FAST HYBRID SOLUTION OF ALGEBRAIC SYSTEMS*

CRAIG C. DOUGLAS†, JAN MANDEL‡ AND WILLARD L. MIRANKER§

Abstract. We propose and analyze the error and timing of solvers consisting of both analog and digital circuitry for sparse linear systems of equations. We obtain high speed, but low precision from the analog circuits. We combine this with low speed, but high precision from the digital circuits. The hybrid circuit should be faster than digital circuits alone. As a preconditioner to standard iterative solution methods, the hybrid circuit makes the cost of the preconditioning step negligible. We also apply the hybrid circuit to a standard multilevel algorithm.

Key words. digital-analog computing, linear systems of equations, multigrid methods

AMS(MOS) subject classifications. 35, 65

1. Introduction. We study a fast equation solver consisting of both analog and digital circuitry. We expect this hybrid combination to give better results than digital techniques alone. The basic idea is to use an analog solver as a preconditioner for a digital iterative process. For a related study, see [12]. Thus, we can obtain both high speed from a fast exchange of information in analog circuitry and high precision from digital circuitry. Eventually, both types of circuits should be integrated onto a single chip.

In §2, we define an analog defect correction algorithm and discuss the sources of error. We also provide an error analysis. In §3, we define a basic model for a simple analog solver. We analyze its response speed and its precision. A general example is examined in detail. In §4, we define and motivate a two-stage analog solver. As in §3, we analyze it and examine an example. Finally, in §5, we define and analyze a multilevel solver. We use the term multilevel in the abstract multigrid solution of partial differential equation sense.

The applicability of the method is essentially limited by the condition $\epsilon\kappa \ll 1$, where ϵ is the relative precision of the analog circuitry and κ is the condition number of the linear system. For the multilevel solver, the condition number involved is the one for the linear system on the coarsest level. The time required for the analog part of the method also depends on the condition number. We conclude that this time is negligible in comparison to that for the digital part of the method when $\epsilon\kappa \ll 1$.

Due to ϵ being technology dependent, the limit of this theory is currently $\kappa \leq 1024$. A sequel to this paper[9] considers preconditionings and modifications to the analog-

* Received by the editors XXXXXXXX XX, 1988; accepted for publication (in revised form) XXXXXXXX XX, 1989.

† Mathematical Sciences Department, IBM Research Division, Thomas J. Watson Research Center, P. O. Box 218, Yorktown Heights, New York 10598. E-mail: *bells@ibm.com*.

‡ Department of Mathematics, Computational Math Group, Department of Mathematics, University of Colorado at Denver, Denver, CO 80202. This research was done while visiting IBM Research Division, Thomas J. Watson Research Center, Yorktown Heights, New York.

§ Mathematical Sciences Department, IBM Research Division, Thomas J. Watson Research Center, P. O. Box 218, Yorktown Heights, New York 10598.

digital algorithm in §2 to apply this theory to problems where $\kappa > 1024$.

We expect the principal benefits of the proposed method to manifest themselves with advances in technology: analog circuitry has the potential to avoid the information exchange bottleneck of massively parallel digital computation. Essentially, we are trading (recoverable) precision for fast dissemination of information between the processors.

We expect that these techniques will be advantageous for large but moderately conditioned positive definite problems with well defined sparsity structures. Systems arising by either finite element or finite difference discretizations of partial differential equation problems are one possible application. For a general, non-sparse system, the number of connections required is prohibitive.

The technique used in the analysis is classical and can be found in any electrical engineering textbook. We believe the defect correction approach to the hybrid method is new and offers, as yet, unexplored possibilities in massive parallel computation. For related studies, see [11, 2, 7].

The precision of current analog circuitry is up to 10 bits[13] using capacitors as basic circuit components. Optical processors have the same or lower precision[10]. A purely optical analog method for solving linear systems was presented in[4]. See[6] for a survey of basic concepts of optical computing.

2. Analog Defect Correction: Error analysis. Consider a system of linear equations in matrix notation:

$$Ax^* = f,$$

where A is an $n \times n$ nonsingular matrix and x^* is the exact solution. We will sometimes require the hypothesis that A is symmetric, positive definite. In the following algorithm, we use a standard residual correction technique to solve the system, except that part of the computation is performed digitally and part is performed with an analog solver.

The Hybrid Algorithm (Iterative)

Step 1. For given x_1, compute $r = f - Ax_1$, digitally.

Step 2. Convert r to n parallel analog signals, and using an analog solver, solve the equation

$$Ay = r.$$

Convert y (in parallel) to digital output.

Step 3. Compute $x_2 = x_1 + y$, digitally.

Step 4. Set $x_1 \leftarrow x_2$ and go to Step 1.

We can analyze the error reduction per step using techniques found in [14]. Assume that the precision of digital computation is very high relative to the analog computation. The quality of the analog computation is described by the following backward characterization:

$$y = (A - E)^{-1} r + e, \tag{2.1}$$

where E is a perturbation matrix and e a perturbation vector such that

$$\|e\| \leq \epsilon \|(A - E)^{-1} r\| \tag{2.2}$$

and

$$\|E\| \leq \epsilon \|A\|. \tag{2.3}$$

Here, $\|\cdot\|$ is a suitable norm (the ℓ_∞ norm norm will be appropriate here) and ϵ is the relative precision of analog circuitry.

