

Local Mode Analysis of Multicolor and Composite Relaxation Schemes

Oren E. Livne ^{†‡} Achi Brandt ^{†§}

August 24, 2002

Dedicated to Antonin Dvôrák for his Symphony No. 8 in G major, Op. 88 (1889)

Abstract

We develop local mode analysis of multicolor (or more generally multicoupling) numerical processes and provide practical procedures for computing their symbols and predicting their convergence and smoothing properties, in the context of multigrid methods. We apply our analysis, including simple numerical examples, to multicolor point relaxation, especially in the context of the novel composite relaxation (A. Brandt, Barriers to achieving textbook multigrid efficiency (TME) in CFD, Tech. Report, Gauss Center Report WI/GC-10,

[†]Department of Computer Science and Applied Mathematics, The Weizmann Institute of Science, Rehovot 76100, Israel. Phone: +972-8-9343545, Fax: +972-8-9344122

[‡]Email address: livneo@wisdom.weizmann.ac.il

[§]Email address: achi@wisdom.weizmann.ac.il

Weizmann Institute, Dec. 1998; O. E. Livne, Multiscale Eigenbasis Algorithms, Ph.D thesis, The Weizmann Institute of Science, Rehovot, Israel, 2000) for products of discrete operators, such as the biharmonic.

Key words. local Fourier mode analysis, multigrid, multicolor relaxation schemes, composite relaxation schemes.

AMS subject classifications. 42A16, 42A38, 65B99, 65G99, 65N06, 65N12, 65N55.

1 Introduction

Fourier analysis, also known as Local Mode Analysis (LMA) in the multigrid community, dates back to the early days of multigrid development [1] and has been well-established since then as the main tool for designing multigrid algorithms and predicting their asymptotic convergence (see for example [1, 2, 4, 14, 16, 15, 17, 18]). The most common variants of LMA are smoothing analysis and two-level multigrid cycle analysis.

Herein, we develop Fourier analysis for general situations that may occur in the practical design of multigrid methods, where many Fourier modes are coupled during the numerical process. Such a process can be a general (k -grid) multigrid cycle (as in [3]; see [17, 18] for a thorough 3-grid Fourier analysis), a complicated relaxation scheme, or any other numerical process with similar coupling characteristics.

In §2 we develop tools for analyzing a general multicoupling process. We restrict ourselves to linear scalar constant-coefficient difference equations on

an infinite uniform grid, which is the standard LMA setting [2, 4, 15, 16]. In the two-level multigrid cycle predictions, we have also assumed a standard 1:2 uniform coarsening in all dimensions, with standard intergrid transfers. However, our results are by no means restricted to such cases, and the quantitatively sharp predictions obtained are applicable to practical problems on general domains, with general boundary conditions and smoothly varying coefficients, provided that supplementary processing is performed at and near boundaries, at a negligible additional cost [3, 4].

The core of Fourier analysis is the computation of the *symbol*, or the amplification matrix of the process under consideration. In §2 we obtain general formulae for the symbol of a multicoupling process, from which its convergence and smoothing rates can be predicted. We extend the discussion to *products* (compositions) of such processes.

The general analysis of §2 is then applied to two important examples, in §3 and §4. Section §3 discusses *multicolor point relaxation schemes*, which generalize the well-known Red-Black (RB) Gauss-Seidel (GS) relaxation [2, 6, 14, 15, 16, 19, 20]. These multicolor schemes couple c^d Fourier modes in d -dimensions, where c is the number of colors. Numerical examples are presented.

Our main application is the analysis of the novel approach of *composite relaxation* for relaxing a *product* of two discrete operators, provided that a relaxation scheme has already been designed for each factor (see §4). Our Fourier analysis and the numerical examples presented shed light on the behavior of

a composite relaxation, and provide guidelines for its proper implementation in practice. In particular it shows that for the biharmonic operator, a suitable composite relaxation is about four times as efficient as the Gauss-Seidel scheme.

Section §5 is devoted to some closing remarks. Although not required in all problems, the general analysis of highly multicoupling processes constitutes an important tool for the design of efficient multigrid algorithms (see for example [10]).

A detailed version of this paper is available as a technical report [11].

2 Fourier Analysis of c -Processes

In this section we develop local mode analysis for general “ c -processes”, which are linear numerical processes that couple c^d Fourier modes in d -dimensional problems. Similar analyses [1, 2, 3, 4, 6, 8, 9, 14, 16, 19, 20] usually refer to the case $c = 2$.

2.1 Notation and Basic Principles

The rigorous foundations of local mode analysis in the context of multigrid are elaborated in [4, 12]. We follow the formulation of [4, 17, 18], start with a simplified framework and discuss its generalizations in §5.

We consider a general linear scalar constant-coefficient system of difference

equations

$$(A^h u^h)(x_\alpha) = \sum_{\gamma \in S} a_\gamma^h u(x_{\alpha+\gamma}) = b^h(x_\alpha), \quad \alpha \in \Omega^h \quad (1)$$

on the infinite uniform grid

$$\Omega^h := \{x_\alpha = \alpha h \mid \alpha = (\alpha_1, \dots, \alpha_d) \in \mathbb{Z}^d\}$$

with meshsize h , where $\{a_\gamma\}_{\gamma \in S}$ are stencil coefficients, and without loss of generality

$$S := [-l_1, l_1] \times \dots \times [-l_d, l_d] \in \mathbb{Z}^d$$

is a compact index set, for some $\{l_j\}_{j=1}^d \in \mathbb{N}$ (a general stencil set may be inscribed in such a d -dimensional cube, by zero-filling the stencil coefficients wherever necessary). We will also use the notation

$$[A^h] = \begin{bmatrix} a_{1,-1}^h & a_{1,0}^h & a_{1,1}^h \\ a_{0,-1}^h & a_{0,0}^h & a_{0,1}^h \\ a_{-1,-1}^h & a_{-1,0}^h & a_{-1,1}^h \end{bmatrix}, \quad S = [-1, 1] \times [-1, 1],$$

with $l_1 = l_2 = 1$, for a compact two-dimensional 9-point stencil $[A^h]$, and similarly for a general stencil. The notation $\theta\alpha$ stands for $\theta \cdot \alpha := \sum_{j=1}^d \theta_j \alpha_j$ whenever there are two adjacent symbols θ, α representing d -vectors. As presented in [17, 18], every bounded infinite-grid function v in

$$\mathcal{B}^h := \{v : \Omega^h \rightarrow \mathbb{C} \mid \|v\| := \langle v, v \rangle^{\frac{1}{2}} < \infty\}$$

is a linear combination of pure Fourier modes of the type

$$\phi(\theta, x) := e^{i\theta x/h} = e^{i\theta\alpha}, \quad x \in \Omega^h, \quad \theta \in \Theta := (-\pi, \pi]^d,$$

which are the eigenfunctions of A^h , and orthonormal with respect to the scaled Euclidean inner product [14, 17, 18]

$$\langle v, w \rangle := \lim_{m \rightarrow \infty} (2m)^{-d} \sum_{|\alpha| \leq m} v(\alpha h) \overline{w(\alpha h)},$$

$$|\alpha| := \max_{1 \leq j \leq d} |\alpha_j|, \quad \forall v, w : \Omega^h \rightarrow \mathbb{C}.$$

In other words, the Fourier space

$$\mathcal{F}^h := \text{span}\{e^{i\theta\alpha} \mid \theta \in \Theta\}$$

contains \mathcal{B}^h .

2.2 Harmonics, Coupling, and c -processes

In the spirit of “classical” LMA, our Fourier analysis provides analytic expressions for the “response” of Fourier modes under some inspected *processes*, that is, linear transformations from \mathcal{F} onto itself. In particular, these expressions can be used, in the context of multigrid methods, to accurately predict the overall convergence of the processes, as well as the convergence of particular sets of modes (e.g., high-frequency modes, related to the “smoothing” property of relaxation schemes; see for example [2, 15, 16]). We will concentrate on processes P that transform any pure mode into a general *periodic modulation* of this mode, with a *mode-independent period* c , where c is a positive integer. For simplicity, we have assumed here that the period c is the same for all the d dimensions; analogous results can be similarly obtained for general periods c_1, \dots, c_d .

