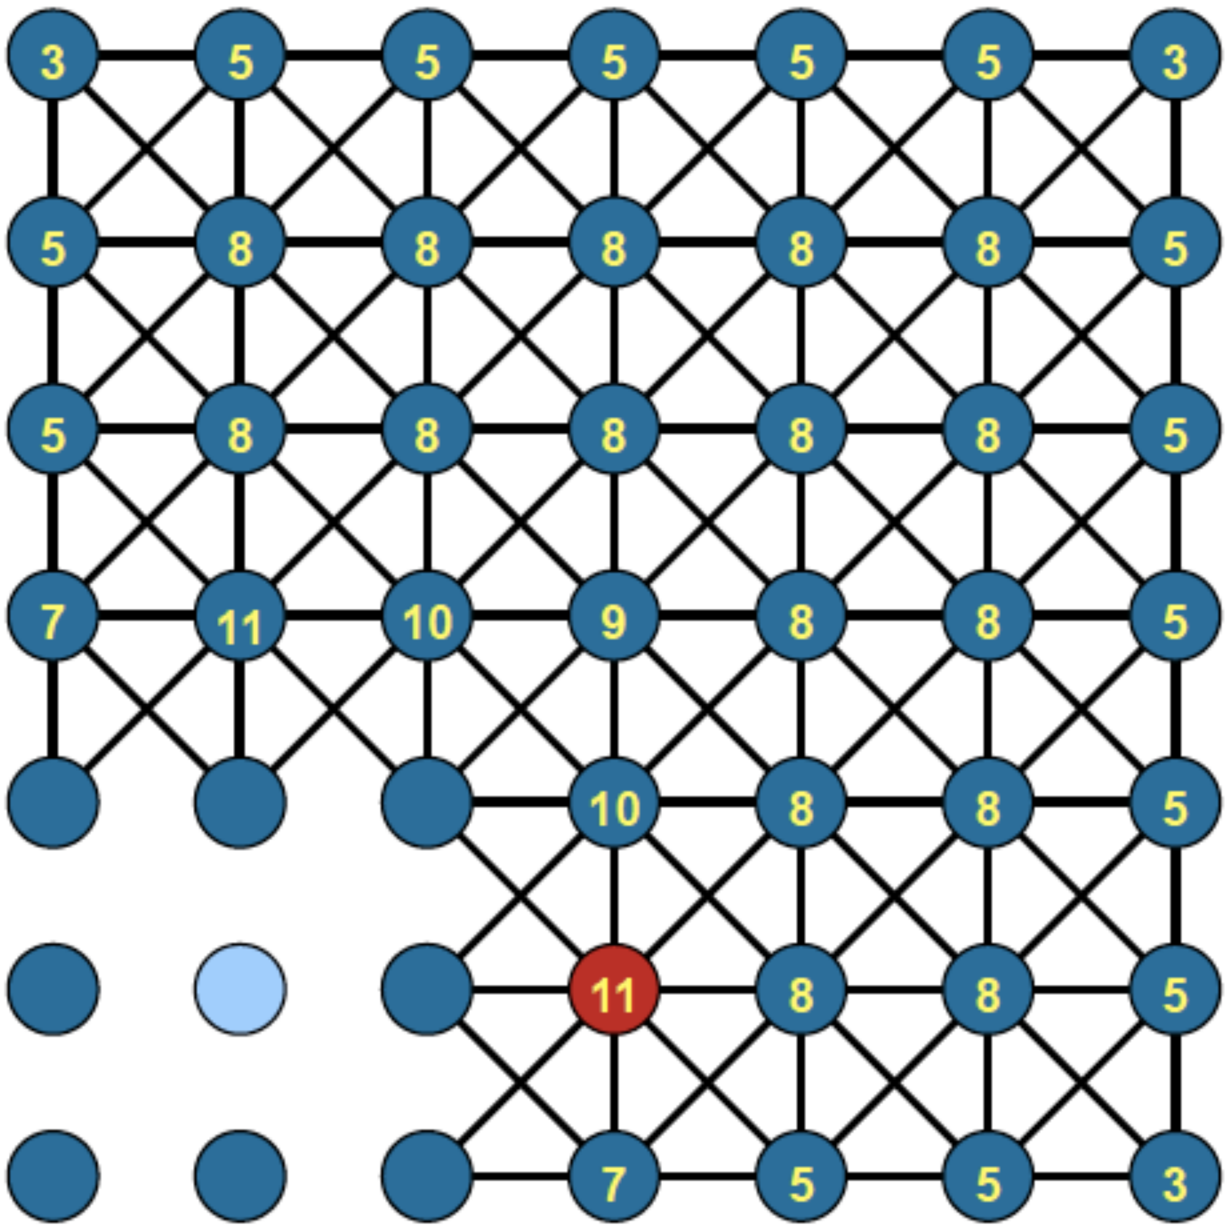

- ➔ select C-pt with maximal measure
- ➔ select neighbors as F-pts
- ➔ update measures of F-pt neighbors

AMG Interpolation

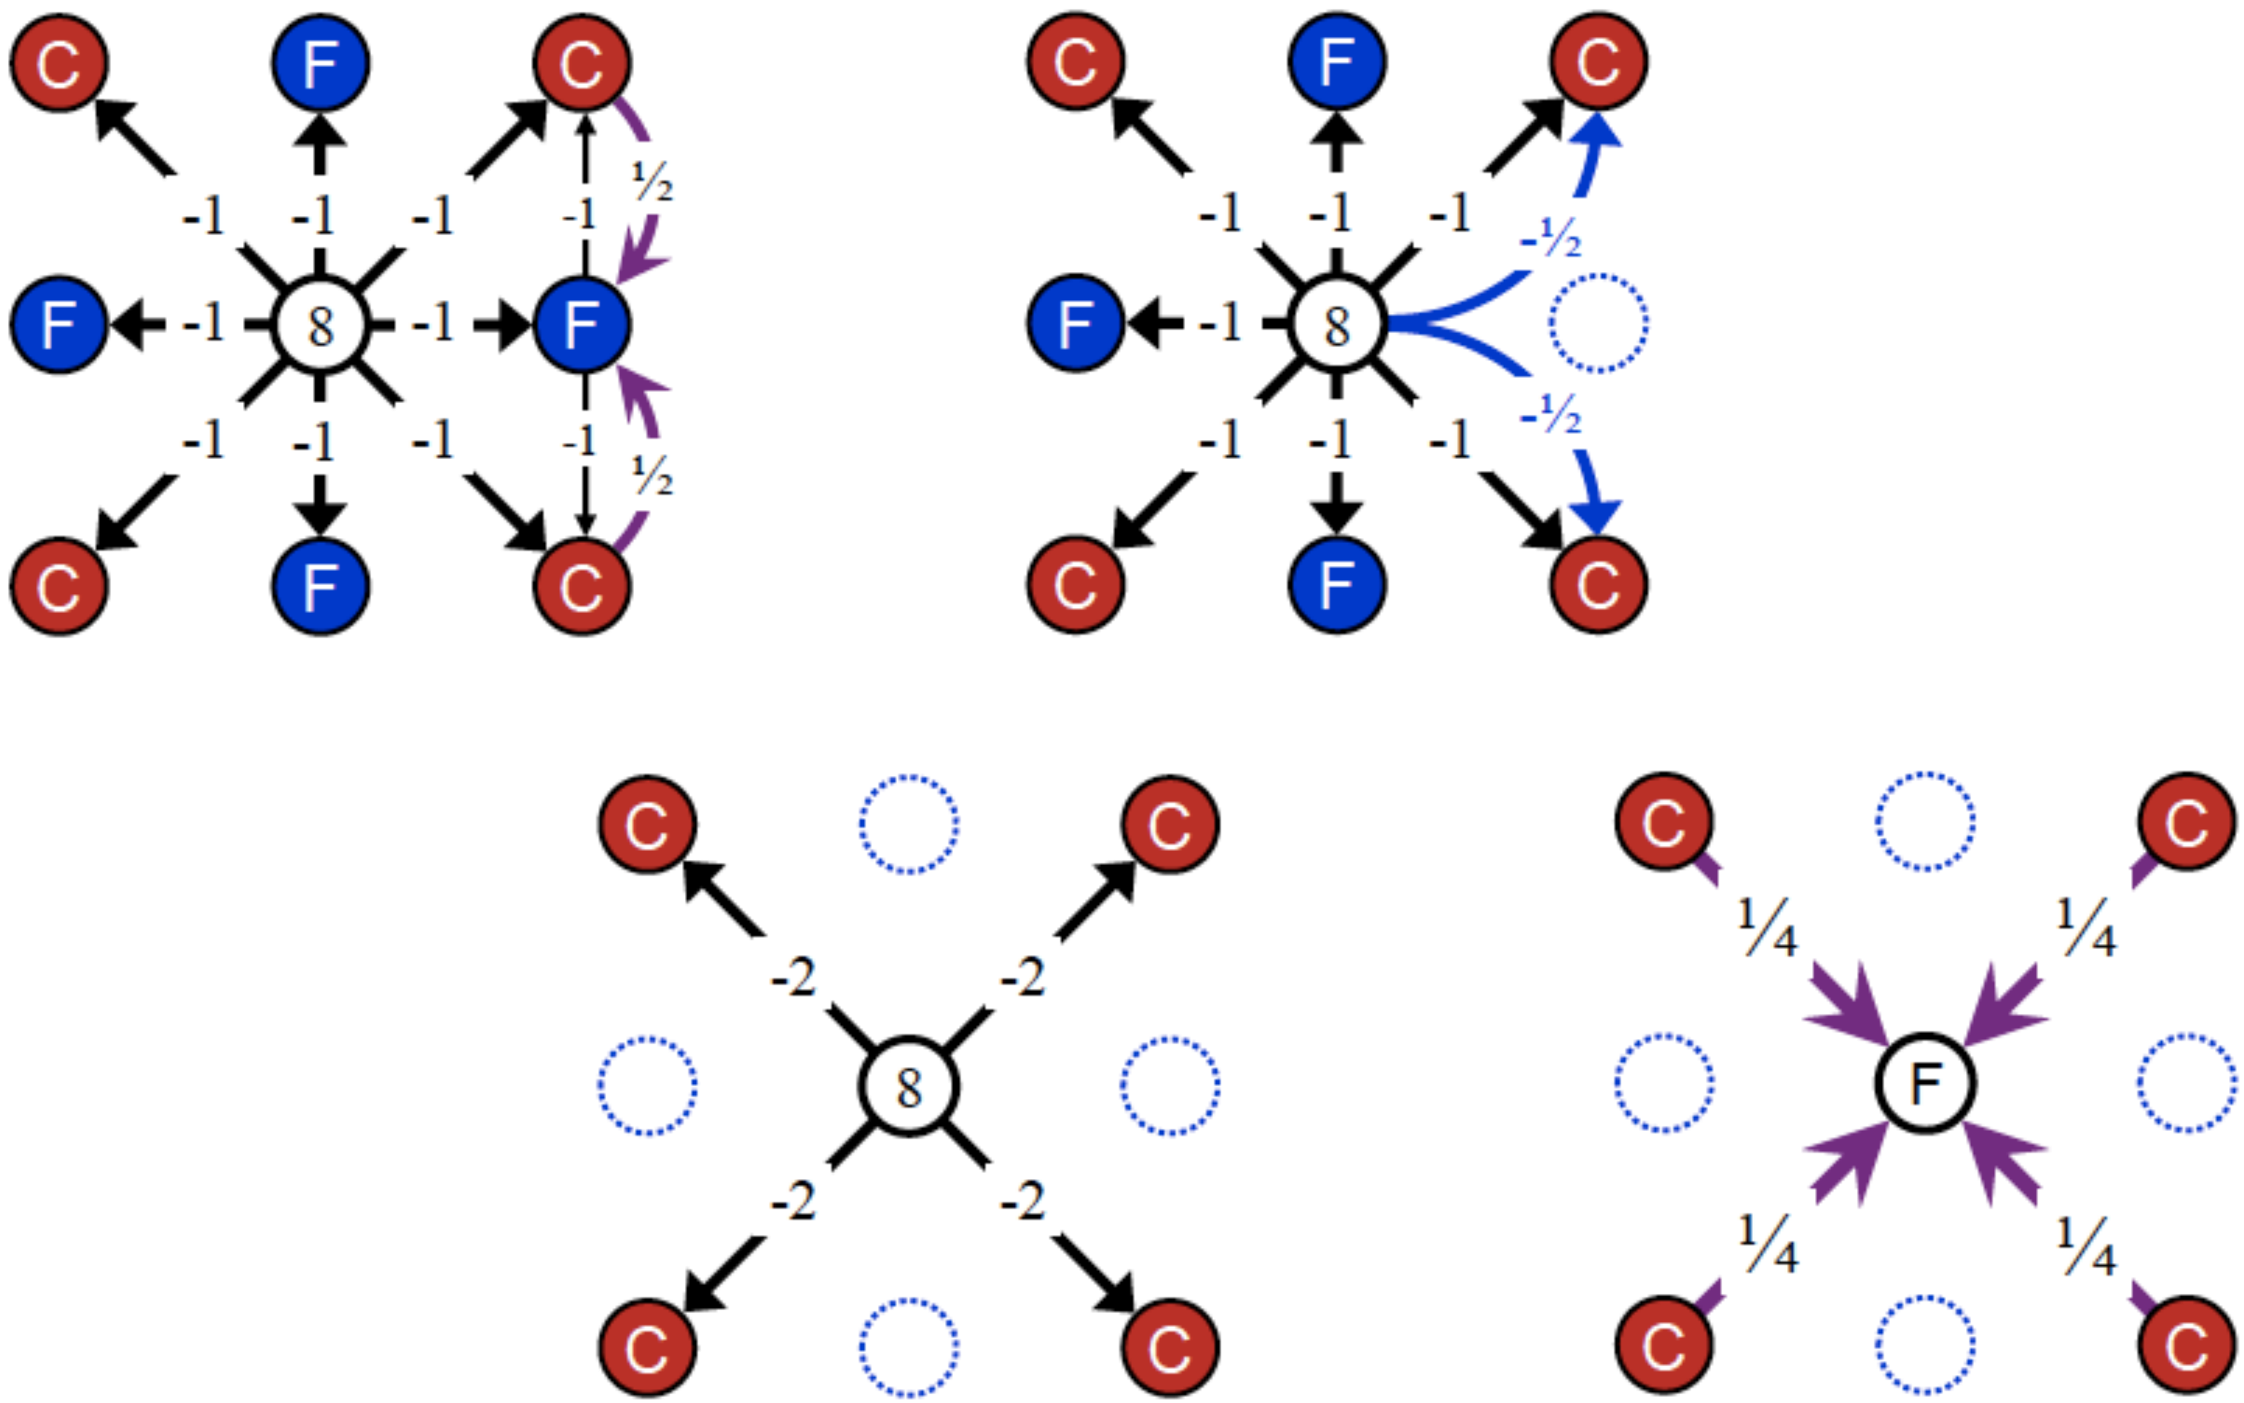

Parallel AMG

