

Supremum–Norm Convergence for Step–Asynchronous Successive Overrelaxation on M-matrices

Sebastiano Vigna*

April 13, 2014

Abstract

Step-asynchronous successive overrelaxation updates the values contained in a single vector using the usual Gauß–Seidel-like weighted rule, but arbitrarily mixing old and new values, the only constraint being temporal coherence—you cannot use a value before it has been computed. We show that given a nonnegative real matrix A , a $\sigma \geq \rho(A)$ and a vector $\mathbf{w} > 0$ such that $A\mathbf{w} \leq \sigma\mathbf{w}$, every iteration of step-asynchronous successive overrelaxation for the problem $(sI - A)\mathbf{x} = \mathbf{b}$, with $s > \sigma$, reduces geometrically the \mathbf{w} -norm of the current error by a factor that we can compute explicitly. Then, we show that given a $\sigma > \rho(A)$ it is in principle always possible to compute such a \mathbf{w} . This property makes it possible to estimate the supremum norm of the absolute error at each iteration without any additional hypothesis on A , even when A is so large that computing the product $A\mathbf{x}$ is feasible, but estimating the supremum norm of $(sI - A)^{-1}$ is not.

Mathematical Subject Classification: 65F10 (Iterative methods for linear systems)

Keywords: Successive overrelaxation; M-matrices; asynchronous iterative solvers

1 Introduction

We are interested in providing *computable absolute bounds in ℓ_∞ norm* on the convergence of a mildly asynchronous version of successive overrelaxation (SOR) applied to problems of the form $(sI - A)\mathbf{x} = \mathbf{b}$, where A is a nonnegative real matrix and $s > \rho(A)$. A matrix of the form $sI - A$ under these hypotheses is called a *nonsingular M-matrix* [BP94].

We stress from the start that there are no other hypotheses on A such as irreducibility, symmetry, positive definiteness or (weak) 2-cyclicity, and that A is assumed to be very large—so large that computing $A\mathbf{x}$ (or performing a SOR iteration) is feasible (maybe streaming over the matrix entries), but estimating $\|(sI - A)^{-1}\|_\infty$ is not.

*The author was supported by the EU-FET grant NADINE (GA 288956).

Our main motivation is the parallel computation with arbitrary guaranteed precision of various kinds of *spectral rankings with damping* [Vig09], most notably Katz’s index [Kat53] and PageRank [PBMW98], which are solutions of problems of the form above with A derived from the adjacency matrix of a very large graph, the only relevant difference being that the rows of A are ℓ_1 -normalized in the case of PageRank.

By “computable” we mean that there must be a finite computational process that provides a bound on $\|\bar{\mathbf{x}} - \mathbf{x}^{(t)}\|_\infty$, where $\bar{\mathbf{x}}$ is the solution and $\mathbf{x}^{(t)}$ is the t -th approximation. Such a bound would make it possible to claim that we know the solution up to some given number of significant fractional digits. For example, without further assumptions on A convergence results based on the spectral radius are not computable in this sense and results concerning the residual are not applicable because of the unfeasibility of estimating $\|(sI - A)^{-1}\|_\infty$.

We are also interested in highly parallel versions for modern multicore systems. While SOR and other iterative methods are apparently strictly sequential algorithms, there is a large body of literature that studies what happens when updates are executed in arbitrary order, mixing old and new values. Essentially, as long as old values come from a finite time horizon (e.g., there is a finite bound on the “oldness” of a value) convergence has been proved for all major standard sequential hypothesis of convergence¹ (for the main results, see the sections about *partial asynchrony* in Bertsekas and Tsitsiklis’s encyclopedic book [BT89]).

Again, however, results are always stated in terms of convergence in the limit, and the speed of convergence, which decays as the time horizon gets larger, often cannot be stated explicitly. Moreover, the theory is modeled around message-passing systems, where processor might actually use very old values due to transmission delays. In the multicore, shared-memory system application we have in mind it is reasonable to assume that after each iteration memory is synchronized and all processors have the same view.

Our main motivation is obtaining (almost) “noise-free” scores to perform accurate comparisons of the induced rankings using Kendall’s τ [Ken45]:

$$\tau(\mathbf{r}, \mathbf{s}) := \frac{\sum_{i < j} \operatorname{sgn}(r_i - r_j) \operatorname{sgn}(s_i - s_j)}{\sqrt{\sum_{i < j} \operatorname{sgn}(r_i - r_j)^2} \sqrt{\sum_{i < j} \operatorname{sgn}(s_i - s_j)^2}}.$$

Computational noise can be quite problematic in evaluating Kendall’s τ because the signum function has no way to distinguish large and small differences—they are all mapped to 1 or -1 [BPSV08].

Suppose, for example, that we have a graph with a large number n of nodes, and some centrality index that assigns score 0 the first $n/2$ nodes and score 1 the remaining nodes. Suppose we have also another index assigning the same scores, and that this new index is defined by an iterative process, which is stopped at some point (e.g., an iterative solver for linear systems). If the computed values include computational random noise and evaluate τ on the two vectors, we will obtain a τ close to $1/\sqrt{2} \approx 0.707$, even if the ranks are perfectly correlated. On the other

¹It is a bit surprising, indeed, that the statement that Gauß-Seidel is difficult to parallelize appears so often in the literature. In a sense, an algorithm updating in arbitrary order using possibly old values is not any longer Gauß-Seidel. On the other hand, this is exactly what one expects when asking the question “is Gauß-Seidel parallelizable”?

hand, with a sufficiently small guaranteed absolute error we can proceed to truncate or round the second set of scores, obtaining a result closer to the real correlation.

