Artificial Parameter Homotopy Methods for the DC Operating Point Problem

Robert C. Melville, Ljiljana Trajković, *Member, IEEE*, San-Chin Fang, *Member, IEEE*,
and Layne T. Watson, *Senior Member, IEEE*

Abstract—Efficient and robust computation of one or more of the operating points of a nonlinear circuit is a necessary first step in a circuit simulator. This paper discusses the application of so-called globally convergent probability-one homotopy methods to various systems of nonlinear equations that arise in circuit simulation. The so-called "coercivity conditions" required for such methods are established using concepts from circuit theory. The theoretical claims of global convergence for such methods are substantiated by experiments with a collection of examples that have proved difficult for commercial simulation packages that do not use homotopy methods. Moreover, by careful design of the homotopy equations, the performance of the homotopy methods can be made quite reasonable. An extension to the steady-state problem in the time domain is also discussed.

I. INTRODUCTION

ROBUST computation of the direct current (dc) operating point(s) of a VLSI circuit is a theoretically and practically difficult problem [1], [18], [25], [32], [34], [35], [39], [58]. For certain circuits—primarily analog designs—engineers [33], [36] spend large amounts of time "fighting" a simulator to obtain an operating point of their circuit. In some cases, the difficulty of obtaining an operating point signals that something is wrong with the circuit, and indicates that a redesign is needed. However, in other cases, the difficulty is simply an artifact of the simulator's algorithm used to find an operating point. In such cases, the designer's time spent finding an operating point has contributed nothing to the design process.

In this paper we discuss *homotopy* methods [9], [10], [13], [14], [19], [38], [49], [50], [59], [64]-[67], [69] for the operating point problem. The homotopy approach involves forming a simplified version of the circuit whose operating point is needed, finding an operating point of this easier circuit, then "sweeping" some quantity to generate a trajectory of solutions. The terminus of the sweep is the operating point of the original circuit. Once the operating point problem has been formulated as the

solution of a system of nonlinear equations, the intuitive idea of sweeping can be made mathematically precise, and numerical algorithms have been developed to follow the solution trajectories. Homotopy methods, also known as *continuation* methods, are not new in circuit simulation [9], [38], [40], [59], [64]-[67]. For example, one approach to the operating point problem is to compute a trajectory of node voltages as the power supply voltage is swept from zero to its final value. Sweeping the supply voltage is a good example of a so-called *natural parameter* homotopy, because the sweeping parameter has an obvious physical interpretation. This paper concentrates on the application of so-called "globally convergent probability-one homotopy methods" to circuit problems. As the name implies, numerical implementations of globally convergent homotopy methods exhibit an exceptionally wide domain of convergence. However, these methods are *not* applicable to arbitrary systems of nonlinear equations. Certain "coercivity" conditions are needed to argue that solutions to the homotopy equations cannot become unbounded. In Section 1.1, we draw a distinction between "qualitative" and "quantitative" failure modes of a nonlinear solver, and state in what sense our proposed methods are globally convergent. In Section II, we demonstrate that the required coercivity conditions are, in fact, quite natural for equations that arise in circuit simulation. In Section III, we describe various homotopies for solving systems of nonlinear equations, and develop a homotopy specific to bipolar designs that performs surprisingly well. In Section IV, we extent these methods to a two-point boundary value problem that formulates the *steady-state* equations for a circuit. In Section V, we describe the practical implementation of a variety of homotopies for finding the direct current bias and the steady-state solutions of electronic circuits. Finally, in Section VI, we cite some related work.

The following are the main contributions of this paper:

- Demonstrate the use of globally convergent probability-one homotopy methods (also called "artificial parameter" methods) to avoid singularities in the curve being followed by an arc length continuation procedure (Sections II and III)
- Import concepts from circuit theory to show that the zero curve of various homotopies cannot diverge to infinity (Section II)

- Give an *electronic interpretation* of a popular homotopy and explain, based on this interpretation, why it performs poorly (Section 3.1)
- Develop alternative homotopies with superior performance (Section 3.2). At least for bipolar designs, our methods come within a speed factor of two to five of faster, but less widely convergent methods
- Extend the methods developed for finding operating points to give a globally convergent procedure for a certain class of two-point boundary value problems (Section IV)
- Present convergence and timing results on a suite of difficult benchmark circuits (Section V)

1.1. Qualitative versus Quantitative Failures

The operating point problem for a dc network is formulated as a system of (nonlinear) equations to be solved:

$$F(x) = 0 \qquad (1)$$

in which x is an n-vector of unknowns. The operating point equations are typically solved by an iterative numerical method that often involves computation of the Jacobian matrix (derivative) of F. Most approaches to the operating point problem fall into one of the following categories:

- Plain Newton's method
- Norm reducing variations of Newton's method (so-called "damped" Newton) that adjust the Newton correction step to assure monotonic decrease of $\|F\|$ [4], [37]
- Continuation methods [9], [10], [13], [14], [38], [49], [50], [59], [64]–[67], [69] in which (1) is augmented with a circuit parameter μ:

$$H(x, \mu).$$

Here, μ is chosen so that the solution of $H(x,0) = 0$ is easy, and $H(x, 1) = F(x)$ identically in x. For example, the value of all independent voltage sources could be multiplied by μ. At $\mu = 0$, the circuit has an obvious operating point at which each node voltage is zero. Starting from this point, small increments in μ are taken, at each point finding an appropriate x to satisfy $H(x, \mu) = 0$, until $\mu = 1$. (A more sophisticated implementation will parameterize the solution path by *arc length*, which allows the curve to turn backward at an intermediate value of μ.)