Definition 1 (*c*-modulation) *A function $v \in \mathcal{F}$ is called a c -modulation of the mode $e^{i\theta\alpha}$, if c is the smallest natural number such that $v(x_\alpha)e^{-i\theta\alpha}$ is c -periodic in $\alpha_1, \dots, \alpha_d$. That is, if*

$$v(x_{\alpha+c\gamma}) = v(x_\alpha)e^{i\theta c\gamma}, \quad \forall \alpha, \gamma \in \mathbb{Z}^d.$$

Definition 2 (*c*-process) *A process $P : \mathcal{F} \rightarrow \mathcal{F}$ that transforms any mode $e^{i\theta\alpha}$ to a c -modulation of $e^{i\theta\alpha}$ is called a c -process.*

Examples of c -processes are as follows:

1. Any point relaxation for (1) in lexicographic ordering is a 1-process.
2. Any point relaxation for (1) in red-black ordering is a 2-process (in fact, it couples only two modes, not 2^d).
3. A two-level coarse grid correction operation [2, 15] is a 2-process, for a 1 : 2 uniform coarsening in all dimensions.
4. A similar three-level coarse grid correction is a 4-process; a q -level coarse grid correction is a 2^{q-1} -process [3, 17, 18].

In §3 and §4 we consider two special classes of relaxation schemes that are c -processes for a general $c \in \mathbb{N}$.

Definition 3 (*c*-harmonics) *The frequency $\theta' \in \Theta$ is a c -harmonic of $\theta \in \Theta$ if $\theta'_j = \theta_j \pmod{\frac{2\pi}{c}}$ for all $j = 1, \dots, d$. The set of all c -harmonics of θ is*

$$\{\theta^{c,(k)} := (\theta + \eta\tau_k) \pmod{2\pi} \mid k = 0, \dots, c^d - 1\}, \quad \eta := \frac{2\pi}{c},$$

where $\theta(\bmod a) := (\theta_1(\bmod a), \dots, \theta_d(\bmod a))$, and $\tau_k := (\tau_{k1}, \dots, \tau_{kd})$ is the base- c representation of the integer k , i.e.

$$k = \sum_{j=1}^d \tau_{kj} c^{d-j}, \quad 0 \leq \tau_{kj} < c, \quad j = 1, \dots, d.$$

The subspace of linear combinations of the modes $\{e^{i\theta^{c,(k)}\alpha}\}_{k=0}^{c^d-1}$ will be denoted by \mathcal{F}_θ^c .

Lemma 1 (Equivalence Lemma) *A function $v \in \mathcal{F}$ is a c -modulation of the mode $e^{i\theta\alpha}$ if and only if $v \in \mathcal{F}_\theta^c$.*

Proof. Consider first the case $d = 1$, and let $B := (B_0, \dots, B_{c-1})^T \in \mathbb{C}^c$. The function

$$v(x_\alpha) := \begin{cases} B_0, & \alpha(\bmod c) = 0, \\ \vdots & \vdots \\ B_{c-1}, & \alpha(\bmod c) = c-1, \end{cases} \quad \alpha \in \mathbb{Z} \quad (2)$$

can be expressed as

$$v(x_\alpha) = \sum_{r=0}^{c-1} A_r e^{i\theta^{c,(k)}\alpha}, \quad \alpha \in \mathbb{Z}, \quad (3)$$

where $A := (A_0, \dots, A_{c-1})^T \in \mathbb{C}^c$ satisfies [11, §2.2]

$$\sum_{r=0}^{c-1} A_r e^{i\eta r \alpha} = B_\alpha, \quad \alpha = 0, \dots, c-1. \quad (4)$$

This is a $c \times c$ linear system of equations

$$\mathcal{V}(1, \xi, \dots, \xi^{c-1})A = B, \quad \xi := e^{i\eta}, \quad (5)$$

where

$$\mathcal{V}(x_0, \dots, x_{c-1}) := \begin{pmatrix} 1 & x_0 & \dots & x_0^{c-1} \\ \vdots & \vdots & \dots & \vdots \\ 1 & x_{c-1} & \dots & x_{c-1}^{c-1} \end{pmatrix}. \quad (6)$$

The Vandermonde matrix $\mathcal{V}(1, \xi, \dots, \xi^{c-1})$ is regular, since all the numbers $\{\xi^r\}_{r=0}^{c-1}$ are different [13]. Moreover, [11, Lemma 3] implies

$$A = \frac{1}{c} \mathcal{V}(\xi^c = 1, \xi^{c-1}, \dots, \xi) B. \quad (7)$$

Thus, there is a 1:1 correspondence between the forms (2) and (3) of the function v . The d -dimensional case can be proved recursively (see [11] for a complete proof); the matrix (6) in (5) is replaced by a matrix $\mathcal{V}^{(d)}$ composed of *blocks* of Vandermonde matrices. \square

Corollary 1 *A c -process P keeps the c -dimensional subspace \mathcal{F}_θ^c invariant.*

Proof. By Lemma 1, there exist functions $\{A_k(\theta)\}_{k=0}^{c^d-1}$ such that

$$(Pv)(x_\alpha) = \sum_{k=0}^{c^d-1} A_k(\theta) e^{i\theta^{c,(k)}\alpha}, \quad \forall \alpha \in \mathbb{Z}^d. \quad (8)$$

Thus, the P -image of a general element

$$v(x_\alpha) = \sum_{k=0}^{c^d-1} C_k e^{i\theta^{c,(k)}\alpha}$$

of \mathcal{F}_θ^c is

$$\begin{aligned} (Pv)(x_\alpha) &= \sum_{k=0}^{c^d-1} \sum_{m=0}^{c^d-1} C_k A_m \left(\theta^{c,(k)} \right) e^{i(\theta^{c,(k)})^{c,(m)}\alpha} \\ &= \sum_{k=0}^{c^d-1} \sum_{m=0}^{c^d-1} C_k A_m \left(\theta^{c,(k)} \right) e^{i\theta^{c,(\frac{k+m}{c})}\alpha}, \end{aligned}$$

since

$$(\theta^{c,(k)})^{c,(m)} = (\theta + \eta(\tau_k + \tau_m)) \pmod{2\pi} = (\theta + \eta\underline{\tau_{k+m}}) \pmod{2\pi} = \theta^{c,\underline{(k+m)}},$$

where $\underline{a} := a \pmod{c^d}$. We interchange the order of summation and substitute $k' := k + m$ to obtain

$$\begin{aligned} (Pv)(x_\alpha) &= \sum_{m=0}^{c^d-1} \left[\sum_{k'=m}^{c^d+m-1} C_{k'-m} A_m \left(\theta^{c,(k'-m)} \right) \right] e^{i\theta^{c,(k')} \alpha} \\ &= \sum_{m=0}^{c^d-1} \left[\sum_{k'=0}^{c^d-1} C_{\underline{k'-m}} A_m \left(\theta^{c,\underline{(k'-m)}} \right) \right] e^{i\theta^{c,(k')} \alpha}. \end{aligned}$$

Interchanging the order of summation again and substituting $m' := k' - m$ yields

$$(Pv)(x_\alpha) = \sum_{k'=0}^{c^d-1} \left[\sum_{m=0}^{c^d-1} C_{m'} A_{\underline{k'-m'}} \left(\theta^{c,(m')} \right) \right] e^{i\theta^{c,(k')} \alpha} =: \sum_{k'=0}^{c^d-1} \tilde{C}_k e^{i\theta^{c,(k')} \alpha},$$

where

$$\begin{pmatrix} \tilde{C}_0 \\ \tilde{C}_1 \\ \vdots \\ \tilde{C}_{c^d-1} \end{pmatrix} = \begin{pmatrix} A_0(\theta^{c,(0)}) & A_{c^d-1}(\theta^{c,(1)}) & \dots & A_1(\theta^{c,(c^d-1)}) \\ A_1(\theta^{c,(0)}) & A_0(\theta^{c,(1)}) & \dots & A_2(\theta^{c,(c^d-1)}) \\ \vdots & \vdots & \dots & \vdots \\ A_{c^d-1}(\theta^{c,(0)}) & A_{c^d-2}(\theta^{c,(1)}) & \dots & A_0(\theta^{c,(c^d-1)}) \end{pmatrix} \begin{pmatrix} C_0 \\ C_1 \\ \vdots \\ C_{c^d-1} \end{pmatrix}.$$

Hence, $Pv \in \mathcal{F}_\theta^c$. ■

We have shown that a c -process *couples sets of c^d modes*, which are additive groups (in coordination with the required symmetry of all the properties of c -processes discussed previously).