This scenario is not artificial: when comparing, for instance, indegree with an index computed iteratively (e.g., Katz’s index, PageRank, etc.), we have a similar situation. Surprisingly, the noise from iterative computations can even *increase* correlation (e.g., between the dominant eigenvector of a graph that is not strongly connected and Katz’s index, as the residual score in nodes whose actual score is zero induces a ranking similar to that induced by Katz’s index).

In this paper, we provide convergence bounds in ℓ_∞ norm for SOR iterations for the problem $(sI - A)\mathbf{x} = \mathbf{b}$, where A is a nonnegative real matrix and $s > \rho(A)$, in conditions of mild asynchrony, without any additional hypothesis on A . Our main results are Theorem 1, which shows that given a $\sigma < s$ and a vector $\mathbf{w} > 0$ such that $A\mathbf{w} \leq \sigma\mathbf{w}$ SOR iterations reduce geometrically the \mathbf{w} -norm of the error (with a computable contraction factor), and Theorem 2, which shows how to compute such a \mathbf{w} using only iterated products of A with a vector. The two results can be viewed as a constructive and computable version of the standard convergence results on SOR iteration based on the spectral radius.

We remark that SOR is actually not useful for PageRank, as shown recently by Greif and Kurokawa [GK11]. The author has found experimentally that the same phenomenon plagues the computation of Katz’s index. However, since generalizing from Gauß–Seidel to SOR does not bring any significant increase in complexity in the proof, we decided to prove our results in the more general setting.

2 Step-asynchronous SOR

We now define *step-asynchronous* SOR for the problem $(sI - A)\mathbf{x} = \mathbf{b}$. In general, *asynchronous* SOR computes new values using arbitrarily old values; in this case, the hypotheses for convergence are definitely stronger. In the *partially asynchronous* case, instead, there is a finite limit on the “oldness” of the values used to compute new values, and while there is a decrease in convergence speed, the hypotheses for convergence are essentially the same of the sequential case (see [BT89] for more details).

Step-asynchronous SOR uses the strictest possible time bound: one step. We thus perform a SOR-like update in arbitrary order:

$$x_i^{(t+1)} = (1 - \omega)x_i^{(t)} + \frac{\omega}{s - a_{ii}} \left(b_i + \sum_{j \in N_i^{(t)}} a_{ij}x_j^{(t+1)} + \sum_{j \in P_i^{(t)} \setminus \{i\}} a_{ij}x_j^{(t)} \right). \quad (1)$$

The only constraint is that for each iteration an *update total preorder*² $\preceq^{(t)}$ of the indices is given: $i \preceq^{(t)} j$ iff x_i is updated before (or at the same time of) x_j at iteration t , and the set $P_i^{(t)}$ of the indices for which we use the *previous* values is such that for all $j \succeq^{(t)} i$ we have $j \in P_i^{(t)}$, whereas $N_i^{(t)} = n \setminus P_i^{(t)}$ is the set indices for which we use the *next* values. Essentially, we *must* use previous values for all variables that are updated at the same time of x_i or after x_i , but we make

²A *total preorder* is a set endowed with a reflexive and transitive total relation. We remark that a choice of a sequence of such preorders is equivalent to a *scenario* in the terminology of [BT89].

no assumption on the remaining variables. In this way we take into account cache incoherence, unpredictable scheduling of multiple threads, and so on.³

Matrixwise, the set $N_i^{(t)}$ induces a nonnegative matrix $L^{(t)}$ given by

$$L_{ij}^{(t)} = \left[j \in N_i^{(t)} \right] a_{ij}$$

and a *regular splitting*

$$sI - A = (D - L^{(t)}) - R^{(t)},$$

where $D = sI - \text{Diag}(A)$ and $R^{(t)}$ is nonnegative with zeros on the diagonal. Then, equation (1) can be rewritten as

$$(D - \omega L^{(t)})\mathbf{x}^{(t+1)} = (1 - \omega)D\mathbf{x}^{(t)} + \omega(\mathbf{b} + R^{(t)}\mathbf{x}^{(t)}).$$

There is of course a permutation of row and columns (depending on t) such that $L^{(t)}$ is strictly lower triangular, but the only claim that can be made about $R^{(t)}$ is that its diagonal is zero: actually, we could have $L^{(t)} = 0$ and $R^{(t)} = sI - A - D$.

In particular, independently from the choice of $L^{(t)}$, if $\bar{\mathbf{x}}$ is a solution we have as usual

$$(D - \omega L^{(t)})\bar{\mathbf{x}} = (1 - \omega)D\bar{\mathbf{x}} + \omega(\mathbf{b} + R^{(t)}\bar{\mathbf{x}})$$

and

$$(D - \omega L^{(t)})(\bar{\mathbf{x}} - \mathbf{x}^{(t+1)}) = (1 - \omega)D(\bar{\mathbf{x}} - \mathbf{x}^{(t)}) + \omega R^{(t)}(\bar{\mathbf{x}} - \mathbf{x}^{(t)}). \quad (2)$$

3 Suitability and convergence in w -norm

We now define suitability of a vector for a matrix, which will be the main tool in proving our results. The idea is implicitly or explicitly at the core of several classical proofs of convergence, and is closely related to that of *generalized diagonal dominance*:

Definition 1 A vector $\mathbf{w} > 0$ is σ -suitable for A if $A\mathbf{w} \leq \sigma\mathbf{w}$.