Let $e_1 = x_1 - x^*$ and $e_2 = x_2 - x^*$. Then from the Hybrid Algorithm,

$$r \;=\; -Ae_1, \tag{2.4}$$

so by (2.1),

$$y \;=\; -(A - E)^{-1} A e_1 + e.$$

Hence

$$e_2 \;=\; e_1 - (A - E)^{-1} A e_1 + e. \tag{2.5}$$

So,

$$\|I - (A - E)^{-1} A\| \;\leq\; \frac{\epsilon \kappa}{1 - \epsilon \kappa}, \tag{2.6}$$

where κ is the condition number of A in the $\|\cdot\|$ norm,

$$\kappa \;=\; \|A^{-1}\| \; \|A\|.$$

Further,

$$\|(A - E)^{-1} A\| \;\leq\; \frac{1}{1 - \epsilon \kappa}.$$

Combining this with (2.4) in (2.2), we get

$$\|e\| \leq \frac{\epsilon}{1 - \epsilon \kappa} \|e_1\|.$$

Using this estimate in (2.5), we obtain from (2.6) that

$$\|e_2\| \;\leq\; \epsilon \frac{\kappa + 1}{1 - \epsilon \kappa} \|e_1\|. \tag{2.7}$$

This estimate furnishes a good error bound if

$$\epsilon\kappa \ll 1. \tag{2.8}$$

Then we can expect each cycle of the algorithm to reduce the error by at least $\epsilon\kappa$ (approximately).

From (2.7) we see that the effect of e on the error propagation is small compared to the effect of E. Then in (2.5) we formally drop e, viewing its effect to be absorbed into E. We shall see that the error matrix E is of the form

$$E = E_\epsilon - \mu_0^{-1} I$$

for a simple analog model. Here E_ϵ is the error in analog representation of the entries of the matrix A and $\mu_0^{-1} I$ models the effect of finite amplification in the circuitry. Incorporating these observations into (2.5) gives

$$e_2 = M e_1,$$

where

$$M = I - [\mu_0^{-1} I + A - E_\epsilon]^{-1} A. \tag{2.9}$$

Remark on sources of error: Analog computation will have three principle sources of error:

1. Digital-analog and analog-digital conversion of input and output (contained in (2.2)).
2. Digital-analog conversion error in representation of the matrix A (contained in (2.3)).
3. Effects of finite amplification and finite time (contributing to both (2.2) and (2.3)).

The third source of error will be analyzed in the next section.

In the following section we show that the time spent on the analog part of the computation is negligible in comparison to the digital part (if $\epsilon\kappa \ll 1$. When A is symmetric, positive definite, the overrelaxed Jacobi method,

$$u \leftarrow u - \omega(Au - f),$$

has a convergence factor of

$$\frac{\kappa - 1}{\kappa + 1}$$

(when the optimal ω is used). When the Hybrid and Jacobi methods are fully parallelized, the cost of one iteration of either method is approximately equal. The interesting cases in practice for the Hybrid method are when both $\epsilon\kappa \ll 1$ and $\kappa \gg 1$, so that the Hybrid method is much faster than Jacobi.

TABLE 1
Contraction Factors for the Hybrid Method

$\epsilon \setminus \kappa$	50	100	200	400
1/50	1.0000	2.0000	4.0000	8.0000
1/100	0.5000	1.0000	2.0000	4.0000
1/200	0.2500	0.5000	1.0000	2.0000
1/400	0.1250	0.2500	0.5000	1.0000
1/800	0.0625	0.1250	0.2500	0.5000

TABLE 2
Contraction Factors for Optimally Overrelaxed Jacobi

$- \setminus \kappa$	50	100	200	400
	0.9608	0.9802	0.9900	0.9950

Tables 1 and 2 contain the contraction factors (the error reduction per iteration factors) for the Hybrid and Jacobi methods, respectively, for some sample values of ϵ and κ. The ratio of logarithms of contraction factors, $\log(\epsilon\kappa)/\log((\kappa-1)/(\kappa+1))$, equals the ratio of the number of iterations required to attain a specified precision. Table 3 contains the ratios when the Hybrid method converges, i.e., when $\epsilon\kappa < 1$. The numbers represent speedup factors of the Hybrid method over a comparable fully parallel (digital) simple iterative method, since an iteration of each method requires the same time. As the table demonstrates, the Hybrid method can be more than a factor of 100 faster than optimally overrelaxed Jacobi. More sophisticated digital methods (e.g., conjugate gradients) are faster than Jacobi, but are not usually fully parallelizable.

A great deal of research in the past twenty years has been devoted to developing useful preconditioners for (digital) iterative methods (see [5]). Generally, an approximation to A is constructed which is close to A (in some sense), but much easier to factor. The corresponding preconditioned iterative method converges faster than the original method, but costs more per iteration. These preconditioners typically reduce the error by considerably less than a digit per iteration. The analog solve step can be thought of as a preconditioning step, where the preconditioner is the original matrix A rounded off to nearly three digits. We can do this because we can prove that the analog step takes almost no time (in comparison to the digital steps). Table 1 demonstrates when we can expect to reduce the error by at least one digit per iteration.