2.3 The symbol of a c -process

Analogously to the definitions of [2, 4, 17, 18, 19, 20], the *symbol* \hat{P} of a c -process P given by (8) denotes its amplification matrix of any $v \in \mathcal{F}_\theta^c$. Specifically,

$$\hat{P}(\theta) := \begin{pmatrix} A_0(\theta^{c,(0)}) & A_{c^d-1}(\theta^{c,(1)}) & \dots & A_1(\theta^{c,(c^d-1)}) \\ A_1(\theta^{c,(0)}) & A_0(\theta^{c,(1)}) & \dots & A_2(\theta^{c,(c^d-1)}) \\ \vdots & \vdots & \dots & \vdots \\ A_{c^d-1}(\theta^{c,(0)}) & A_{c^d-2}(\theta^{c,(1)}) & \dots & A_0(\theta^{c,(c^d-1)}) \end{pmatrix}. \quad (9)$$

In the context of multigrid, we are usually interested in computing

$$\mu_P := \sup_{\theta \in \Theta} \rho(\hat{P}(\theta)) \quad (10)$$

$$\bar{\mu}_{P,\nu} := \left(\sup_{\theta \in \Theta} \rho(C^{(c)}(\theta) \hat{P}(\theta)^\nu) \right)^{\frac{1}{\nu}}, \quad (11)$$

where

$$C^{(c)}(\theta) := \text{diag} \{c(\theta^{c,(0)}), \dots, c(\theta^{c,(c^d-1)})\}, \quad c(\theta) := \begin{cases} 0, & |\theta| \leq \frac{\pi}{2}, \\ 1, & \text{otherwise.} \end{cases}$$

The quantity (10) is the *asymptotic convergence factor* of P . (11) represents the *smoothing factor* per sweep, for ν consecutive sweeps of P ; that is, the reduction of high-frequency modes that cannot be represented on a coarse grid with meshsize $2h$. $C^{(c)}(\theta)$ is a projection on this set, and reflects an ideal coarse grid correction operator. Thus, $\bar{\mu}_{P,\nu}$ predicts the ideal performance expected from a two-level multigrid cycle containing ν relaxation sweeps, when P represents the relaxation scheme at the fine level (see [2, 15, 16]).

We will also use in 4.3 the following analogous definitions for the L_2 norm.

Definition 4 (L_2 -factors) *The L_2 -convergence and L_2 -smoothing factors of a c -process P are defined by (see [4])*

$$M_P := \sup_{\theta \in \Theta} \|\hat{P}(\theta)\|_2 \quad (12)$$

$$\overline{M}_{P,\nu} := \left(\sup_{\theta \in \Theta} \|C^{(c)}(\theta) \hat{P}(\theta)^\nu\|_2 \right)^{\frac{1}{\nu}}, \quad (13)$$

where $\|A\|_2 := (\rho(A^T A))^{\frac{1}{2}}$ is the maximal singular value of A .

We will also briefly denote $\overline{M}_P := \overline{M}_{P,1}$. Note that $M_P \geq \mu_P$, $\overline{M}_{P,\nu} \geq \overline{\mu}_{P,\nu}$, and for a uni-coupling process, e.g., lexicographic point relaxation ($c = 1$), $M_P = \mu_P$, $\overline{M}_{P,\nu} = \overline{\mu}_{P,\nu}$.

The supremum in (10),(11) can be analytically computed for some special cases (see [19, 20]). However, in general, it is computed numerically by sampling the Fourier space Θ by a sufficiently fine lattice, with meshsize δ , and explicitly computing \hat{P} over that lattice. δ need not be too small, as noted in [3]; all the numerical results presented here use $\delta = \frac{1}{32}$, that gives an adequate resolution of the symbols discussed. However, the total number of lattice points is $O\left(\left(\frac{2\pi}{\delta}\right)^d\right)$, which can be quite large for $d = 3$, for instance. When the power iteration is used to compute $\rho\left(\hat{P}(\theta)\right)$ [3, §10.3], the total cost of computing (10) is $O\left(\left(\frac{2\pi}{\delta}\right)^d c^{3d} = \left(\frac{2\pi c^3}{\delta}\right)^d\right)$, and a similar, somewhat larger cost is required for computing (11) for the ν -values of interest, which are small [2]. This cost can be reduced using the following result.

Lemma 2 (Symmetricity Lemma) *For any c -process P , $\rho\left(\hat{P}(\theta)\right)$ and $\rho\left(C^{(c)}(\theta) \hat{P}(\theta)^\nu\right)$ are $\frac{2\pi}{c}$ -periodic in $\theta_1, \dots, \theta_d$.*

The lemma follows by observing that the sets of c -harmonics of $(\theta_1, \dots, \theta_d)$ and of $(\theta_1 + \frac{2\pi}{c}, \dots, \theta_d)$ are identical, up to a permutation. See [11] for the complete proof. \square

As a result, we can restrict the supremum search in (10) and (11) to the basic cell $\Theta_c := (-\frac{\pi}{c}, \frac{\pi}{c}]^d$, thereby reducing the cost of computing $\mu_P, \bar{\mu}_{P,\nu}$ from $O\left(\left(\frac{2\pi c^3}{\delta}\right)^d\right)$ to $O\left(\left(\frac{2\pi c^2}{\delta}\right)^d\right)$. Even for practical values (e.g., $c \leq 4, d \leq 3$), this improvement may prove valuable. It is certainly significant when $\mu_P, \{\bar{\mu}_{P,\nu}\}_\nu$ need to be repeatedly computed.

2.4 Products of c -Processes

Often we wish to analyze the behavior of a *composition* of processes, that is, the compound action of consecutive processes. For instance, a two-level multigrid cycle consists of several pre-relaxation sweeps, followed by a coarse grid correction, and finally, some post-relaxations. Another important example occurs in the study of composite relaxations (see §4). Here we analyze a *product* $P = P_1 P_2$ of a general c_1 -process P_1 and a general c_2 -process P_2 .

Theorem 1 (Product Color Theorem) *P is an $\text{lcm}(c_1, c_2)$ -process.*

Here $c := \text{lcm}(c_1, c_2)$ is the least common multiple of c_1, c_2 .

Proof. A mode $e^{i\theta\alpha}$ is transformed by P_2 to an element $v_2 \in \mathcal{F}_\theta^{c_2}$. P_1 is linear, thus $P_1 v_2$ is a linear combination of $P_1\left(\exp(i\theta^{c_2, (k)}\alpha)\right)$ for all $\theta^{c_2, (k)} \in \mathcal{F}_\theta^{c_2}$.

Since P_1 is a c_1 -process, $P_1 \left(\exp(i\theta^{c_2, (k)} \alpha) \right) \in \mathcal{F}_{\theta^{c_2, (k)}}^{c_1}$, thus

$$P_1 v_2 \in \bigcup_{k \in \mathcal{F}_\theta^{c_2}} \mathcal{F}_{\theta^{c_2, (k)}}^{c_1}.$$

The last set can easily be shown to equal $\mathcal{F}_\theta^{\text{lcm}(c_1, c_2)}$ (see for example Fig. 1).

■




Figure 1: An example of confounded coupling ($d = 1, c_2 = 6, c_1 = 4$). Given an original frequency θ , we denote a coupled frequency θ' by the integer $c_1 c_2 ((\theta' - \theta) / (2\pi) \pmod{2\pi})$. θ (denoted by 0 on the left) is first coupled by P_2 (the solid arrows) with the six harmonics $\theta + (0, 4, 8, 12, -8, -4) \times \frac{2\pi}{c_1 c_2}$. Then, each of these frequencies is coupled (through the P_1 process, denoted by the dashed arrows) with four harmonics. However, there are only $\text{lcm}(c_1, c_2) = 12$ different harmonics on the right matrix.