The usefulness of suitable vectors is that they induce norms in which the decrease of the error caused by a SOR iteration for of the problem $(sI - A)\mathbf{x} = \mathbf{b}$ can be controlled if $s > \sigma$. If A is irreducible, for instance, the dominant eigenvector is suitable for the spectral radius, but it is exactly this kind of hypotheses that we want to avoid.

Definition 2 Given a vector $\mathbf{w} > 0$, the w -norm is defined by

$$\|\mathbf{x}\|_{\infty}^{\mathbf{w}} = \max_i \frac{|x_i|}{w_i}.$$

The notation $\|\cdot\|_{\infty}^{\mathbf{w}}$ is used also for the operator norm induced in the usual way. We note a few useful properties—many others can be found in [BT89]:

³For example, if we have exactly n parallel updates at the same time we would have, in fact, a Jacobi iteration: in that case, $N_i^{(t)} = \emptyset$ for all i .

Proposition 1 *Given a vector \mathbf{w} that is σ -suitable for a nonnegative matrix A , the following statements are true for all vectors \mathbf{x} :*

1. $|x_i| \leq w_i \|\mathbf{x}\|_\infty^{\mathbf{w}}$;
2. $\min_i w_i \|\mathbf{x}\|_\infty^{\mathbf{w}} \leq \|\mathbf{x}\|_\infty$;
3. $\max_i w_i \|\mathbf{x}\|_\infty^{\mathbf{w}} \geq \|\mathbf{x}\|_\infty$;
4. $\|\mathbf{w}\|_\infty^{\mathbf{w}} = 1$;
5. $\|A\|_\infty^{\mathbf{w}} = \|A\mathbf{w}\|_\infty^{\mathbf{w}}$;
6. if $\mathbf{x} \geq 0$, $\|\mathbf{x}\|_\infty^{\mathbf{w}} = \min\{\alpha \geq 0 \mid \mathbf{x} \leq \alpha\mathbf{w}\}$.
7. $\|A\mathbf{x}\|_\infty^{\mathbf{w}} \leq \sigma \|\mathbf{x}\|_\infty^{\mathbf{w}}$; in particular, $\rho(A) \leq \|A\|_\infty^{\mathbf{w}} \leq \sigma$.

Proof. The first claims are immediate from the definition of \mathbf{w} -norm. For the last claim,

$$\|A\mathbf{x}\|_\infty^{\mathbf{w}} = \max_i \left| \frac{\sum_j a_{ij} x_j}{w_i} \right| \leq \max_i \frac{\sum_j a_{ij} |x_j|}{w_i} = \max_i \frac{\sum_j a_{ij} w_j \|\mathbf{x}\|_\infty^{\mathbf{w}}}{w_i} \leq \sigma \|\mathbf{x}\|_\infty^{\mathbf{w}}.$$

■

The next theorem is based on the standard proof by induction of convergence for SOR, but we make induction on the update time of a component rather than on its index, and we use suitability to provide bounds to the norm of the error.

Theorem 1 *Let A be a nonnegative matrix and let \mathbf{w} be σ -suitable for A . Then, given $s > \sigma$ step-asynchronous SOR for the problem $(sI - A)\mathbf{x} = \mathbf{b}$ converges for*

$$0 < \omega < \frac{2}{1 + \max_k \frac{\sigma - a_{kk}}{s - a_{kk}}}$$

and letting $\bar{\mathbf{x}} = (sI - A)^{-1}\mathbf{b}$ we have

$$\|\bar{\mathbf{x}} - \mathbf{x}^{(t+1)}\|_\infty^{\mathbf{w}} \leq r \|\bar{\mathbf{x}} - \mathbf{x}^{(t)}\|_\infty^{\mathbf{w}},$$

where

$$r = |1 - \omega| + \omega \max_k \frac{\sigma - a_{kk}}{s - a_{kk}} < 1.$$

Moreover,

$$\|\bar{\mathbf{x}} - \mathbf{x}^{(t+1)}\|_\infty^{\mathbf{w}} \leq \frac{r}{1 - r} \|\mathbf{x}^{(t+1)} - \mathbf{x}^{(t)}\|_\infty^{\mathbf{w}}.$$

Proof. Let $\preceq^{(t)}$ be a sequence of update orders, and $P_i(t)$ a sequence of previous-value sets, one for each step t and variable index i , compatible with the respective update orders. We work by induction on the order $\preceq^{(t)}$, proving the statement

$$|e_i^{(t+1)}| \leq \left(|1 - \omega| + \omega \frac{\sigma - a_{ii}}{s - a_{ii}} \right) w_i \|e^{(t)}\|_\infty^{\mathbf{w}}, \quad (3)$$

where $\mathbf{e}^{(t)} = \bar{\mathbf{x}} - \mathbf{x}^{(t)}$, assuming it is true for all $k \prec^{(t)} i$.