3. A Simple Analog Solver.

3.1. Basic model. The basic component of an analog circuit is an amplifier:

Table 3

Ratio of Logarithms of Contraction Factors

$\epsilon \setminus \kappa$	50	100	200	400
1/50	-	-	-	-
1/100	17	-	-	-
1/200	35	35	-	-
1/400	52	69	69	-
1/800	69	104	139	139

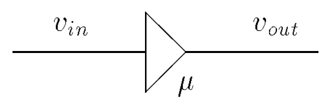

In the simplest approximation, the signals v_{in} and v_{out} satisfy the differential equation

$$\left(c\,\frac{d}{dt}+1\right) v_{\text{out}} \;=\; -\mu_0 v_{\text{in}}, \quad t \geq t_0 \tag{3.1}$$

where μ_0 is the steady state gain, c is the time constant of the amplifier and t_0 is the starting time. (Solid state amplifiers with $c < 10^{-7}$s and $\mu_0 > 10^4$ are currently available.)

Remark: In more general models, the differential operator $c\,\frac{d}{dt}+1$ in (3.1) is replaced by a polynomial in $\frac{d}{dt}$ or, more generally, by an analytic function of $\frac{d}{dt}$, i.e., a pseudo-differential operator.

The transmission function μ of the amplifier is defined as

$$\mu \;=\; \mu(\omega) \;=\; \frac{v_{\text{out}}}{v_{\text{in}}}, \quad v_{\text{in}} \;=\; e^{j\omega t},$$

where ω is the angular frequency and j is the imaginary unit. For (3.1),

$$\mu \;=\; \frac{-\mu_0}{1+cj\omega},$$

the so-called *one-pole* transmission function. (This relates our approach to conventional engineering terminology.)

Consider the network in Figure 1. Here, the amplifier part consists of n identical amplifiers acting on n signals in parallel. Each amplifier is assumed to have the same transmission function $\mu = \mu(\omega)$ corresponding to a linear differential operator M:

$$\mu(\omega) \;=\; \frac{-\mu_0 e^{j\omega t}}{M e^{j\omega t}}.$$

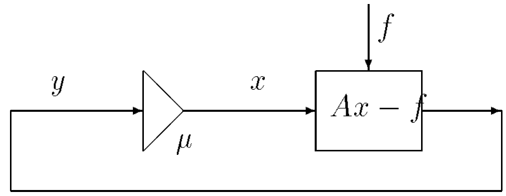

FIG. 1. *One Component in a One-Pole Network*

The output x of the amplifiers is processed by a passive network implementing multiplication by A, and then the residual $Ax - f$ is fed back to the input. In fact, this residual determination will be merged with the amplifiers into one circuit; Figure 1 presents just a convenient equivalent model.

Remark: Note that the magnitude of the elements of A is limited since A cannot contain any amplification.

The state variable x satisfies the system of ordinary differential equations

$$Mx \;=\; -\mu_0(Ax - f).$$

In particular, for the one pole model (3.1), we have

$$\left(c\frac{d}{dt} + 1\right) x \;=\; -\mu_0(Ax - f). \tag{3.2}$$

Now it is obvious that for $f = 0$ and $Ax(t_0) = \lambda x(t_0) \neq 0$, we have

$$Mx = -\mu_0 \lambda x, \quad t \geq t_0.$$

Let λ is an eigenvalue of A. We can thus reduce stability considerations to stability of the circuits of the form

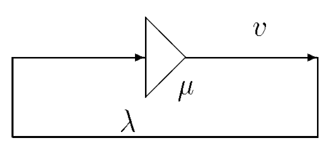

For the one-pole model, we get the equation for the state variables

$$c\dot{v} + v \;=\; -\mu_0 \lambda v,$$

where λ is an eigenvalue of A. Thus the one-pole model will be stable for all $\mu_0 > 0$ if and only if all real parts of the eigenvalues of A are positive.

3.2. Response Speed and Precision. Here we consider the one-pole model only. The state variable x satisfies the first-order system (3.2). To estimate the response, suppose that f is constant and that $x(0) = 0$. Then (3.2) has the solution

$$x(t) \;=\; \bar{x} - \exp[-\frac{1}{c}(I + \mu_0 A)t]\bar{x},$$

where $\bar{x}$ is the steady state,

$$\bar{x} \;=\; \left(\frac{1}{\mu_0}I + A\right)^{-1} f. \tag{3.3}$$

Now assume that A is symmetric, positive definite and let $\|\cdot\|$ be the ℓ_2 norm. Then for the transient part of x, we have

$$\|\exp[-\frac{1}{c}(I + \mu_0 A)t]\bar{x}\| \le e^{-\frac{1}{c}(1+\mu_0\lambda_{\min})t} \|\bar{x}\|.$$

The response time t_ϵ for relative precision ϵ is given by

$$-\frac{1}{c}(1 + \mu_0\lambda_{\min})t_\epsilon \;=\; \ln \epsilon,$$

or

$$t_\epsilon \;\approx\; \frac{c \ln \frac{1}{\epsilon}}{\mu_0\lambda_{\min}}.$$

Because c is the time constant of the fast analog, meaningfully fast response is assured by the requirement $\mu_0\lambda_{\min} \gg 1$