The symbol of P is obtained by “embedding” \hat{P}_1, \hat{P}_2 into a common matrix size, and then multiplying the embedded symbols. For the example of Fig. 1, \hat{P} is computed as follows:

$$\mathcal{P}_{1,KM} := \text{diag} \left\{ \hat{P}_1 \left(\theta^{12,(0)} \right)_{KM}, \hat{P}_1 \left(\theta^{12,(1)} \right)_{KM}, \hat{P}_1 \left(\theta^{12,(2)} \right)_{KM} \right\}, \quad (14)$$

$$K, M = 0, \dots, 3,$$

$$\check{P}_1(\theta) := \begin{pmatrix} \mathcal{P}_{1,00} & \mathcal{P}_{1,01} & \mathcal{P}_{1,02} & \mathcal{P}_{1,03} \\ \mathcal{P}_{1,10} & \mathcal{P}_{1,11} & \mathcal{P}_{1,12} & \mathcal{P}_{1,13} \\ \mathcal{P}_{1,20} & \mathcal{P}_{1,21} & \mathcal{P}_{1,22} & \mathcal{P}_{1,23} \\ \mathcal{P}_{1,30} & \mathcal{P}_{1,31} & \mathcal{P}_{1,32} & \mathcal{P}_{1,33} \end{pmatrix}; \quad (15)$$

$$\mathcal{P}_{2,KM} := \text{diag} \left\{ \hat{P}_2 \left(\theta^{12,(0)} \right)_{KM}, \hat{P}_2 \left(\theta^{12,(1)} \right)_{KM} \right\}, \quad (16)$$

$$K, M = 0, \dots, 5,$$

$$\check{P}_2(\theta) := \begin{pmatrix} \mathcal{P}_{2,00} & \mathcal{P}_{2,01} & \mathcal{P}_{2,02} & \mathcal{P}_{2,03} & \mathcal{P}_{2,04} & \mathcal{P}_{2,05} \\ \mathcal{P}_{2,10} & \mathcal{P}_{2,11} & \mathcal{P}_{2,12} & \mathcal{P}_{2,13} & \mathcal{P}_{2,14} & \mathcal{P}_{2,15} \\ \mathcal{P}_{2,20} & \mathcal{P}_{2,21} & \mathcal{P}_{2,22} & \mathcal{P}_{2,23} & \mathcal{P}_{2,24} & \mathcal{P}_{2,25} \\ \mathcal{P}_{2,30} & \mathcal{P}_{2,31} & \mathcal{P}_{2,32} & \mathcal{P}_{2,33} & \mathcal{P}_{2,34} & \mathcal{P}_{2,35} \\ \mathcal{P}_{2,40} & \mathcal{P}_{2,41} & \mathcal{P}_{2,42} & \mathcal{P}_{2,43} & \mathcal{P}_{2,44} & \mathcal{P}_{2,45} \\ \mathcal{P}_{2,50} & \mathcal{P}_{2,51} & \mathcal{P}_{2,52} & \mathcal{P}_{2,53} & \mathcal{P}_{2,54} & \mathcal{P}_{2,55} \end{pmatrix}; \quad (17)$$

$$\hat{P}(\theta) = \check{P}_1(\theta)\check{P}_2(\theta). \quad (18)$$

2.4.1 Convergence and Smoothing

An important question is whether the convergence and smoothing of the product process may be apriori estimated in terms of the convergence and smoothing of each factor. Although general upper bounds can be obtained, as shown in [11], they may be far above the realistic values that can be directly computed using the procedure described in §2.3 for the symbol (18).

An excellent example is provided by the two-level multigrid cycle: although both the smoother (P_1) and the coarse grid correction operator (P_2) usually have $\mu_{P_1}, \mu_{P_2} \approx 1$, their composition has far better convergence than $\mu_{P_1}\mu_{P_2}$, which is (roughly) the general upper bound.

Moreover, by using the sparse structure of $\check{P}_1(\theta), \check{P}_2(\theta)$ in the power iteration, we can effectively reduce the cost of computing the convergence and smoothing of $P = P_1P_2$, as explained in [3, §10.3].

For a detailed discussion and numerical illustrations of the upper bounds for the convergence and smoothing of P_1P_2 , see [11].

Having acquired the general tools for Fourier analysis of c -processes, we can now apply them to two important classes of c -processes: multicolor point relaxations (§3) and composite relaxations (§4), and discuss some practical implications of the theoretical analyses.

3 Multicolor Relaxations

To show an example of analyzing multicoupling processes, we will analyze a particular generalization of the well known Red-Black Gauss-Seidel (GS) relaxation scheme [2], namely, the multicolor GS scheme.

3.1 Definition

A c -color point GS relaxation sweep (denoted GS_c) for (1) is defined as a GS sweep [7] in which the variables are updated in the following order:

For $r = 0$ to $c - 1$,

Color sweep r : Update all $u(x_\alpha)$ for which $\Sigma_\alpha \pmod{c} = r$.

End For

In these sweeps, $\Sigma_\alpha := \sum_{j=1}^d \alpha_j$ for $\alpha \in \mathbb{Z}^d$.

The motivation for using such “colored ordering” appears, for instance, in [2, 15, 16], where a two-color (red-black) GS proves to be a better smoother for the standard discretization of the Laplace operator than the lexicographic GS: points of the same color do not appear in each other’s equation; hence a high-frequency Fourier mode in the error is effectively converted to a low-frequency mode. However, the red-black (RB) ordering is not always better than the lexicographic ordering.

In general, stencils at points of the same color may *overlap*; therefore, the result of a multicolor relaxation also depends on the ordering *within* each color (unless c is chosen large enough, see §3.2). For simplicity, we will assume Jacobi ordering within each color (i.e., all points of the same color are simultaneously updated). Of course, the definition of the color sweeps and their internal ordering may be organized differently for various practical problems (e.g., [11] considers lexicographic ordering within each color); here we consider a particular setting, to demonstrate our analysis tools.

3.2 The Symbol

The underlying characteristic of GS_c is summarized in the following result.

Theorem 2 GS_c is a c -process.

Proof. Fix θ and let $v(x_\alpha) := e^{i\theta\alpha}$. After applying the color sweep 0 to v , we obtain a new grid-function $v^{(0)}$ that satisfies

$$e^{-i\theta\alpha}v^{(0)}(x_\alpha) = \begin{cases} J_0(\theta), & \Sigma_\alpha(\bmod c) = 0, \\ 1, & \text{otherwise,} \end{cases} \quad J_0(\theta) := j(\theta, \{D_{0,\alpha}\}_{\alpha \in S}),$$

where

$$j(\theta, \{d_\alpha\}_{\alpha \in S}) := -\frac{\sum_{\alpha \in S \setminus \{0\}} d_\alpha a_\alpha^h e^{i\theta s}}{a_0^h}, \quad \underline{0} := \underbrace{(0, \dots, 0)}_d, \quad D_{0,\alpha} = 1, \forall \alpha \in S.$$

By following the results obtained after subsequent color sweeps, it can be shown [11] that the final result $v^{(c-1)}$ obtained after GS_c is

$$e^{-i\theta\alpha}v^{(c-1)}(x_\alpha) = \begin{cases} J_0(\theta), & \Sigma_\alpha(\bmod c) = 0, \\ J_1(\theta), & \Sigma_\alpha(\bmod c) = 1, \\ \vdots & \vdots \\ J_{c-1}(\theta), & \Sigma_\alpha(\bmod c) = c-1, \end{cases} \quad (19)$$

$$J_r(\theta) := j(\theta, \{D_{r,\alpha}\}_{\alpha \in S}), \quad D_\alpha^{(r)} = \begin{cases} J_{r-1}(\theta), & \Sigma_\alpha(\bmod c) = c-1, \\ D_{r-1,\sigma(\alpha)}, & \text{otherwise,} \end{cases}$$

where $\sigma(\alpha) := ((\alpha_1 + 1) \bmod c, \alpha_2, \dots, \alpha_d)$ and $D_{-1,\alpha} := 0$ for all $\alpha \in S$ (see Fig. 2).