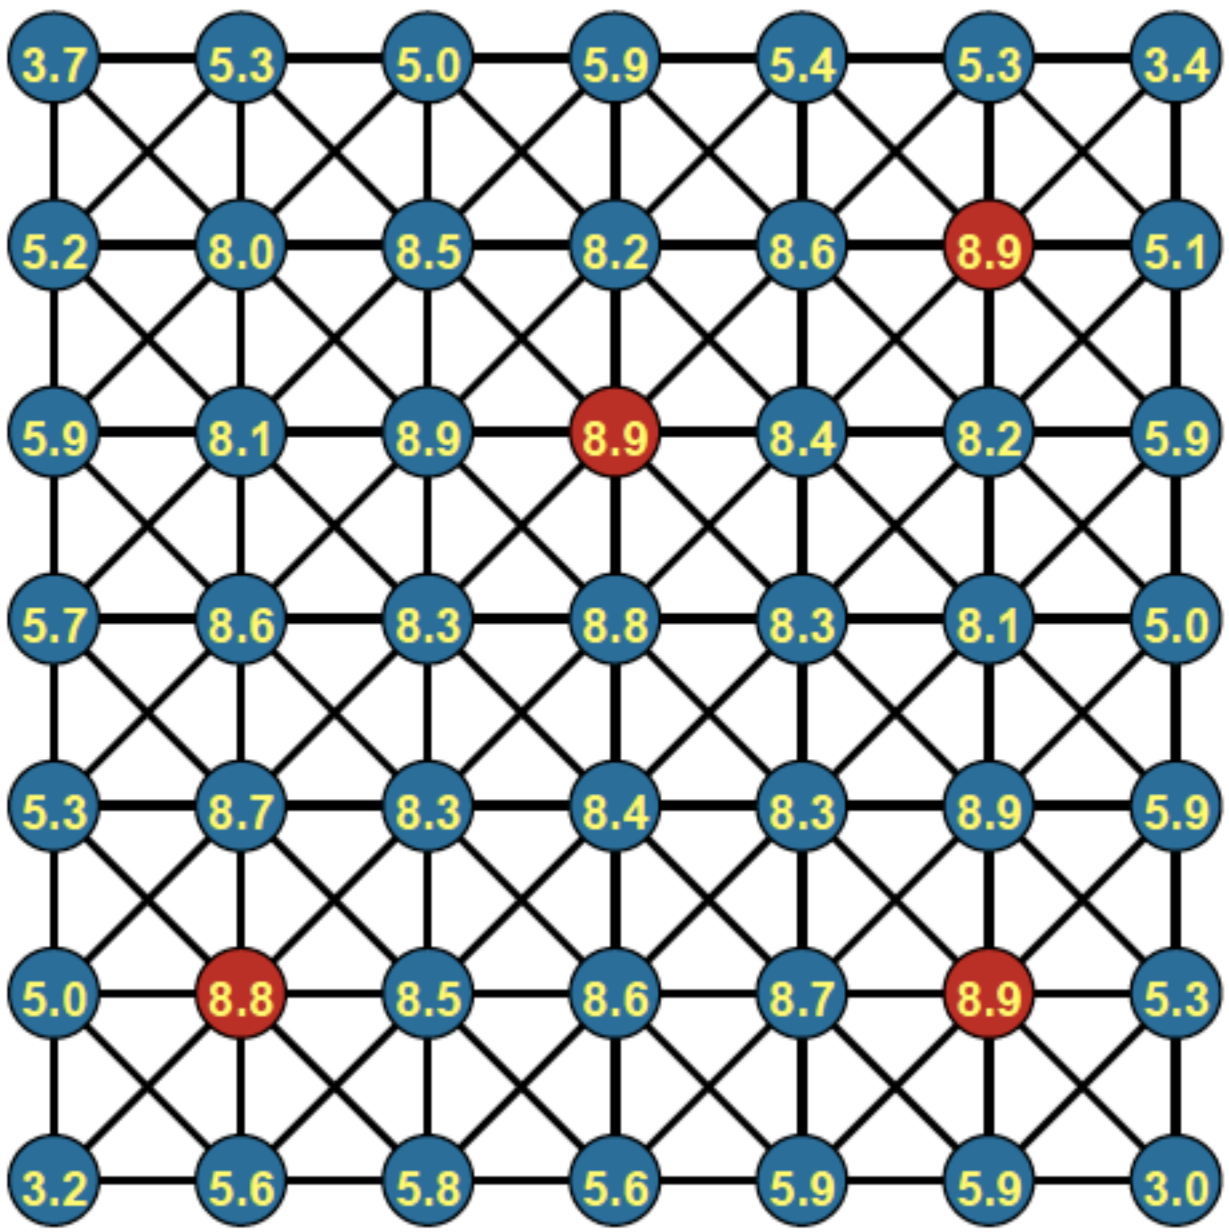

- ➔ select C-pts with maximal measure locally
- ➔ remove neighbor edges
- ➔ update neighbor measures

Parallel AMG

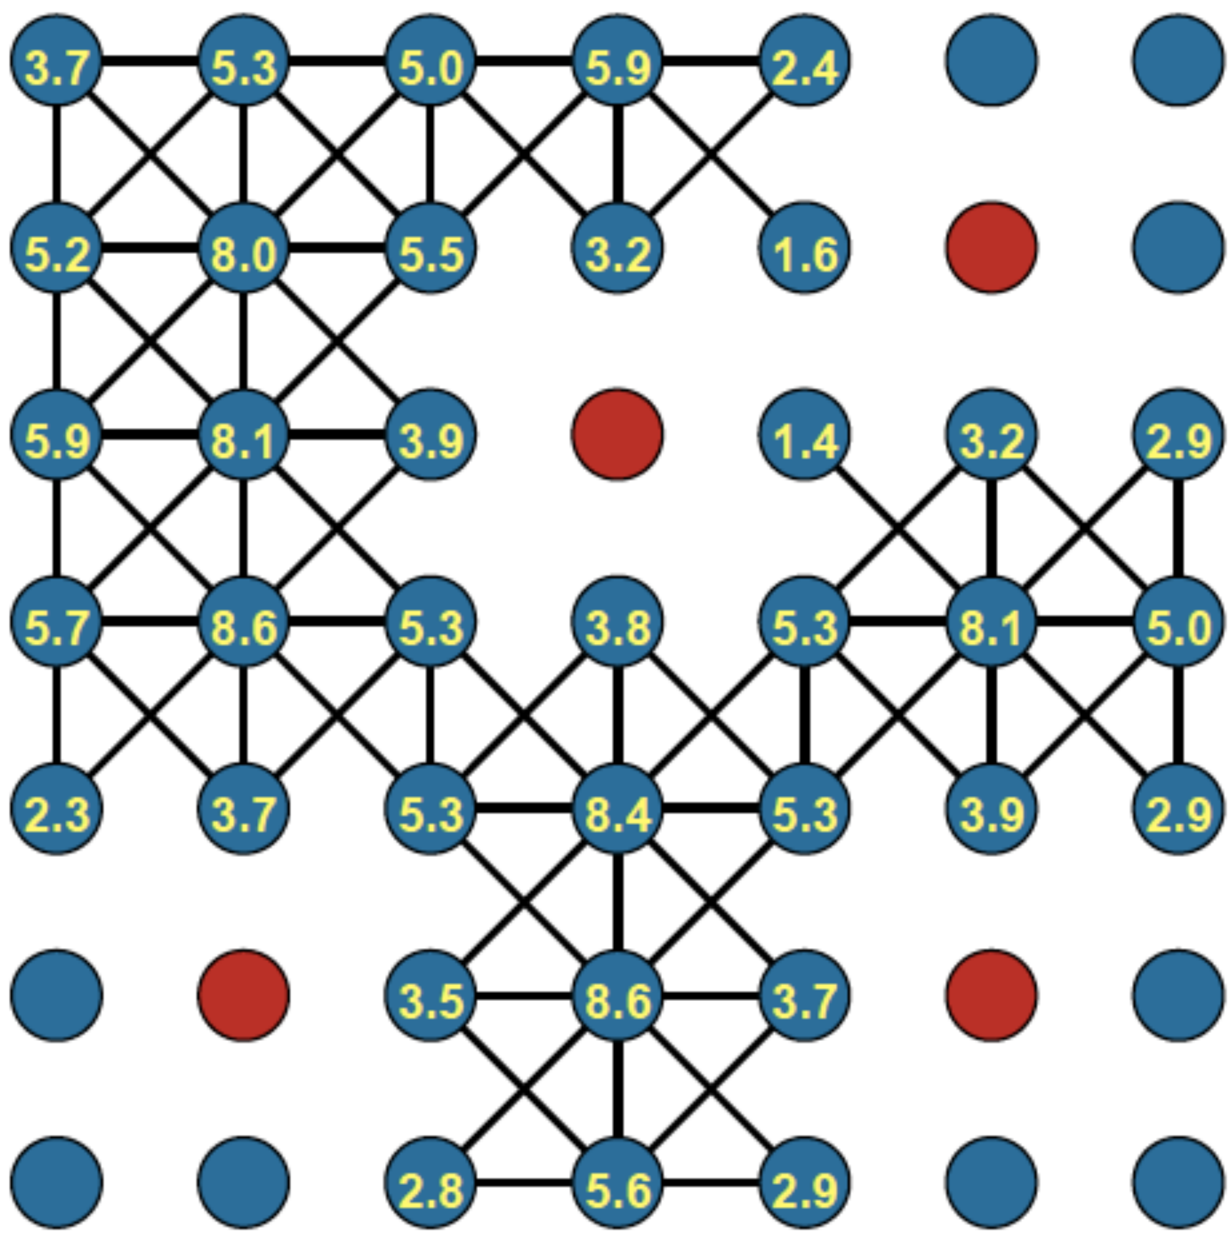