Example: Consider the typical values $\lambda_{\min} = 10^{-3}$, $\epsilon = 10^{-4}$, $\mu = 10^5$, and $c = 10^{-7}$s . Then

$$t_\epsilon \;\approx\; \frac{10^{-7} \cdot 9}{10^5 \cdot 10^{-3}}\text{s} \;=\; 9 \cdot 10^{-9} s.$$

If $\|A\|\;\|A^{-1}\| \approx 10^3$ in the ℓ_∞ norm, then we get from (2.8) that the error reduction per iteration will be bounded approximately by

$$\epsilon\kappa \;\approx\; 10^{-1}.$$

Conclusion: The total time required per iteration will be the sum of time for residual computation, digital-analog conversion time, a multiple of the time t_ϵ, analog-digital conversion time, and the time for addition. Because the time required for a digital operation can be assumed to be approximately c, the time constant of an analog circuit, we see that the time required for the analog solution itself is insignificant under the assumption that

$$(3.4) \qquad \mu_0 \lambda_{\min} \gg 1.$$

However, we expect the contraction factor limitation (2.8), viz

$$(3.5) \qquad \epsilon\kappa \ll 1,$$

to be critical, because the attainable precision ϵ of analog circuits is limited. The latter basically imposes a lower bound on $\lambda_{\min}$: assuming that

$$\kappa \;\approx\; \lambda_{\max}/\lambda_{\min}$$

and

$$\lambda_{\max} \;\approx\; 1,$$

we can write (3.5) as $\dfrac{\epsilon}{\lambda_{\min}} \ll 1$, or

$$(3.6) \qquad \epsilon \ll \lambda_{\min}.$$

Note that we further need $\|\dfrac{1}{\mu_0} I\| \leq \epsilon$, cf. (3.3) and (2.3). This condition is satisfied for typical parameters of contemporary devices. Thus, we require that

$$(3.7) \qquad \mu_0 \epsilon \tilde{>} 1.$$

Note that (3.6) and (3.7) implies (3.4). Thus the time t_ϵ will never be significant when our estimates for the hybrid method are applicable.

3.3. Sample Embodiment. We now consider a specific idealized circuit schema which embodies the analog part of the hybrid algorithm. We make use of classical devices: programmable resistors and operational amplifiers. The analog computational network will consist of n identical nodes as in Figure 2. The resistors should be capable of attaining the values 0 and ∞. Each node has $n+1$ inputs $x_1, \ldots, x_n$ (the components of x) and f_i (one component of f). The output is x_i. It is connected to all inputs x_i of all nodes. In a practical implementation, most of the connections and most of the resistors will be missing. A fixed sparsity structure of A will be assumed. Such a sparsity structure may correspond to the discretization of a problem on a 2-dimensional or 3-dimensional mesh, or it might be a band structure.

The output x_i is given by the transmission function of the operational amplifier,

$$x_i \;=\; \mu(v_+ - v_-).$$

Assuming zero output independence and infinite input independence of the operational amplifiers, the current balance at the inputs of the operational amplifier is

$$\frac{f_i - v_+}{R_{0i}} + \sum_j \frac{x_j - v_+}{R_{ji}} \;=\; \frac{v_+}{R_0}.$$

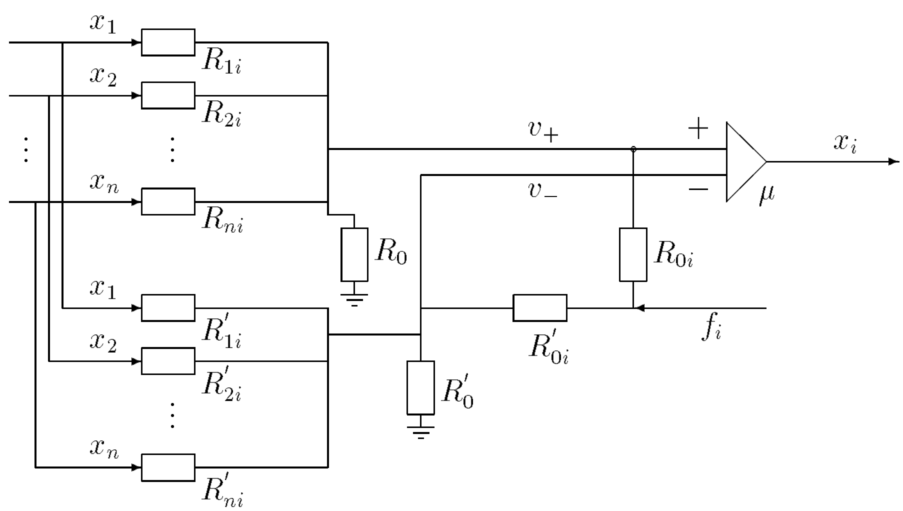

All quantities $x_1, \ldots, x_n, v_+, v_-, f_i$ are voltages.

FIG. 2. *One-Pole Sample Embodiment*

Then

$$v_+ \;=\; \frac{\sum_j \dfrac{x_j}{R_{ji}} + \dfrac{f_i}{R_{0i}}}{\dfrac{1}{R_0} + \sum_j \dfrac{1}{R_{ji}} + \dfrac{1}{R_{0i}}},$$

and similarly for v_-. We can thus implement the transmission function of the node:

$$x_i \;=\; -\mu\left(\sum_j a_{ij}x_j - f_i\right).$$

By expressing a_{ij} in terms of R_{ji} and R_0, and noting that resistances are nonnegative, we can show that

$$\sum_j \mid a_{ij} \mid < \; 2.$$

For a practical realization, a network using capacitors instead of resistors might be required. (Capacitors of high precision are easy to construct using MOS 1C technique.)