Equation (19) implies that the result of the GS_c action on $e^{i\theta\alpha}$ is a c -modulation of $e^{i\theta\alpha}$. Let

$$\begin{aligned} \mathcal{K}_k &:= v^{(c-1)}(x_{\alpha+\gamma})e^{-i\theta\alpha}, \quad \forall \gamma = (\gamma_1, \dots, \gamma_d), \\ 0 \leq \gamma_j &< c, \quad 0 < k := \sum_{j=0}^{c-1} \gamma_j c^{d-j} < c^d - 1, \end{aligned} \quad (20)$$




Figure 2: A two-dimensional example ($d = 2$) of a $c = 3$ -color relaxation on a 5×5 stencil. After the full sweep on a Fourier mode $e^{i\theta\alpha}$, we obtain $v^{(2)}$; the figure presents the values $v^{(2)}(x_\alpha)e^{-i\theta\alpha}$ on the infinite spatial grid $\{\alpha = (\alpha_1, \alpha_2) \in \mathbb{Z}^2\}$. The hexagon designates a reference point $\alpha = (0, 0)$. For example, in the last color sweep, the point $(0, 2)$ (denoted by a circle) was updated using $v^{(1)}$ values from the relevant stencil S inscribed in the large square. Note that $v^{(2)}$ is a c -modulation whose periodicity cell is marked by the dashed square. The values $\{\mathcal{K}_k\}_{k=0}^{c^d-1}$ are computed from this basic cell by (20).

which are α -independent values (but depend on θ). By the equivalence lemma and (9) (see §2.2), GS_c is a c -process, whose symbol is

$$G\hat{S}_c(\theta) := (\mathcal{V}^{(d)})^{-1} \begin{pmatrix} \mathcal{K}_0(\theta^{c,(0)}) & \mathcal{K}_{c^d-1}(\theta^{c,(1)}) & \dots & \mathcal{K}_1(\theta^{c,(c^d-1)}) \\ \mathcal{K}_1(\theta^{c,(0)}) & \mathcal{K}_0(\theta^{c,(1)}) & \dots & \mathcal{K}_2(\theta^{c,(c^d-1)}) \\ \vdots & \vdots & \dots & \vdots \\ \mathcal{K}_{c^d-1}(\theta^{c,(0)}) & \mathcal{K}_{c^d-2}(\theta^{c,(1)}) & \dots & \mathcal{K}_0(\theta^{c,(c^d-1)}) \end{pmatrix}, \quad (21)$$

where $\mathcal{V}^{(d)}$ is a block Vandermonde matrix (see [11, §2] for its definition). Thus, we can compute (10),(11),(12),(13) for any GS_c using (21), given the stencil coefficients $\{a_\gamma\}_{\gamma \in S}$. ■

A proper choice of the number of colors c may be an important feature of relaxation design, although in many problems $c \leq 2$ satisfies all practical requirements, and does not require special divisibility properties on the number of gridpoints to avoid boundary effects.

For some problems, however, enlarging c so that points of the same color do not appear in each other's equation, may significantly improve the smoothness properties of the relaxation scheme (see for example [6, 10]). Another helpful tool in relaxation design is the choice of ordering *within each color sweep*, e.g., Jacobi, colored GS, or even distributive GS. A detailed example of optimizing a multicolor relaxation scheme using these strategies appears in [10].

3.3 Numerical Examples

As a simple illustration, we applied our analysis to multicolor GS relaxation for various colors, for the two operators

$$\Delta^h = \frac{1}{h^2} \begin{bmatrix} & & 1 & & \\ & 1 & & & \\ 1 & -4 & 1 & & \\ & & & & \\ & & & & 1 \end{bmatrix}; \quad (\Delta^2)^h = \frac{1}{h^4} \begin{bmatrix} & & & & 1 & & \\ & & & & 2 & -8 & 2 \\ & & & 1 & -8 & 20 & -8 & 1 \\ & & & & 2 & -8 & 2 & \\ & & & & & & & 1 \end{bmatrix} \quad (22)$$

in $d = 2$ dimensions. For each relaxation, we computed the smoothing factor (for $\nu = 1, 2$ consecutive sweeps), and the asymptotic convergence factors per work unit of a two-level multigrid cycle containing $\nu = 1, 2$ fine grid relaxation sweeps, with linear and cubic intergrid transfers (see for instance [4, §3.2] for the definition of the symbol of a two-level multigrid cycle). The convergence factor in all cases was $\mu_P = 1$. The results are summarized in Table 1.

The results indicate that $c = 1$ is the worst case for both operators, and there was no improvement when c was enlarged to more than 3. The case of $c = 2$ has the best performance ($c = 3$ yields similar results). Note that the smoothing factor always predicts quite accurately the asymptotic convergence of a two-level multigrid cycle.

For $A^h = (\Delta^2)^h$, the smoothing of the second and third relaxations is as efficient as for the first one, for *all* c (for the Laplace operator, the third sweep already “lags” behind). However, the smoothing per sweep is much

worse than for the Laplace operator. This finding may be related to the results of [19] for RB GS for the Laplace operator: the smoothing factor turns out to be the maximum between a term that is independent of the number of sweeps, and one that depends *only* on the number of sweeps. The number of sweeps becomes important only when the latter term becomes dominant. Since a point GS is less efficient for $(\Delta^2)^h$, the “lagging” effect arises here too, but after more sweeps.

In the next section we present an alternative and more efficient approach for relaxing the biharmonic operator.

4 Composite Relaxations

In this section we analyze the novel “composite relaxation” (CR) scheme for products of operators, which was first introduced in [5] and applied in [10] for the 1D Schrödinger eigenproblem. This scheme may be applied to many problems, e.g., for the compressible and incompressible Euler and Navier-Stokes equations [5].

4.1 Definition

Suppose our discrete operator A^h can be written as the product $A^h = A_1 A_2$ (similar constructions can be made for an arbitrary number of factors, $A^h = A_1 \dots A_q$). Let R_t be a relaxation scheme for A_t , $t = 1, 2$. A way to obtain good smoothing for A^h (provided that $R_{1,2}$ have good smoothing factors) is

by regarding $Au = f$ (omitting the superscript h) as a 2×2 system

$$A_1 y = f \tag{23}$$

$$A_2 u = y, \tag{24}$$

which can be written as

$$\begin{pmatrix} A_1 & 0 \\ -I & A_2 \end{pmatrix} \begin{pmatrix} y \\ u \end{pmatrix} = \begin{pmatrix} f \\ 0 \end{pmatrix}. \tag{25}$$

According to the general principles of relaxation design for a system of equations [2], the smoothing factor will be $\max\{\bar{\mu}_{R_1,1}, \bar{\mu}_{R_2,1}\}$. However, this method requires twice the variables (u, y) instead of the original u throughout the multigrid solver, and a more complicated overall setup is required for the system (25) instead of the original scalar (1). Instead, a second and novel approach that was stated in [5] and tested in [10] is much more convenient. It might have slightly worse smoothing than system relaxation, but it typically has far better smoothing properties than with an optimized multicolor distributive relaxation (see §3.2 and [10]) for the original operator A .

A single sweep of the composite relaxation scheme $R := \text{CR}(\nu_1, \nu_2)$ consists of the following steps:

1. Let u^0 be the initial grid-function. Set $y^0 = A_2 u^0$.
2. Perform ν_1 R_1 -relaxation sweeps on (23), starting with $y = y^0$; denote the result by y^1 .
3. Perform ν_2 R_2 -relaxation sweeps on $A_2 u = y^1$, starting with $u = u^0$; denote the result by u^1 .

That is, the result of relaxing u^0 by R is u^1 . We focus on computing the symbol of a composite relaxation, and obtain some apriori bounds on its convergence and smoothing, in terms of convergence and smoothing for each factor.

4.2 The Symbol

In the rest of this section we assume that R_1 is a c_1 -process and R_2 is a c_2 -process.

Theorem 3 *R is an $\text{lcm}(c_1, c_2)$ -process.*

Proof. For simplicity, consider first the case of $c_1 = c_2 =: c$. Let $\hat{A}_2(\theta)$ be the symbol of the operator A_2 , let

$$\check{A}_2(\theta) := \text{diag}\{\hat{A}_2(\theta^{c,(0)}), \dots, \hat{A}_2(\theta^{c,(c^d-1)})\}, \quad \theta \in \Theta, \quad (26)$$

and let $\hat{R}_{1,2}(\theta)$ be the symbols of their relaxation schemes $R_{1,2}$. We denote the exact solution of $A_1 A_2 u = f$ by \bar{u} and $\bar{y} := A_2 \bar{u}$.