- ➔ **select C-pts with maximal measure locally**
- ➔ **remove neighbor edges**
- ➔ **update neighbor measures**

Parallel AMG

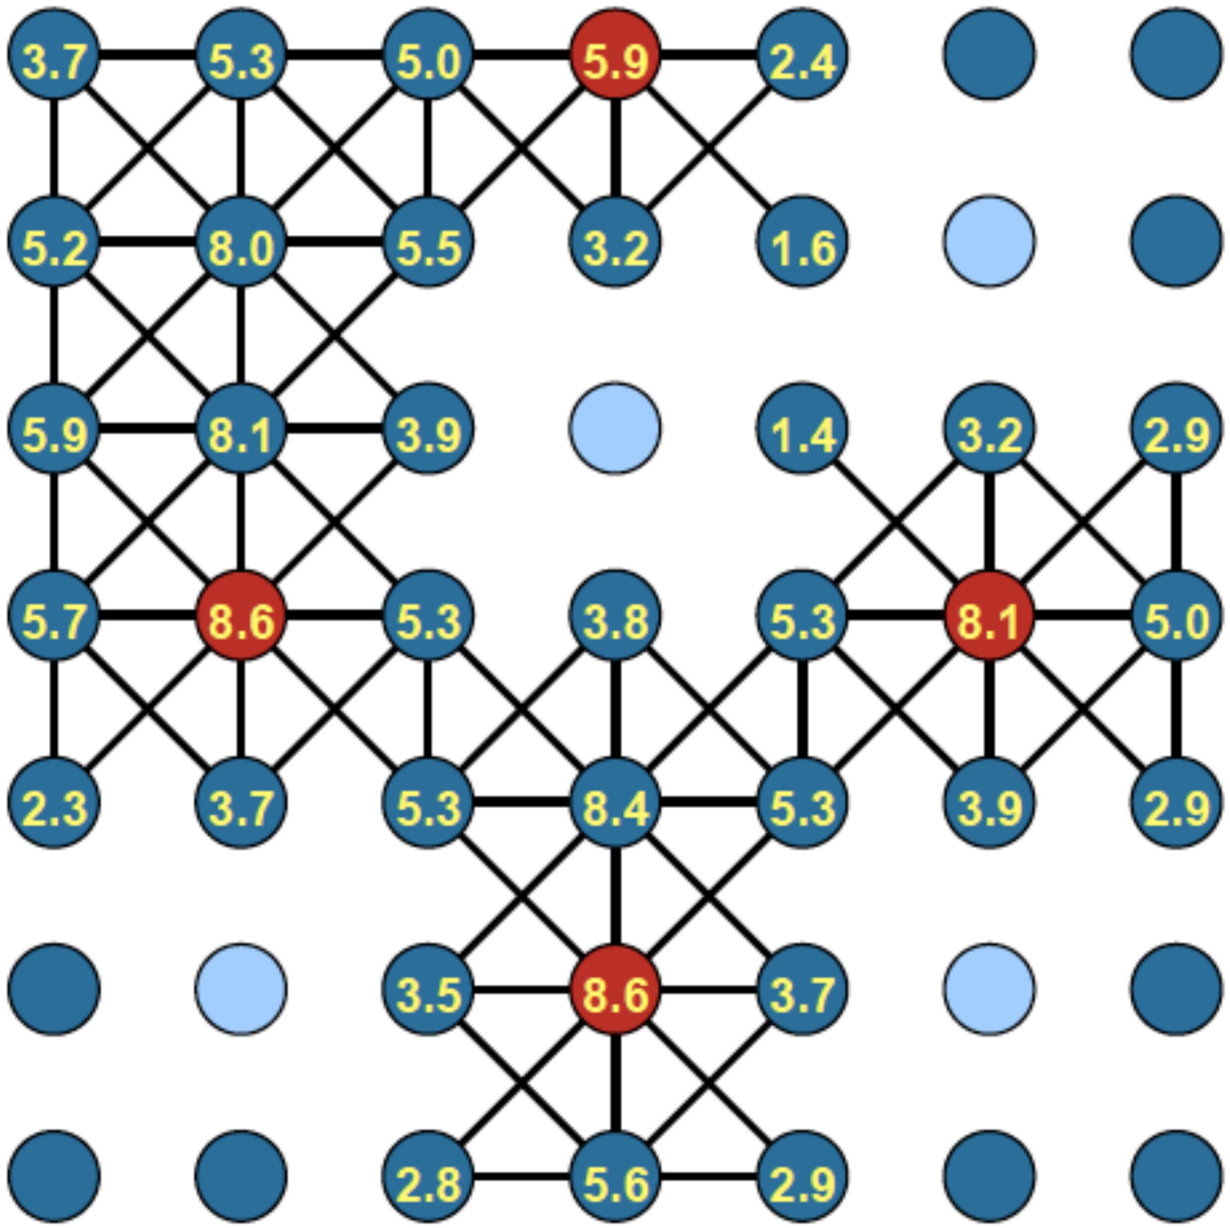

- ➔ **select C-pts with maximal measure locally**
- ➔ **remove neighbor edges**
- ➔ **update neighbor measures**

05/09	Class 9	Dense direct solvers	Understand the principle of LU decomposition and the optimization and parallelization techniques that lead to the LINPACK benchmark.
05/12	Class 10	Dense eigensolvers	Determine eigenvalues and eigenvectors and understand the fast algorithms for diagonalization and orthonormalization.
05/16	Class 11	Sparse direct solvers	Understand reordering in AMD and nested dissection, and fast algorithms such as skyline and multifrontal methods.
05/19	Class 12	Sparse iterative solvers	Understand the notion of positive definiteness, condition number, and the difference between Jacobi, CG, and GMRES.
05/23	Class 13	Preconditioners	Understand how preconditioning affects the condition number and spectral radius, and how that affects the CG method.
05/26	Class 14	Multigrid methods	Understand the role of smoothers, restriction, and prolongation in the V-cycle.
05/30	Class 15	Fast multipole methods, H-matrices	Understand the concept of multipole expansion and low-rank approximation, and the role of the tree structure.

Title