An entire hybrid circuit using 2-D mesh geometry could look similar to the one in Figure 3. An overflow/underflow detection two-line bus must be added to adjust the scaling of the residual fed into the analog solver. If the size of the analog output is too large, then the node which detects the overflow condition,

$$\mid x_i \mid > \; v_{\max},$$

will send a signal on bus line 1. Similarly, if a node detects the condition

$$\mid x_i \mid > v_{\min}, \quad v_{\min} < v_{\max}/a, \quad a > 1,$$

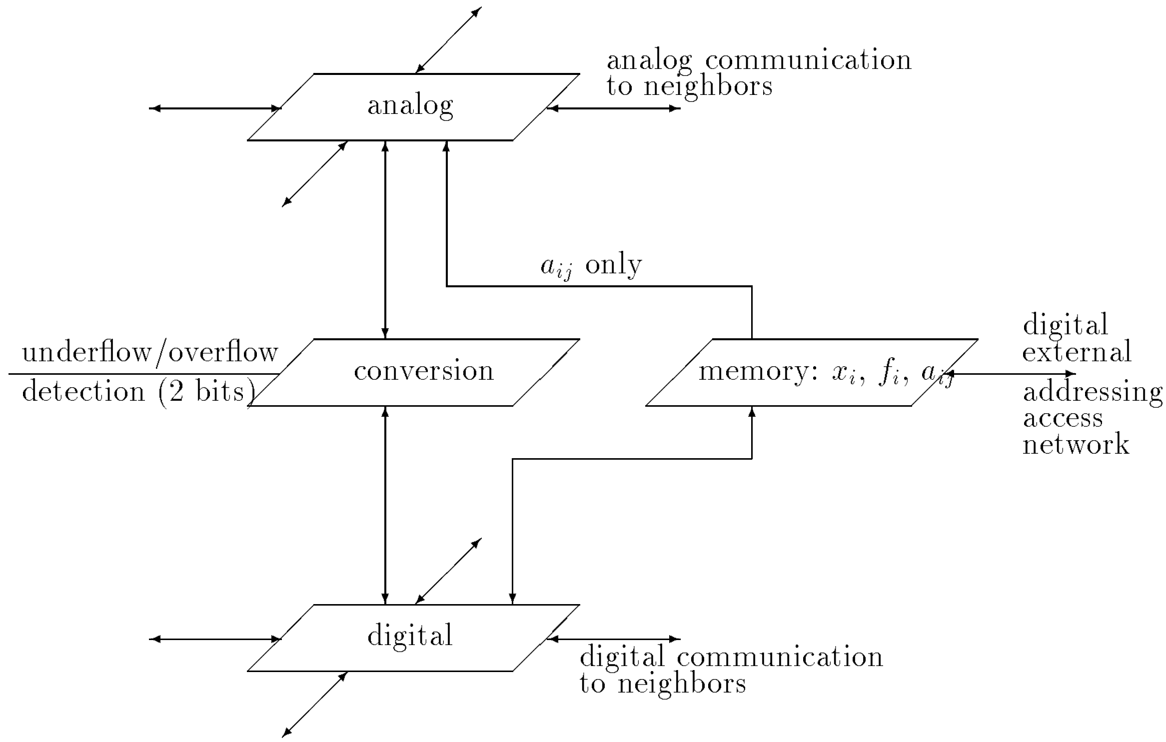

This is part of a repeating chip pattern.

FIG. 3. *Sample Hybrid Circuit for 2-D Mesh Geometry*

then it sends a signal on bus line 0. These bus lines are sensed by all conversion interface units. If bus line 1 is on, then all units decrease the analog right hand side by the factor $1/a$, $a > 1$. If bus line 0 is off, then all units increase it by the factor a. This will guarantee that no $\mid x_i \mid$ is larger than $v_{\max}$ and at least one is larger than $v_{\min}$, thus making full use of available precision. The scaling factor is then used in the output analog/digital conversion and stored for the next iteration as a good initial guess.

This scaling can be easily implemented by a voltage multiplier/attenuator at the output of the conversion unit. An analog bus, using continuous adjustments of the scale, could be also considered.

4. A Two-Stage Analog Circuit.

4.1. Motivation. In the one-stage circuit, the scale of the output x and of the input f is, in general, different. Since $x \approx A^{-1}f$, x will be much larger than f. This can be a source of errors. Therefore, we consider an implementation of the product Ax using another amplifier. The right hand side f is then easily combined with the output using a differential amplifier. The scaling of A can be used also to increase the speed of the circuit if necessary.