Starting from an initial error mode $v^0(x_\alpha) = \{e^{i\theta\alpha}\}$, our first approximation to \bar{u} is

$$u^0(x_\alpha) = \bar{u}(x_\alpha) + v^0(x_\alpha) = \bar{u}(x_\alpha) + e^{i\theta\alpha}.$$

In the first step of the composite relaxation we obtain

$$y^0(x_\alpha) = \bar{y}(x_\alpha) + \hat{A}_2(\theta)e^{i\theta\alpha}.$$

After ν_1 R_1 -relaxation sweeps on (23) we obtain

$$y^1(x_\alpha) = \bar{y}(x_\alpha) + \hat{R}_1(\theta)^{\nu_1} \check{A}_2(\theta) e^{i\theta\alpha}.$$

Thus, the equation to be relaxed at step 3 of the composite relaxation can be written as

$$A_2 v = z, \quad z(x_\alpha) := \hat{R}_1(\theta)^{\nu_1} \check{A}_2(\theta) e^{i\theta\alpha}, \quad v := u - \bar{u}. \quad (27)$$

The exact solution of this equation is $\bar{v}(x_\alpha) = \check{A}_2(\theta)^{-1} \hat{R}_1(\theta)^{\nu_1} \check{A}_2(\theta) e^{i\theta\alpha}$.

Therefore, we start the ν_2 R_2 -relaxation sweeps on (27) with the function

$$v^0(x_\alpha) = \bar{v} + \left(I_{c^d} - \check{A}_2(\theta)^{-1} \hat{R}_1(\theta)^{\nu_1} \check{A}_2(\theta) \right) e^{i\theta\alpha},$$

corresponding to u^0 , where I_{c^d} is the $c^d \times c^d$ identity matrix, and after these relaxation sweeps we obtain the new function

$$v^1(x_\alpha) = \bar{v} + \hat{R}_2(\theta)^{\nu_2} \left(I_{c^d} - \check{A}_2(\theta)^{-1} \hat{R}_1(\theta)^{\nu_1} \check{A}_2(\theta) \right) e^{i\theta\alpha}$$

which corresponds to the final result u^1 after the full R -sweep. The error in u^1 is

$$\left(\check{A}_2(\theta)^{-1} \hat{R}_1(\theta)^{\nu_1} \check{A}_2(\theta) + \hat{R}_2(\theta)^{\nu_2} \left(I_{c^d} - \check{A}_2(\theta)^{-1} \hat{R}_1(\theta)^{\nu_1} \check{A}_2(\theta) \right) \right) e^{i\theta\alpha}.$$

Hence,

$$\hat{R}(\theta) = \check{A}_2(\theta)^{-1} \hat{R}_1(\theta)^{\nu_1} \check{A}_2(\theta) + \hat{R}_2(\theta)^{\nu_2} \left(I_{c^d} - \check{A}_2(\theta)^{-1} \hat{R}_1(\theta)^{\nu_1} \check{A}_2(\theta) \right). \quad (28)$$

This implies in particular that R is a c -process. Note that (28) is also valid for general c_1, c_2 , with $\hat{\cdot}$ being replaced with the appropriate $\check{\cdot}$ (see (14)–(17))

and $c := \text{lcm}(c_1, c_2)$. That is,

$$\hat{R}(\theta) = \check{A}_2(\theta)^{-1} \check{R}_1(\theta)^{\nu_1} \check{A}_2(\theta) + \check{R}_2(\theta)^{\nu_2} \left(I_{c^d} - \check{A}_2(\theta)^{-1} \check{R}_1(\theta)^{\nu_1} \check{A}_2(\theta) \right). \quad \blacksquare \quad (29)$$

4.3 Convergence and Smoothing

The convergence and smoothing of the composite relaxation can be directly calculated from (29). However, the following theorem provides apriori upper bounds that may serve as “rules of thumb”, based on the convergence and smoothing of $R_{1,2}$.

To simplify the notation, we have assumed here that the rows and columns in the symbols of R, R_1, R_2 have been properly interchanged for all $\theta \in \Theta$, so that all high frequencies $\theta \in \mathcal{F}_\theta^c$ (for which $c(\theta) = 1$, denoted by “H”) appear before all low frequencies $\theta \in \mathcal{F}_\theta^l$ (with $c(\theta) = 0$, denoted by “L”). Thus, we write

$$C^{(c)}(\theta) = \begin{pmatrix} I_n & 0 \\ 0 & 0 \end{pmatrix}, \quad \check{R}_t(\theta)^{\nu_t} = \begin{pmatrix} R_t^{HH}(\theta) & R_t^{LH}(\theta) \\ R_t^{HL}(\theta) & R_t^{LL}(\theta) \end{pmatrix}, \quad t = 1, 2,$$

where n denotes the number of high frequencies in \mathcal{F}_θ^c and $m := c^d - n$ denotes the number of low frequencies.

Theorem 4 (CR Convergence Theorem) *Assume $c_1 = 1$. Then*

$$M_R \leq M_{R_1}^{\nu_1} + M_{R_2}^{\nu_2} + M_{R_1}^{\nu_1} M_{R_2}^{\nu_2} \quad (30)$$

$$\overline{M}_R \leq \overline{M}_{R_1}^{\nu_1} + \overline{M}_{R_2}^{\nu_2} + \overline{M}_{R_1}^{\nu_1} \overline{M}_{R_2}^{\nu_2} + \varepsilon, \quad (31)$$

where

$$\varepsilon := \sup_{\theta \in \Theta} \|R_2^{LH}(\theta) (I_m - R_1^{LL}(\theta))\|_2.$$

In other words, the smoothing of R for $c_1 = 1$ is the sum of the smoothing performed on each factor, plus two “cross-terms”, resulting from the interference of high frequencies with high frequencies ($M_{R_1}^{\nu_1} M_{R_2}^{\nu_2}$), and from the feedback (ε) of the remaining low frequencies after the R_1 -action ($I_m - R_1^{LL}(\theta)$) through the action of R_2 ($R_2^{LH}(\theta)$).

Proof of Theorem 4. For $c_1 = 1$, (29) becomes

$$\hat{R}(\theta) = \check{R}_1(\theta)^{\nu_1} + \check{R}_2(\theta)^{\nu_2} (I_{c^d} - \check{R}_1(\theta)^{\nu_1}).$$

Thus

$$\|\hat{R}(\theta)\|_2 \leq \|\check{R}_1(\theta)\|_2^{\nu_1} + \|\check{R}_2(\theta)\|_2^{\nu_2} + \|\check{R}_1(\theta)\|_2^{\nu_1} \|\check{R}_2(\theta)\|_2^{\nu_2}$$

for all $\theta \in \Theta$, which implies (30), even for a general matrix norm $\|\cdot\|$ instead of $\|\cdot\|_2$. Furthermore,

$$C^{(c)}(\theta) \hat{R}(\theta) = C^{(c)}(\theta) \check{R}_1(\theta)^{\nu_1} + C^{(c)}(\theta) \check{R}_2(\theta)^{\nu_2} (I_{c^d} - \check{R}_1(\theta)^{\nu_1}).$$

The norm of the first term is bounded by $\overline{M}_1^{\nu_1}$; the second (denoted by T) can be explicitly written as

$$\begin{aligned} T &= \begin{pmatrix} R_2^{HH}(\theta) & R_2^{LH}(\theta) \\ 0 & 0 \end{pmatrix} \begin{pmatrix} I_n - R_1^{HH}(\theta) & 0 \\ 0 & I_m - R_1^{LL}(\theta) \end{pmatrix} \\ &= \begin{pmatrix} R_2^{LH}(\theta) (I_n - R_1^{HH}(\theta)) & R_2^{LH}(\theta) (I_m - R_1^{LL}(\theta)) \\ 0 & 0 \end{pmatrix}. \end{aligned}$$

Hence,

$$\begin{aligned}
 \|T\|_2 &\leq \left\| \begin{pmatrix} R_2^{LH}(\theta) (I_n - R_1^{HH}(\theta)) & 0 \\ 0 & 0 \end{pmatrix} \right\|_2 \\
 &\quad + \left\| \begin{pmatrix} 0 & R_2^{LH}(\theta) (I_m - R_1^{LL}(\theta)) \\ 0 & 0 \end{pmatrix} \right\|_2 \\
 &= \|R_2^{LH}(\theta) (I_n - R_1^{HH}(\theta))\|_2 + \|R_2^{LH}(\theta) (I_m - R_1^{LL}(\theta))\|_2 \\
 &\leq \overline{M}_{R_2}^{\nu_2} + \overline{M}_{R_1}^{\nu_1} \overline{M}_{R_2}^{\nu_2} + \varepsilon,
 \end{aligned}$$

yielding (31). ■

Remark 1 For $c_1 = c_2 = 1$, we obtain

$$\mu_R \leq \mu_{R_1}^{\nu_1} + \mu_{R_2}^{\nu_2} + \mu_{R_1}^{\nu_1} \mu_{R_2}^{\nu_2} \quad (32)$$

$$\overline{\mu}_R \leq \overline{\mu}_{R_1}^{\nu_1} + \overline{\mu}_{R_2}^{\nu_2} + \overline{\mu}_{R_1}^{\nu_1} \overline{\mu}_{R_2}^{\nu_2}, \quad (33)$$

which seems to be the ideal performance that can be expected from a CR. The optimal smoothing per work unit, defined by $\overline{\mu}_R^{\frac{m}{\nu_1 + \nu_2}}$, where m is the ratio between the number of points in A 's stencil and A_1, A_2 's stencils (assuming that the latter ones are comparable and a work unit to be the work of one relaxation sweep on A) is attained when both smoothing $\overline{\mu}_1^{\nu_1}, \overline{\mu}_2^{\nu_2}$ are comparable and $\ll 1$.