Note that for all i

$$0 < \frac{\sigma - a_{ii}}{s - a_{ii}} < 1,$$

so for $0 < \omega \leq 1$

$$|1 - \omega| + \omega \frac{\sigma - a_{ii}}{s - a_{ii}} = 1 - \omega \left(1 - \frac{\sigma - a_{ii}}{s - a_{ii}} \right) < 1,$$

and analogously for

$$1 \leq \omega < \frac{2}{1 + \max_k \frac{\sigma - a_{kk}}{s - a_{kk}}}$$

we have

$$|1 - \omega| + \omega \frac{\sigma - a_{ii}}{s - a_{ii}} = \omega \left(1 + \frac{\sigma - a_{ii}}{s - a_{ii}} \right) - 1 < \frac{2}{1 + \max_k \frac{\sigma - a_{kk}}{s - a_{kk}}} \left(1 + \frac{\sigma - a_{ii}}{s - a_{ii}} \right) - 1 < 1.$$

Writing explicitly (2) for the i -th coordinate, we have

$$|e_i^{(t+1)}| = \left| (1 - \omega)e_i^{(t)} + \frac{\omega}{s - a_{ii}} \left(\sum_{j \in N_i^{(t)}} a_{ij}e_j^{(t+1)} + \sum_{j \in P_i^{(t)} \setminus \{i\}} a_{ij}e_j^{(t)} \right) \right|.$$

Since $j \in N_i^{(t)}$ implies by definition $j \prec^{(t)} i$, we can apply the induction hypothesis on $e_j^{(t+1)}$ to state that $e_j^{(t+1)} \leq w_j \|\mathbf{e}^{(t)}\|_{\infty}^{\mathbf{w}}$. The same bound applies to $e_j^{(t)}$ using the first statement of Proposition 1.

We now notice that σ -suitability implies

$$(A - \text{Diag}(A))\mathbf{w} \leq (\sigma I - \text{Diag}(A))\mathbf{w},$$

which in coordinates tells us that

$$\sum_{j \neq i} a_{ij}w_j \leq (\sigma - a_{ii})w_i.$$

Thus,

$$\begin{aligned} |e_i^{(t+1)}| &\leq \left(|1 - \omega|w_i + \omega \frac{1}{s - a_{ii}} \left(\sum_{j \in N_i^{(t)}} a_{ij}w_j + \sum_{j \in P_i^{(t)} \setminus \{i\}} a_{ij}w_j \right) \right) \|\mathbf{e}^{(t)}\|_{\infty}^{\mathbf{w}} \\ &\leq \left(|1 - \omega| + \omega \frac{\sigma - a_{ii}}{s - a_{ii}} \right) w_i \|\mathbf{e}^{(t)}\|_{\infty}^{\mathbf{w}}. \end{aligned}$$

By the very definition of \mathbf{w} -norm, (3) yields

$$\|\mathbf{e}_i^{(t+1)}\|_{\infty}^{\mathbf{w}} \leq \left(|1 - \omega| + \omega \max_k \frac{\sigma - a_{kk}}{s - a_{kk}} \right) \|\mathbf{e}^{(t)}\|_{\infty}^{\mathbf{w}}.$$

For the second statement, we have

$$\begin{aligned} \|\mathbf{x} - \mathbf{x}^{(t)}\|_{\infty}^{\mathbf{w}} - \|\mathbf{x}^{(t+1)} - \mathbf{x}^{(t)}\|_{\infty}^{\mathbf{w}} &\leq \|\mathbf{x} - \mathbf{x}^{(t+1)} + \mathbf{x}^{(t)} - \mathbf{x}^{(t)}\|_{\infty}^{\mathbf{w}} \\ &= \|\mathbf{x} - \mathbf{x}^{(t+1)}\|_{\infty}^{\mathbf{w}} \leq r \|\mathbf{x} - \mathbf{x}^{(t)}\|_{\infty}^{\mathbf{w}}, \end{aligned}$$

whence

$$\|\mathbf{x} - \mathbf{x}^{(t+1)}\|_{\infty}^{\mathbf{w}} \leq r \|\mathbf{x} - \mathbf{x}^{(t)}\|_{\infty}^{\mathbf{w}} \leq \frac{r}{1-r} \|\mathbf{x}^{(t+1)} - \mathbf{x}^{(t)}\|_{\infty}^{\mathbf{w}}.$$

■

We remark that the smallest contraction factor is obtained when $\omega = 1$, that is, with no relaxation. This does not mean, however, that relaxation is not useful: convergence might be faster with $\omega \neq 1$; it is just that the error bound we provide features the best constant when $\omega = 1$.

Corollary 1 *With the same hypotheses and notation of Theorem 1, step-asynchronous Gauß–Seidel iterations converge and*

$$\|\bar{\mathbf{x}} - \mathbf{x}^{(t+1)}\|_{\infty}^{\mathbf{w}} \leq \frac{\max_k \frac{\sigma - a_{kk}}{s - a_{kk}}}{1 - \max_k \frac{\sigma - a_{kk}}{s - a_{kk}}} \|\mathbf{x}^{(t+1)} - \mathbf{x}^{(t)}\|_{\infty}^{\mathbf{w}}.$$

Corollary 2 *Let A be an irreducible nonnegative matrix and \mathbf{w} its dominant eigenvector. Then the statement of Theorem 1 is true in \mathbf{w} -norm with $\sigma = \rho(A)$.*

A simple consequence is that if we know a σ -suitable vector \mathbf{w} for A we can just behave as if the step-asynchronous SOR is converging in the standard supremum norm, but we have a reduction in the strength of the bound given by the ratio between the maximum and the minimum component of \mathbf{w} :

Corollary 3 *With the same hypotheses and notation of Theorem 1, step-asynchronous Gauß–Seidel iterations converge and*

$$\|\bar{\mathbf{x}} - \mathbf{x}^{(t+1)}\|_{\infty} \leq \frac{\max_i w_i}{\min_i w_i} \frac{\max_k \frac{\sigma - a_{kk}}{s - a_{kk}}}{1 - \max_k \frac{\sigma - a_{kk}}{s - a_{kk}}} \|\mathbf{x}^{(t+1)} - \mathbf{x}^{(t)}\|_{\infty}.$$

Proof. An application of Proposition 1.2 and 1.3. ■

We remark that

$$\max_k \frac{\sigma - a_{kk}}{s - a_{kk}} \leq \frac{\sigma}{s},$$

so it is possible to restate all results in a simplified (but less powerful) form.