4.2. Basic Model. Consider the network consisting of two amplifier arrays, feedback array a, and a passive network for Ax:

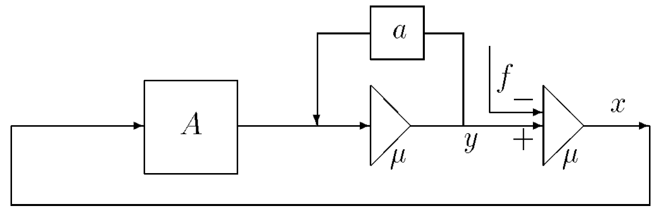

Assuming that all amplifiers have the same one pole transmission function

$$\mu = \frac{-\mu_0}{1 + cj\omega},$$

the constitutive equations of the network are

$$\left\{\begin{array}{rcl} c\dot{y} + y &=& -\mu_0(Ax + ay) \\ \\ c\dot{x} + x &=& -\mu_0(y + f). \end{array}\right. \tag{4.1}$$

Here, x, y, and f consist of n parallel analog signals on separate lines. The feedback factor a is assumed to be the same for all components.

4.3. Analysis. Equations (4.1) have the steady state solution (for constant f) given by

$$\left\{\begin{array}{rcl} \bar{y} &=& -\mu_0(A\bar{x} + a\bar{y}) \\ \\ \bar{x} &=& -\mu_0(\bar{y} + f). \end{array}\right.$$

The first equation gives $\bar{y} = \dfrac{-\mu_0}{1 + \mu_0 a} A\bar{x}$. Then

$$\bar{x} = \left(\frac{1}{\mu_0} I + \frac{\mu_0 A}{1 + \mu_0 a} \right)^{-1} f. \tag{4.2}$$

Thus $\lim_{\mu_0 \to \infty} \bar{x} = aA^{-1}f$. The transient part of the solution is a solution of the homogeneous system

$$\left\{\begin{array}{rcl} c\dot{y} &=& -(1 + \mu_0 a)y - \mu_0 Ax \\ \\ c\dot{x} &=& \mu_0 y - x, \end{array}\right. \tag{4.3}$$

with the initial condition being the error at $t = 0$. Introducing the matrices

$$G = \begin{pmatrix} aI & A \\ -I & 0 \end{pmatrix}, \quad H = -\frac{1}{c}[I + \mu_0 G],$$

we can write (4.3) as

$$\begin{pmatrix} \dot{y} \\ \dot{x} \end{pmatrix} = H \begin{pmatrix} y \\ x \end{pmatrix}.$$

If λ is an eigenvalue of A, then G has two corresponding eigenvalues φ obtained as the eigenvalues of the matrix $\begin{pmatrix} a & \lambda \\ -1 & 0 \end{pmatrix}$. In particular,

$$\varphi = \frac{a}{2} \pm \sqrt{\frac{a^2}{4} - \lambda}.$$

Thus the spectrum of H is

$$\sigma(H) = \left\{ -\frac{1}{c} \left[1 + \mu_0 \left(\frac{a}{2} \pm \sqrt{\frac{a^2}{4} - \lambda} \right) \right]; \quad \lambda \in \sigma(A) \right\}.$$

Suppose that A is symmetric and positive definite, and let $\lambda_{\min}$ denote its least eigenvalue. Then the element χ of $\sigma(H)$ with the largest real part (thus determining the response time) is equal to

$$\chi = -\frac{1}{c} \left[1 + \mu_0 \left(\frac{a}{2} - \sqrt{\frac{a^2}{4} - \lambda_{\min}} \right) \right].$$

The choice $a = 2\lambda_{\min}$ gives the transient response (assuming $\lambda_{\min} \leq 1$) for the slowest component

$$e^{Re\ \chi t} = e^{-\frac{1}{c}(1+\mu_0 \lambda_{\min})t}.$$

This is the same as for the one-stage circuit. However, by (4.2) the scaling of $\bar{x}$ is now

$$\bar{x} \approx aA^{-1}f = 2\lambda_{\min} A^{-1} f.$$

Thus we can expect $\|x\| \approx 2\|f\|$. The response speed and precision analysis of §3.2 holds for this case. The speed and precision are identical.

To speed up the transient response, we could use a larger feedback factor a. The choice

$$a = \sqrt{2\lambda_{\min}} \tag{4.4}$$

gives

$$-Re\ \chi = \frac{1}{c} \left[1 + \mu_0 \frac{a}{2} \right] = \frac{1}{c} \left[1 + \mu_0 \sqrt{\frac{\lambda_{\min}}{2}} \right].$$

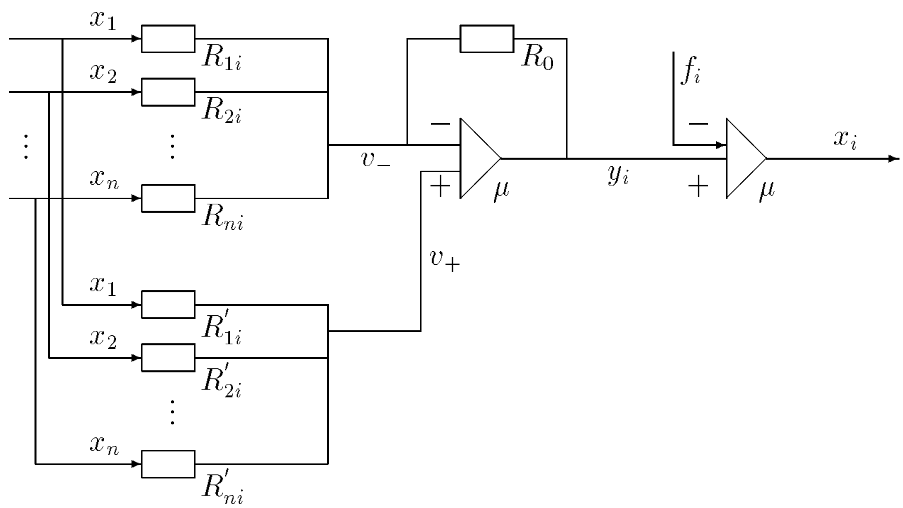

All quantities $x_1, \ldots, x_n, v_+, v_-, f_i$ are voltages.
All resistors are programable and can assume the values 0 and ∞.