Remark 2 The bounds (30),(31) are *not* valid for *general* c_1 , and in §4.4 it will be shown that if both $c_1, c_2 > 1$, even ideal smoothing for A_1, A_2 cannot guarantee $\overline{\mu}_R < 1$. This observation and our numerical examples (see §4.4)

led to the conclusion that a necessary condition for good smoothing of R is that at least one relaxation (R_1 or R_2) should be uni-coupling ($c_1 = 1$), e.g., a lexicographic or Jacobi type.

4.4 Numerical Examples

We considered the following four combinations of P_1, P_2 :

- (a) $P_1 =$ lexicographic (LEX) GS relaxation ($c_1 = 1$), $P_2 =$ LEX GS relaxation ($c_2 = 1$);
- (b) $P_1 =$ LEX GS relaxation ($c_1 = 1$), $P_2 =$ RB GS relaxation ($c_2 = 2$);
- (c) $P_1 =$ RB GS relaxation ($c_1 = 2$), $P_2 =$ LEX GS relaxation ($c_2 = 1$);
- (d) $P_1 =$ RB GS relaxation ($c_1 = 2$), $P_2 =$ RB GS relaxation ($c_2 = 2$),

for the biharmonic operator $A^h := (\Delta^2)^h$, as a product of $A_1 = A_2 := \Delta^h$ (see (22)). For each case, we computed for various ν_1, ν_2 values the smoothing factor e per work unit (with $m = 13/5 = 2.6$, see (22); m may change, depending on the implementation, but will most likely still be between 2 and 3), and the asymptotic convergence factor C_κ per work unit of a two-level multigrid cycle containing κ fine grid relaxation sweeps, for $\kappa = 1, 2$. The results are summarized in Table 2.

Table 2 can be used to choose the $c_1, c_2, \nu_1, \nu_2, \kappa$ that yield the optimal multigrid convergence. Although the results depend on the specific definition of a

“work unit”, the following observations provide general guidelines for choosing the parameters:

- Case (d) ($c_1 = 2, c_2 = 2$) yields no smoothing. Indeed, it can be shown that if we start with a highly oscillatory mode (e.g., with frequency $(\theta_1, \theta_2) \approx (\pi, \pi)$), it will be coupled through R_1 with a very smooth mode. This smooth harmonic enters the *right-hand side* of (24), contributing a large smooth mode to the error in u of *that equation*. After the R_2 action, this smooth error regenerates the oscillatory mode, exactly at the amplitude of the original oscillatory mode. This unwanted feedback can be prevented by restricting at least one of the relaxation schemes to be lexicographic, although R_1 and R_2 will still be somewhat confounded.
- The best multigrid convergence per work unit (0.374) is attained when $\kappa = 1$ and both $\nu_1, \nu_2 \geq 2$, but not excessively large. That is, it is better to first perform 2 – 3 relaxation sweeps on A_1 , and only then perform 2 – 3 sweeps on A_2 , rather than alternate between the two. The latter corresponds to the case $\nu_1 = \nu_2 = 1$, whose best error reduction per work unit is $C_2 = 0.528$.
- Since an error reduction of about 0.1 per cycle is adequate for practical purposes (e.g., for adaptation to an FMG algorithm [2]), the cases $c_1 = 1, c_2 = 2, \nu_1 = 3, \nu_2 = 2, \kappa = 1$ and $c_1 = 2, c_2 = 1, \nu_1 = 2, \nu_2 = 3, \kappa = 1$ seem to be optimal.

- The smoothing factor e predicts the two-level convergence rate C_1 quite accurately in most cases. A second cycle (C_2) is usually slower, since the coarse grid correction rate of moderately smooth components lags behind the smoothing rate of high-frequency components [2].
- The “rule of thumb” b (based on the left hand side of (33)) gives an upper bound for e almost in all cases, although Remark 1 is theoretically limited to $c_1 = c_2 = 1$. Moreover, b usually gives a good approximation for the smoothing factor and for the convergence factor of multigrid cycles, especially for large ν_1, ν_2 (where $\bar{\mu}_1^{\nu_1}, \bar{\mu}_2^{\nu_2} \ll 1$). Hence, once the smoothing factors of $R_{1,2}$ are computed, b can be employed as a practical estimate for the performance of composite relaxation.
- It is also possible to alternate between the LEX and RB ordering *within* the sequence of ν_1 relaxation sweeps on A_1 (and similarly for A_2). However, this technique did not improve the results presented in Table 2.

Note that a multigrid cycle that employ such CR relaxation sweeps is 4.4 times faster than a multigrid cycle with GS sweeps for the compound biharmonic operator (see Table 1). This justifies the CR approach for the biharmonic operator; similar accelerations are expected for more general product operators.

5 Concluding Remarks

In this paper we have presented formal local mode analysis for general *multicoupling processes*. As discussed in [4, 15, 16], the predicted convergence and smoothing factors are quantitatively sharp and can be obtained, at least for the case of discretized elliptic PDEs, provided that proper supplementary processing is performed at and near the boundaries [4]. For generalizations of our analysis to staggered grids, piecewise smooth stencil coefficients, systems of equations, etc., we refer the reader to [3, 4].

The presented local mode analysis follows [4, 17, 18] in putting both “grid functions” and “Fourier modes” in the same space. This formulation is by no means unique; for instance, one may instead utilize the Fourier *transform* (using the continuum of Fourier frequencies instead of a discrete set) to define symbols of processes. While the latter formulation enjoys several advantages (e.g., connecting the convergence factor μ_P to an operator-norm of P by Parseval’s identity), the resulting formulae for symbols are of course identical in all the equivalent formulations.

Already at the formal level, the analysis of multicolor processes provides important tools for the general design of relaxation. Moreover, the novel approach of composite relaxation provides a simple scheme for relaxing products of discrete operators. Its smoothing per work unit has proved to be good enough, provided that the relaxations of the factors are good smoothers, and their confounded coupling of Fourier modes has a small effect. The latter

requirement was shown by our analysis to be fulfilled by using lexicographic ordering for the factor whose relaxation is performed first.

In general, the convergence and smoothing theorems, such as those presented here, provide some insights; however, they may be skipped by the practitioner by directly computing the desired performance from the *process symbol*.

The discussion in this paper relates only to problems on grids. For disordered problems there is no natural “ordering” of variables, and we cannot use “colors” there. Any ordering would result in the coupling of modes, and the design of relaxation (in general, and composite relaxation, in particular) should be done with caution.

6 Acknowledgments

The authors are thankful to Prof. Irad Yavneh from the Department of Computer Science, Technion, Israel, for his important comments and suggestions concerning this paper. This work was supported in part by Grant No. 295/01 from the Israel Science Foundation, by AFOSR Contract F33615-97-D5405, and by the Carl F. Gauss Minerva Center for Scientific Computation at the Weizmann Institute of Science.

References

- [1] A. BRANDT, *Multi-level adaptive solutions to boundary-value problems*,

- Math. Comp., 31 (1977), pp. 333–390.
- [2] ———, *Guide to multigrid development*, in Multigrid Methods, W. Hackbusch and U. Trottenberg, eds., vol. 960 of Lecture Notes in Mathematics, Springer-Verlag, Berlin, 1982, pp. 220–312.
- [3] A. BRANDT, *Rigorous quantitative analysis of multigrid, II: Extensions and practical implications*, manuscript, 1991. Appeared together with [4] in Preliminary Proceedings of the 4th Copper Mountain Conference on Multigrid Methods, 1989.
- [4] A. BRANDT, *Rigorous quantitative analysis of multigrid, I: Constant coefficients two-level cycle with L_2 -norm*, SIAM J. Numer. Anal., 31 (1994), pp. 1695–1730.
- [5] A. BRANDT, *Barriers to achieving textbook multigrid efficiency (TME) in CFD*, tech. report, Gauss Center Report WI/GC–10, Weizmann Institute, Dec. 1998. Appears as Appendix C in [15].
- [6] B. DISKIN, *Multigrid algorithm with conditional coarsening for the non-aligned sonic flow*, Elect. Trans. Numer. Anal., 6 (1997), pp. 106–119.
- [7] G. E. FORSYTHE, *Finite-difference methods for partial differential equations*, Wiley, New York, 1960.