4 Practical issues

In principle it is always better to compute the actual \mathbf{w} -norm, rather than using the rather crude bound of Corollary 3.⁴ On the other hand, computing the \mathbf{w} -norm requires storing and accessing \mathbf{w} , which could be expensive.

In practice, it is convenient to restrict oneself to vectors \mathbf{w} satisfying $\|\mathbf{w}\|_{\infty} = 1$, as in that case $\|\mathbf{x}\|_{\infty} \leq \|\mathbf{x}\|_{\infty}^{\mathbf{w}}$, and for some \mathbf{x} we actually have equality. Then, we

⁴The bound is actually *very* crude, in particular on reducible matrices when σ is close to $\rho(A)$.

can store in few bits an approximate vector $\mathbf{w}' \leq \mathbf{w}$, which can be used to estimate $\|\bar{\mathbf{x}}^{(t+1)} - \mathbf{x}^{(t)}\|_\infty^{\mathbf{w}}$, as we have, using the notation of Theorem 1,

$$\|\bar{\mathbf{x}} - \mathbf{x}^{(t+1)}\|_\infty \leq \|\bar{\mathbf{x}} - \mathbf{x}^{(t+1)}\|_\infty^{\mathbf{w}} \leq \frac{r}{1-r} \|\mathbf{x}^{(t+1)} - \mathbf{x}^{(t)}\|_\infty^{\mathbf{w}} \leq \frac{r}{1-r} \|\mathbf{x}^{(t+1)} - \mathbf{x}^{(t)}\|_\infty^{\mathbf{w}'}$$

A reasonable choice is that of keeping in memory $\lceil -\log_2 w_i \rceil$. Using a byte of storage we can keep track of w_i 's no smaller than 2^{-8} . Moreover, during the evaluation of the norm we just have to multiply by a power of two, which can be done very quickly in IEEE 754 format.

5 Choosing a suitable vector

We now come to the main result: given a nonnegative matrix A and a $\sigma > \rho(A)$, it is possible (constructively) to compute a vector \mathbf{w} that is σ -suitable for A . In essence, the computation of a σ -suitable vector for A “tames” the non-normality of the iterative process, at the price of a reduction of the convergence range.

Theorem 2 *Let A be nonnegative and $\sigma > \rho(A)$. Let*

$$\mathbf{w}_\sigma^{(k)} = \sum_{i=0}^k \left(\frac{A}{\sigma}\right)^i \mathbf{1}$$

and

$$\mathbf{w}_\sigma = \lim_{k \rightarrow \infty} \mathbf{w}_\sigma^{(k)}$$

Then, $A\mathbf{w}_\sigma < \sigma\mathbf{w}_\sigma$. In particular, \mathbf{w}_σ is σ -suitable for A , and there is a k such that

$$A\mathbf{w}_\sigma^{(k)} \leq \sigma\mathbf{w}_\sigma^{(k)},$$

so $\mathbf{w}_\sigma^{(k)}$ is σ -suitable for A .

Proof. Consider the matrix $A + \delta\mathbf{1}\mathbf{1}^*$, where $\delta > 0$. Since it is strictly positive, the Perron–Frobenius Theorem tells us that there is a dominant eigenvector $\mathbf{w}_\delta > 0$. Moreover, since for $\delta \rightarrow \infty$ we have $\rho(A + \delta\mathbf{1}\mathbf{1}^*) \rightarrow \infty$, and the spectral radius is continuous in the matrix entries, there must be a δ_σ such that

$$\rho(A + \delta_\sigma\mathbf{1}\mathbf{1}^*) = \sigma.$$

We have

$$\begin{aligned} (A + \delta_\sigma\mathbf{1}\mathbf{1}^*)\mathbf{w}_{\delta_\sigma} &= \sigma\mathbf{w}_{\delta_\sigma} \\ A\mathbf{w}_{\delta_\sigma} + \delta_\sigma\|\mathbf{w}_{\delta_\sigma}\|_1\mathbf{1} &= \sigma\mathbf{w}_{\delta_\sigma} \\ \frac{\delta_\sigma\|\mathbf{w}_{\delta_\sigma}\|_1}{\sigma}\mathbf{1} &= \left(1 - \frac{A}{\sigma}\right)\mathbf{w}_{\delta_\sigma} \\ \mathbf{w}_{\delta_\sigma} &= \frac{\delta_\sigma\|\mathbf{w}_{\delta_\sigma}\|_1}{\sigma} \sum_{i=0}^{\infty} \left(\frac{A}{\sigma}\right)^i \mathbf{1}. \end{aligned}$$