FIG. 4. *Two-Stage Sample Embodiment*

Using*the technique from §3.2.

$$t_\epsilon \;\approx\; \frac{c \ln \epsilon}{\mu_0 \sqrt{\lambda_{\min}/2}}.$$

Meaningfully fast response is assured by the requirement $\mu_0\sqrt{\lambda_{\min}/2} \gg 1$.
Example: With the choice $a = \sqrt{2\lambda_{\min}}$ and considering the typical values $\lambda_{\min} = 10^{-3}$, $\epsilon = 10^{-4}$, $\mu = 10^5$, and $c = 10^{-7}$s once again, we have

$$t_\epsilon \;\approx\; \frac{10^{-7} \cdot 9}{10^5 \cdot 22.4}\text{s} \;\approx\; 4 \cdot 10^{-11}\text{s}.$$

However, for $a = 2\lambda_{\min}$, we obtain $t_\epsilon = 9 \cdot 10^{-9}$s. Thus choosing the feedback factor a according to (4.4) yields a significant speedup.
Remark: For the one-pole model of the amplifiers, the circuit is stable; however, compensation of the amplifiers so that they are well approximated by the one-pole model may be more critical here, because of both much stronger feedback and the presence of two amplifiers in the feedback loop.

4.4. Sample Embodiment. An analog node will have inputs $x_1, \ldots, x_n$, f_i, and output x_i. In classical devices, such nodes were implemented using programmable resistors and two operational amplifiers (as in Figure 4). Assuming infinite input and zero output impedances, Kirchhoff's laws yield

$$v_- = \frac{\sum_k \frac{x_k}{R_{ki}} + \frac{y_i}{R_0}}{\sum_k \frac{1}{R_{ki}} + \frac{1}{R_0}},$$

and

$$v_+ = \frac{\sum_k \frac{x_k}{R'_{ki}}}{\sum_k \frac{1}{R'_{ki}}}.$$

For the amplifiers, we have

$$y_i = \mu(v_+ - v_-)$$

$$x_i = \mu(y_i - f_i).$$

Thus

$$y_i = \mu \left[\frac{\sum_k \frac{x_k}{R'_{ki}}}{\sum_k \frac{1}{R'_{ki}}} - \frac{\sum_k \frac{x_k}{R_{ki}} + \frac{y_i}{R_0}}{\sum_k \frac{1}{R_{ki}} + \frac{1}{R_0}} \right].$$

Then, setting

$$a = \frac{\frac{1}{R_0}}{\sum_k \frac{1}{R_{ki}} + \frac{1}{R_0}},$$

we have

$$y_i = \frac{\mu}{1 + a\mu} \sum a_{ik} x_k,$$

with

$$a_i k = \frac{\frac{1}{R'_{ki}}}{\sum_k \frac{1}{R'_{ki}}} - \frac{\frac{1}{R_{ki}}}{\sum_k \frac{1}{R_{ki}} + \frac{1}{R_0}}.$$

We see then that the hybrid realization here is similar to the one-stage solver.

5. A Two Level Algorithm. A more powerful version of the residual correction technique is the multilevel variant (see[1, 8]). In this section we consider the two level version of the latter (see[3]), and we show how to employ the hybrid, analog/digital methods to it.

The two level method for solving the system $Ax = f$ is described as follows:

Step 1. Repeat p times: $x \leftarrow x - G(Ax - b)$

Step 2. $x \leftarrow x - PB^{-1}R(Ax - b)$.

Here the first step consists of p smoothing iterations using a scaled iterative procedure G (e.g., Jacobi, symmetric Gauss-Seidel, or conjugate gradients). The matrix R interpolates from the solution space onto a coarser space and P interpolates from the coarse space into the solution space. Typically, R is a linear interpolation method and P is R^T. A customary choice for B is $B = RAP$, with the dimension of B being considerably less than that of A.

Rewrite step 2 as

$$\begin{aligned} r &= R(Ax - b) \\ y &= B^{-1}r \\ x &\leftarrow x - Py. \end{aligned}$$

Our hybrid approach produces an analog solution $\bar{y}$ of $By = r$, where

$$\bar{y} = (B - E)^{-1}r, \quad \|E\| \leq \epsilon\|B\|. \tag{5.1}$$

Note the the error e in (2.1) is incorporated into E here, and E will vary from step to step. We assume that

$$\|(I - PB^{-1}RA)(I - GA)^p\| = a < 1, \tag{5.2}$$

a condition which the original two level method requires for convergence.

Now we study the effect of the limited precision (as characterized by ϵ) of the analog implementation of (5.1) which replaces y by $\bar{y}$.