- [8] C. C. J. KUO AND T. F. CHAN, *Two-color Fourier analysis of iterative algorithms for elliptic problems with red/black ordering*, SIAM J. Sci. Stat. Comput., 11 (1990), pp. 767–793.
- [9] C.-C. J. KUO AND B. C. LEVY, *Two-color Fourier analysis of the multigrid method with red-black Gauss-Seidel smoothing*, Appl. Math. Comput., 29 (1989), pp. 69–87.
- [10] O. E. LIVNE, *Multiscale Eigenbasis Algorithms*, PhD thesis, The Weizmann Institute of Science, Rehovot, Israel, 2000.
- [11] O. E. LIVNE AND A. BRANDT, *Formal Fourier Analysis of Multicolor and Composite Relaxation Schemes*, tech. report, Gauss Center Report WI/GC–14, Weizmann Institute, Feb. 2002. Available at <http://www.wisdom.weizmann.ac.il/~livneo/publications.html>.
- [12] R. STEVENSON, *On the validity of local mode analysis of multi-grid methods*, PhD thesis, Rijksuniversiteit te Utrecht, The Netherlands, 1991.
- [13] J. STOER AND R. BULIRSCH, *Introduction to Numerical Analysis*, Translation of *Einführung in die Numerische*, Springer–Verlag, New York, 1980.
- [14] K. STÜBEN AND U. TROTTENBERG, *Multigrid methods: Fundamental algorithms, model problem analysis and applications*, in *Multigrid Meth-*

- ods, W. Hackbusch and U. Trottenberg, eds., vol. 960 of Lecture Notes in Mathematics, Berlin, 1982, Springer-Verlag, pp. 1–176.
- [15] U. TROTTENBERG, C. W. OOSTERLEE, AND A. SCHÜLLER, *Multigrid*, Academic Press, 2001.
- [16] P. WESSELING, *An Introduction to Multigrid Methods*, Wiley, Chichester, 1992. Reprinted by www.MGNet.org.
- [17] R. WIENANDS, *Extended Local Fourier Analysis for Multigrid: Optimal Smoothing, Coarse Grid Correction, and Preconditioning*, PhD thesis, Universität zu Köln, 2001.
- [18] R. WIENANDS AND C. W. OOSTERLEE, *On three-grid Fourier analysis for multigrid*, SIAM J. Sci. Comput., 23(2) (2001), pp. 651–671.
- [19] I. YAVNEH, *Smoothing factors of two-color Gauss-Seidel relaxation for a class of elliptic operators*, SIAM J. Num. Anal., 32(4) (1995), pp. 1126–1138.
- [20] ———, *On red black SOR smoothing in multigrid*, SIAM J. Sci. Comput., 17(1) (1996), pp. 180–192.

Table 1: Smoothing per sweep ($\bar{\mu}_{P,\nu}, \nu = 1, 2, 3$), and two-level multigrid cycle (containing ν fine grid relaxations) asymptotic convergence factors per work unit with linear and cubic transfers (denoted $C_{P,\nu}^{(2)}, \nu = 1, 2, 3$ and $C_{P,\nu}^{(4)}, \nu = 1, 2, 3$, respectively) of multicolor GS relaxation for Poisson and biharmonic operators, for various numbers of colors

Measure	$A^h = \Delta^h$				$A^h = (\Delta^2)^h$			
	$c = 1$	$c = 2$	$c = 3$	$c = 4$	$c = 1$	$c = 2$	$c = 3$	$c = 4$
$\bar{\mu}_{P,1}$	0.500	0.250	0.300	0.364	0.800	0.700	0.737	0.750
$C_{P,1}^{(2)}$	0.400	0.250	0.266	0.310	0.785	0.800	0.779	0.777
$C_{P,1}^{(4)}$	0.400	0.250	0.273	0.310	0.785	0.800	0.785	0.777
$\bar{\mu}_{P,2}$	0.500	0.250	0.305	0.366	0.800	0.700	0.737	0.750
$C_{P,2}^{(2)}$	0.437	0.270	0.302	0.349	0.784	0.745	0.755	0.866
$C_{P,2}^{(4)}$	0.416	0.261	0.263	0.290	0.784	0.745	0.758	0.758
$\bar{\mu}_{P,3}$	0.500	0.321	0.314	0.366	0.800	0.700	0.737	0.750
$C_{P,3}^{(2)}$	0.491	0.371	0.398	0.440	0.785	0.732	0.752	0.752
$C_{P,3}^{(4)}$	0.436	0.385	0.378	0.356	0.785	0.732	0.752	0.752

Table 2: The biharmonic: the theoretical bound $b := (\bar{\mu}_{R_1}^{\nu_1} + \bar{\mu}_{R_2}^{\nu_2} + \bar{\mu}_{R_1}^{\nu_1} \bar{\mu}_{R_2}^{\nu_2})^{\frac{m}{\nu_1 + \nu_2}}$ (see (33)) versus the actual smoothing per work unit of composite relaxation ($e := \bar{\mu}_{P_1}^{\frac{m}{\nu_1 + \nu_2}}$, computed by (29)), and the asymptotic convergence factor per work unit C_κ of a two-level multigrid cycle containing $\kappa = 1, 2$ fine grid $\text{CR}(\nu_1, \nu_2)$ sweeps, for various c_1, c_2, ν_1, ν_2 values. The convergence factor in all cases is 1. In case (d), $e = C_1 = C_2 = 1$ and $b \leq .491$ for all c_1, c_2, ν_1, ν_2

ν_1, ν_2	(a) LEX/LEX				(b) LEX/RB				(c) RB/LEX			
	b	e	C_1	C_2	b	e	C_1	C_2	b	e	C_1	C_2
1,1	1.34	.787	.722	.724	.841	.543	.638	.528	.841	.543	.638	.528
1,2	.891	.705	.617	.617	.636	.567	.545	.538	.607	.452	.596	.506
2,1	.891	.705	.617	.617	.607	.452	.596	.506	.636	.567	.545	.538
1,3	.784	.686	.633	.633	.657	.639	.606	.605	.557	.447	.638	.533
2,2	.688	.628	.531	.554	.485	.422	.410	.448	.485	.422	.410	.477
3,1	.784	.686	.633	.633	.557	.447	.638	.533	.657	.639	.606	.605
1,4	.763	.703	.668	.668	.702	.697	.662	.661	.560	.483	.696	.576
2,3	.626	.590	.514	.576	.506	.487	.450	.513	.428	.361	.374	.527
3,2	.626	.590	.514	.576	.428	.361	.374	.486	.506	.487	.450	.532
4,1	.763	.703	.668	.668	.560	.483	.696	.576	.702	.697	.662	.661
2,4	.617	.581	.526	.603	.553	.548	.503	.563	.412	.330	.403	.565
3,3	.563	.553	.447	.590	.430	.401	.374	.532	.430	.401	.387	.559
4,2	.617	.581	.526	.603	.412	.330	.403	.511	.553	.548	.503	.576
3,4	.545	.533	.476	.617	.468	.461	.416	.570	.390	.347	.414	.583
4,3	.545	.533	.476	.617	.390	.347	.377	.550	.468	.461	.422	.588
3,5	.551	.530	.499	.640	.510	.508	.457	.602	.371	.339	.448	.617
4,4	.514	.512	.487	.634	.415	.403	.418	.589	.415	.403	.445	.615
5,3	.551	.530	.499	.640	.371	.339	.394	.578	.510	.508	.459	.618
4,5	.508	.501	.504	.649	.451	.448	.454	.621	.381	.364	.470	.641
5,4	.508	.501	.504	.649	.381	.364	.430	.610	.451	.448	.474	.642
4,6	.517	.500	.522	.670	.487	.486	.487	.648	.360	.386	.489	.664