We now observe that the scaling factor is irrelevant: $\mathbf{w}_{\delta_\sigma}$ is an eigenvector, so it is defined up to a multiplicative constant. We can thus just write

$$\mathbf{w}_\sigma = \sum_{i=0}^{\infty} \left(\frac{A}{\sigma}\right)^i \mathbf{1}$$

and state that

$$(A + \delta_\sigma \mathbf{1}\mathbf{1}^*) \mathbf{w}_\sigma = \sigma \mathbf{w}_\sigma,$$

which implies

$$A \mathbf{w}_\sigma = \sigma \mathbf{w}_\sigma - \delta_\sigma \|\mathbf{w}_\sigma\|_1 \mathbf{1} < \sigma \mathbf{w}_\sigma.$$

Thus, as $\mathbf{w}_\sigma^{(k)} \rightarrow \mathbf{w}_\sigma$ when $k \rightarrow \infty$, for some k we must have

$$A \mathbf{w}_\sigma^{(k)} \leq \sigma \mathbf{w}_\sigma^{(k)}.$$

■

The previous theorem suggests the following procedure. Under the given hypotheses, start with $\mathbf{w}^{(0)} = \mathbf{1}$, and iterate

$$\begin{aligned} \mathbf{z} &= A \mathbf{w}^{(t)} \\ \mathbf{w}^{(t+1)} &= \mathbf{z} / \sigma + \mathbf{1}. \end{aligned}$$

Note that this is just a Jacobi iteration for the problem $(I - A/\sigma)\mathbf{x} = \mathbf{1}$, which is natural, as \mathbf{w}_σ is just its solution. The iteration stops as soon as

$$\max_i \frac{z_i}{w_i^{(t)}} \leq \sigma, \tag{4}$$

and at that point $\mathbf{w}^{(t)}$ is by definition σ -suitable for A , so we can apply Theorem 1. In practice, it is useful to keep the current vector $\mathbf{w}^{(t)}$ normalized: just set $s^{(0)} = 1$ at the start, and then iterate

$$\begin{aligned} \mathbf{z} &= A \mathbf{w}^{(t)} \\ \mathbf{u} &= \mathbf{z} / \sigma + s^{(t)} \mathbf{1} \\ s^{(t+1)} &= s^{(t)} / \|\mathbf{u}\|_\infty \\ \mathbf{w}^{(t+1)} &= \mathbf{u} / \|\mathbf{u}\|_\infty. \end{aligned}$$

We remark that, albeit used for clarity in the statement of Theorem 2, the (exact) knowledge of $\rho(A)$ is not strictly necessary to apply the technique above: indeed, if the procedure terminates $\sigma \geq \rho(A)$ by Proposition 1.

There are a few useful observations about the behavior of the normalized version of the procedure. First, if $\sigma < \rho(A)$ necessarily $s^{(t)} \rightarrow 0$ as $t \rightarrow \infty$. Second, by Collatz's classical bound [Col42], the maximum in (4) is an upper bound to $\rho(A)$. This happens without additional hypotheses⁵ on A because whenever $A\mathbf{x} \leq \gamma\mathbf{x}$ with $\mathbf{x} > 0$ we have

$$\rho(A) \leq \|A\|_\infty^{\mathbf{x}} = \|A\mathbf{x}\|_\infty^{\mathbf{x}} \leq \|\gamma\mathbf{x}\|_\infty^{\mathbf{x}} = \gamma.$$

⁵We report the following two easy proofs as in most of the literature Collatz's bounds are proved for irreducible matrices using Perron–Frobenius theory.

If, moreover, we compute also the minimum ratio

$$\min_i \frac{z_i}{w_i^{(t)}}, \quad (5)$$

this is a lower bound to $\rho(A)$, again without additional hypotheses on A . Indeed, note that whenever $\beta \mathbf{x} \leq A \mathbf{x}$ with $\mathbf{x} \geq 0$, for every $\delta > 0$ if \mathbf{w} is a positive eigenvector of $A + \delta \mathbf{1}\mathbf{1}^*$ we have

$$\beta \mathbf{x} \leq A \mathbf{x} \leq (A + \delta \mathbf{1}\mathbf{1}^*) \mathbf{x} \leq (A + \delta \mathbf{1}\mathbf{1}^*) \|\mathbf{x}\|_{\infty}^{\mathbf{w}} \mathbf{w} = \rho(A + \delta \mathbf{1}\mathbf{1}^*) \|\mathbf{x}\|_{\infty}^{\mathbf{w}} \mathbf{w}.$$

The last inequality implies $\beta \leq \rho(A + \delta \mathbf{1}\mathbf{1}^*)$ by Proposition 1.6, and since the inequality is true for every δ it is true by continuity also for $\delta = 0$.

These properties suggest that in practice iteration should be stopped if $s^{(t)}$ goes below the minimum representable floating-point number: in this case, either $\sigma < \rho(A)$, or the finite precision at our disposal is not sufficient to compute a suitable vector because we cannot represent correctly a transient behavior of the powers of A .

If instead the minimum (5) becomes larger than σ , we can safely stop: unfortunately, the latter event cannot be guaranteed to happen when $\sigma < \rho(A)$ without additional hypotheses on A (e.g., irreducibility): for instance, if A has a null row the minimum (5) will always be equal to zero.

Of course, there ain't no such thing as a free lunch. The termination of the process above is guaranteed if $\sigma > \rho(A)$, but we have no indication of how many step will be required. Moreover, in principle some of the coordinates of the suitable vector could be so small to make Theorem 1 unusable. For σ close to $\rho(A)$ convergence can be very slow, as it is related to the convergence of Collatz's lower and upper bounds for the dominant eigenvalue.