Let e_1 be the error before step 1, e_2 the error after step 1, and e_3 the error after step 2 in the two level method. Then

$$e_2 = (1 - GA)^p e_1,$$

and

$$e_3 = (I - PB^{-1}RA)e_2.$$

The corresponding error, $\bar{e}_3$, with the analog process invoked is

$$\bar{e}_3 = (I - P(B - E)^{-1}RA)e_2.$$

Assume that $\|z\| = \|Pz\|$ for all z and

$$\|(I - GA)^p\| \leq C \tag{5.3}$$

(in most cases, $C = 1$). The following theorem quantifies the degree of degradation of the estimate (5.2) in the hybrid version:

THEOREM 1. If (5.2) and (5.3) are satisfied, then

$$\|(I - P(B-E)^{-1}RA)(I - GA)^p\| \;\leq\; a + \frac{(C+a)\kappa(B)\epsilon}{1-\kappa(B)\epsilon},$$

or, equivalently,

$$\|\bar{e}_3\| \;\leq\; [a + \frac{(C+a)\kappa(B)\epsilon}{1-\kappa(B)\epsilon}]\|e_1\|.$$

Proof: First consider $\bar{y} - y$. We have

$$\begin{aligned} \bar{y} - y &= [(B-E)^{-1} - B^{-1}]r \\ &= [(B-E)^{-1}B - I]y. \end{aligned}$$

Then (cf. (2.6)),

$$\|\bar{y} - y\| \;\leq\; \frac{\kappa(B)\epsilon}{1-\kappa(B)\epsilon}\|y\|.$$

Now

$$\|y\| \;=\; \|Py\| \;=\; \|e_2 - e_3\| \;\leq\; (C+a)\|e_1\|.$$

Hence

$$\|\bar{y} - y\| \;\leq\; \frac{(C+a)\kappa(B)\epsilon}{1-\kappa(B)\epsilon}\|e_1\|.$$

Noting that

$$\bar{e}_3 \;=\; e_3 - P(\bar{y} - y),$$

we get

$$\|\bar{e}_3\| \;\leq\; \|e_3\| + \|\bar{y} - y\| \;\leq\; [a + \frac{(C+a)\kappa(B)\epsilon}{1-\kappa(B)\epsilon}]\|e_1\|.$$

□

6. Conclusions. We have analyzed a hybrid digital/analog algorithm and shown that it reduces the error per iteration by a considerable amount. In fact, most preconditioned iterative methods reduce the error by a fraction of this amount, moreover at a greater cost. The cost of the analog step has been shown to be negligible in comparison to the that of the digital step. Further, the cost of one iteration of the hybrid method is comparable to that of one iteration of a fully parallelized (digital) optimally overrelaxed Jacobi method. However, the hybrid method can be over 100 times faster than the corresponding Jacobi method for a fixed accuracy requirement. The technology exists now to build such a hybrid machine, either as a standalone computer or as a coprocessor board for a workstation.

REFERENCES

[1] R. E. Bank and C. C. Douglas, *Sharp estimates for multigrid rates of convergence with general smoothing and acceleration*, SIAM J. Numer. Anal., 22 (1985), pp. 617–633.

[2] A. Brandt, *Binary-input analog-computing (biac) hardware for numerical solution of elliptic problems*. Preliminary report, Weizmann Institute of Sciences, Rehovot, Isreal, 1969.

[3] F. Chatelin and W. L. Miranker, *Acceleration by aggregation of successive approximation methods*, J. Lin. Alg. Applic., 24 (1980), pp. 17–47.

[4] W. K. Cheng and H. J. Caulfield, *Fully-parallel relaxation algebraic operations for optical computers*, Opt. Commun., 43 (1982), pp. 251–254.

[5] L. A. Hageman and D. M. Young, *Applied Iterative Analysis*, Academic Press, New York, 1981.

[6] J. Jahns, *Concepts of optical digital computing-a survey*, Optik, 57 (1980), pp. 429–449.

[7] W. J. Karplus and R. A. Russell, *Increasing digital computer efficiency with the aid of error-correcting analog subroutines*, IEEE Trans. Computer, C-20 (1971), pp. 831–837.

[8] J. Mandel, *Multigrid convergence for nonsymmetric, indefinite variational problems and one smoothing step*, Appl. Math. Comput., 19 (1986), pp. 201–216.

[9] J. Mandel and W. L. Miranker, *New techniques for fast hybrid solutions of systems of equations*, Tech. Report RC 14157, IBM Research Division, Yorktown Heights, NY, 1988.

[10] D. Psaltis and R. A. Athale, *High accuracy computation with linear analog optical systems: a critical study*, Appl. Opt., 25 (1986), pp. 3071–3077.

[11] SPIE, *Bimodal optical computers*, 1988.

[12] ———, *Performance analysis of matrix preconditioning algorithms on parallel optical processors*, 1988.

[13] K. Watanabe and G. C. Temes, *A switched-capacitor digital multiplier*, in 1983 IEEE International Symposium on Circuits and Systems, Volume 3, New York, 1983, IEEE, pp. 1270–1273.

[14] J. Wilkinson, *The Algebraic Eigenvalue Problem*, Oxford University Press, Oxford, England, 1965.