Nonetheless, albeit all of the above must happen in pathological cases, we show on a few examples that, actually, in real-world cases computing a σ -suitable vector is not difficult.

We remark that in principle any dyadic product $\mathbf{u}\mathbf{v}^*$ such that $A + \mathbf{u}\mathbf{v}^*$ is irreducible will do the job in the proof of Theorem 2. There might be choices (possibly depending on A) for which the computation above terminates more quickly.

6 Examples

6.1 Bounding the error of $(I - A)\mathbf{x} = \mathbf{b}$

If A is nonnegative matrix with $\rho(A) < 1$, then $I - A$ is invertible and the problem $(I - A)\mathbf{x} = \mathbf{b}$ has a unique solution, and in the limit we have convergence geometric in $\rho(A)$. However, if we choose a $1 > \sigma > \rho(A)$ (say, $\sigma = (1 + \rho(A))/2$) and a σ -suitable vector \mathbf{w} , the bounds of Theorem 1 will be valid, so we will be able to control the error in \mathbf{w} -norm.

6.2 Katz's index

Let M be a nonnegative matrix (in the standard formulation, the adjacency matrix of a graph). Then, given $\alpha < 1/\rho(M)$ Katz's index is defined by

$$\mathbf{k}^* = \mathbf{v}^*(1 - \alpha M)^{-1} = \mathbf{v}^* \sum_{k \geq 0} \alpha^k M^k,$$

where \mathbf{v} is a *preference vector*, which is just $\mathbf{1}$ in Katz's original definition [Kat53].⁶ If we want to apply Theorem 1, we must choose a $\sigma > \rho(A)$ and a σ -suitable vector \mathbf{w} for A . The vector can then be used to accurately estimate the computation of Katz's index for all $\alpha < 1/\sigma$. This property is particularly useful, as it is common to estimate the index for different values of α , and to that purpose it is sufficient to compute once for all a σ -suitable vector for a σ chosen sufficiently close to $\rho(A)$.

6.3 PageRank

The case of PageRank is similar to Katz's index. We have

$$\mathbf{r}^* = (1 - \alpha)\mathbf{v}^*(1 - \alpha P)^{-1} = (1 - \alpha)\mathbf{v}^* \sum_{k=0}^{\infty} \alpha^k P^k,$$

where \mathbf{v} is the preference vector, and $P = \bar{G} + \mathbf{d}\mathbf{u}^*$ is a stochastic matrix; \bar{G} is the adjacency matrix of a graph G , normalized so that each nonnull row adds to one, \mathbf{d} is the characteristic vector of *dangling nodes* (nodes without outlinks, i.e., null rows), and \mathbf{u} is the dangling-node distribution, used to redistribute the rank lost through dangling nodes. It is common to use a uniform \mathbf{u} , but most often $\mathbf{u} = \mathbf{v}$, and in that case we speak of *strongly preferential* PageRank [BSV09].

We remark that in the latter case it is well known that the *pseudorank*

$$\mathbf{p}^* = (1 - \alpha)\mathbf{v}^* \sum_{k=0}^{\infty} \alpha^k \bar{G}^k$$

satisfies

$$\mathbf{r} = \frac{\mathbf{p}}{\|\mathbf{p}\|_1}.$$

That is, PageRank and the pseudorank are parallel vectors. This is relevant for the computation of several strongly preferential PageRank vectors: just compute a σ -suitable vector for \bar{G} (rather than one for each $\bar{G} + \mathbf{d}\mathbf{v}^*$, depending on \mathbf{v}), and compute pseudoranks instead of ranks.

The case of PageRank is however less interesting because, as David Gleich made the author note, assuming the notation of Section 2 and $\omega = 1$

$$\begin{aligned} (1 - \alpha P^T)\mathbf{x}^{(t+1)} - (1 - \alpha)\mathbf{v} &= (D - L^{(t)} - R^{(t)})\mathbf{x}^{(t+1)} - (1 - \alpha)\mathbf{v} \\ &= (D - L^{(t)})\mathbf{x}^{(t+1)} - R^{(t)}\mathbf{x}^{(t+1)} - (1 - \alpha)\mathbf{v} \\ &= R^{(t)}\mathbf{x}^{(t)} + (1 - \alpha)\mathbf{v} - R^{(t)}\mathbf{x}^{(t+1)} - (1 - \alpha)\mathbf{v} \\ &= R^{(t)}(\mathbf{x}^{(t)} - \mathbf{x}^{(t+1)}). \end{aligned}$$

⁶We must note that actually Katz's index is $\mathbf{v}^*(1 - \alpha M)^{-1}M$. This additional multiplication by M is somewhat common in the literature; it is probably a case of *horror vacui*.

	Wikipedia	.uk
nodes	4 206 785	105 896 555
arcs	101 355 853	3 738 733 648
avg. degree	24.093	35.306
giant component	89.00%	64.76%
harmonic diameter	5.24	22.78
dominant eigenvalue	191.11	5676.63

Table 1: Basic structural data about our two datasets.

Since $\|R^{(t)}\|_1 \leq \alpha$, we can ℓ_1 -bound the residual

$$\|(1 - \alpha P^T)^{-1}\|_1 = \left\| \sum_{k=0}^{\infty} \alpha^k (P^T)^k \right\|_1 \leq \frac{1}{1 - \alpha}$$

we conclude that

$$\|\bar{\mathbf{x}} - \mathbf{x}^{(t+1)}\|_1 \leq \frac{\alpha}{1 - \alpha} \|\mathbf{x}^{(t+1)} - \mathbf{x}^{(t)}\|_1.$$

It is thus possible, albeit wasteful, to bound the supremum norm of the error using its ℓ_1 norm.

7 Experiments

In this section we discuss some computational experiments involving the computation of PageRank and Katz’s index on real-world graphs. We focus on a snapshot of the English version of Wikipedia taken in 2013 (about four million nodes and one hundred million arcs) and a snapshot of the .uk web domain taken in may 2007 (about one hundred million nodes and almost four billion arcs).⁷ These two graphs have some structural differences, which we highlight in Table 1.

We applied the procedure described in Section 5 to the system associated with PageRank and Katz’s index, with $\sigma \in \{1/(1 - 2^{-i}) \mid 1 \leq i \leq 10\}$ for PageRank and $\sigma \in \{\lambda/(1 - 2^{-i}) \mid 1 \leq i \leq 10\}$ for Katz’s index.

In Figure 1 we report the number of iterations that are necessary to compute the i -th suitable vector. The two datasets show the same behavior in the case of PageRank—an exponential increase in the number of iterations as we get exponentially closer to the limit value. The case of Katz is more varied: whereas Wikipedia has a significant growth in the number of iterations (but clearly slower than the PageRank case), .uk has a minimal variation across the range (from 2 to 6).

In Figure 2 we draw the (exponentially binned) distribution of values of suitable vectors for a choice of four equispaced values of i . The vectors are normalized in ℓ_∞ norm, that is, the largest value is one.

The shape of the distribution depends both on the graph and on the type of centrality computed, but two features are constant: first, as we approach λ the distribution

⁷Both datasets are publicly available at the site of the Laboratory for Web Algorithmics (<http://law.di.unimi.it/>) under the identifiers `enwiki-2013` and `uk-2007-05`.




Figure 1: Number of iterations that are necessary to compute a $\lambda/(1-2^{-i})$ -suitable vector.

contains smaller and smaller values; second, the smallest value in the PageRank case is several orders of magnitude smaller.

Smaller values imply a larger \mathbf{w} -norm: indeed, one can think of the elements of an ℓ_∞ -normalized suitable vector \mathbf{w} as weights that “slow down” the convergence of problematic nodes by inflating their raw error. The intuition we gather from the distribution of values is that bounding the convergence of PageRank is more difficult.

8 Conclusions

We have presented results that make it possible to bound the supremum norm of the absolute error of SOR iterations an M -matrix $sI - A$ even when estimating $\|(sI - A)^{-1}\|_\infty$ is not feasible. Rather than relying on additional hypotheses such as positive definiteness, irreducibility and so on, our results suggest to compute first a σ -suitable positive vector \mathbf{w} with the property that SOR iterations converge geometrically in \mathbf{w} -norm by a computable factor.

While we cannot bound without additional hypotheses the resources (number of iterations and precision) that are necessary to compute \mathbf{w} , in practice the computation is not difficult, and given an M -matrix $sI - A$ the associated σ -suitable \mathbf{w} can be used for all $s > \sigma$.

References

- [BP94] Abraham Berman and Robert J. Plemmons. *Nonnegative Matrices in the Mathematical Sciences*. Classics in Applied Mathematics. SIAM, 1994.
- [BPSV08] Paolo Boldi, Roberto Posenato, Massimo Santini, and Sebastiano Vigna. Traps and pitfalls of topic-biased PageRank. In William Aiello, Andrei Broder, Jeannette Janssen, and Evangelos Milios, editors, *WAW*




Figure 2: Exponentially binned frequency plots of the values of $\lambda/(1 - 2^{-i})$ -suitable vectors, $i = 1, 4, 7$ and 10 .

2006. *Fourth Workshop on Algorithms and Models for the Web-Graph*, volume 4936 of *Lecture Notes in Computer Science*, pages 107–116. Springer–Verlag, 2008.

- [BSV09] Paolo Boldi, Massimo Santini, and Sebastiano Vigna. PageRank: Functional dependencies. *ACM Trans. Inf. Sys.*, 27(4):1–23, 2009.
- [BT89] Dimitri P. Bertsekas and John N. Tsitsiklis. *Parallel and Distributed Computation: Numerical Methods*. Prentice Hall, Englewood Cliffs NJ, 1989.
- [Col42] Lothar Collatz. Einschließungssatz für die charakteristischen Zahlen von Matrizen. *Mathematische Zeitschrift*, 48(1):221–226, 1942.
- [GK11] Chen Greif and David Kurokawa. A note on the convergence of SOR for the PageRank problem. *SIAM J. Sci. Computing*, 33(6):3201–3209, 2011.
- [Kat53] Leo Katz. A new status index derived from sociometric analysis. *Psychometrika*, 18(1):39–43, 1953.
- [Ken45] Maurice G. Kendall. The treatment of ties in ranking problems. *Biometrika*, 33(3):239–251, 1945.
- [PBMW98] Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. The PageRank citation ranking: Bringing order to the web. Technical Report SIDL-WP-1999-0120, Stanford Digital Library Technologies Project, Stanford University, 1998.

[Vig09] Sebastiano Vigna. Spectral ranking. *CoRR*, abs/0912.0238, 2009.