- Transient methods [1], [58] that attempt to simulate the actual time domain response of the circuit while "powering up." In some cases, additional reactive elements (either capacitors or inductors) are added to the circuit as an aid to convergence. The time domain response of the circuit is simulated until the currents through the added capacitors, or the voltages across the added inductors, become zero.

Why might a numerical method fail to find a zero of the operating point equations? Here are some possible failure modes:

- An error tolerance cannot be met
- A matrix, such as the Jacobian matrix for F, becomes numerically singular
- Numerical noise or inexact computation of derivatives causes nonconvergence
- A user stops the computation in frustration because of excessive computer time

We classify such failures as *quantitative*. All of the preceding could, at least in principle, be avoided with more accurate arithmetic, more careful device modeling, or allocation of more computer time.

In contrast, consider the following failure modes:

- Plain Newton's method cycles, diverges to infinity, or requests the inversion of a genuinely singular matrix
- Damped Newton's method reduces the norm of F to a local minimum that is not zero, and is unable to progress beyond this point
- A continuation method encounters a singularity at a critical value of μ, and is unable to continue past this value of μ, or oscillates in some neighborhood of the critical value. Alternatively, the path being tracked by the continuation process may diverge to infinity as μ approaches one
- A transient method oscillates so that the currents through the added capacitors, or the voltages across the added inductors, never converge to zero
- One of the first three preceding methods finds a physically unstable operating point. Such a solution is a perfectly good zero of the operating point equations, but would never be observed in practice

We argue that this class of failures is rather different than the first class. More computer time or better arithmetic would not help! The same *qualitative behavior* would be observed (perhaps with more accuracy).

This distinction is not meant to imply that dealing with quantitative failures is "easy"—indeed, in a large practical simulator, quantitative issues may be the dominant concern. However, the distinction forms a basis for discussion.

In the sequel, we claim a procedure for finding a dc operating point of a circuit that avoids the above-mentioned qualitative failures (with probability one). Simplicial methods [64]–[67] can make similar claims but are difficult to apply to circuits with several hundred transistors, especially when using complex and highly accurate mathematical transistor models. The assumptions that our techniques place on the transistor models (smoothness and passivity) are natural from a standpoint of transistor physics.

II. Coercivity Conditions for Circuit Equations

A number of elegant mathematical results concern solutions to systems of equations that satisfy certain "boundedness" conditions [42], [51]. Perhaps the best example is the Brouwer fixed point theorem [8], which states that a continuous map f from a convex compact set

into itself must have a fixed point; i.e., for some x^* in the set, $f(x^*) = x^*$. Why might one think that the Brouwer fixed point theorem would be applicable to the dc operating point problem? The intuition that such is the case is based on the following fact about nonlinear resistive circuits (at least, those that arise in practical integrated circuit designs). At the dc operating point of such a circuit, each node voltage is bounded in absolute value by the sum of the absolute values of the voltage sources in the circuit [61]. A circuit with this property is called *no-gain*. In other words, if the circuit has n nodes, and we imagine the sum of the absolute values of the voltage sources normalized to the range [0, 1], then the operating point is an element of the unit n-cube. This fact is no surprise to designers of electronic circuits, although a rigorous proof of this assertion is not trivial [61].

The no-gain property is intrinsic to real transistors, but may or may not be preserved in a circuit simulator, depending on transistor models. As an example, consider the circuit of Fig. 1, which shows a typical bias configuration for a single transistor amplifier stage.

Fig. 2(a) shows the circuit of Fig. 1 redrawn with an overly simple model of a transistor that does not capture the *saturation behavior* of a real transistor. Correct bias voltage results from the use of a more accurate model shown in Fig. 2(b).

This kind of boundedness property extends to the time domain behavior of electronic circuits. Consider the output waveform of Fig. 3, as the amplifier goes into saturation (Vin = 200 mV). Although the output *clips*, its peak-to-peak amplitude remains bounded by the supply voltage.

Despite the historical appeal of the Brouwer fixed point theorem, and the body of knowledge about the no-gain property, we have found the following theorem [29] easier to apply to circuit equations.

Theorem 1: Let F be a continuous mapping from $\mathbf{R}^n$ into $\mathbf{R}^n$. Suppose that $x^T F(x) \geq 0$ for all $x \in \mathbf{R}^n$ such that $\|x\| = r$, where r is a fixed positive real number. Then F has a zero x^* such that $\|x^*\| \leq r$.

A proof of this theorem for the one-dimensional case is obvious, since it simply states that F is zero at an endpoint of the interval $[-r, +r]$ or has a sign change on the interval. A proof in higher dimensions may be obtained by reduction to the Brouwer fixed point theorem [29].

The nodal formulation of circuit equations [27] specifies a sum of currents for each node. The result is a system

$$F_1(x_1, \cdots, x_n) = 0$$

$$F_2(x_1, \cdots, x_n) = 0$$

$$\vdots$$

$$F_n(x_1, \cdots, x_n) = 0$$

where the dimension of x_i is voltage and the dimension of F_i is current. Thus, the dimension of the inner product $x^T F(x)$ is *power*.
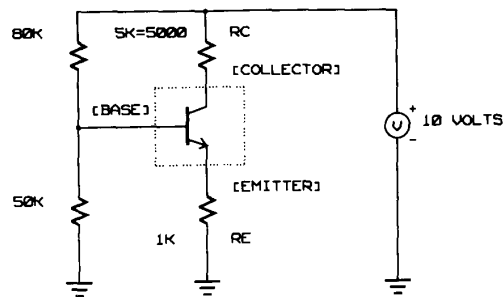

Fig. 1. One transistor bias circuit.
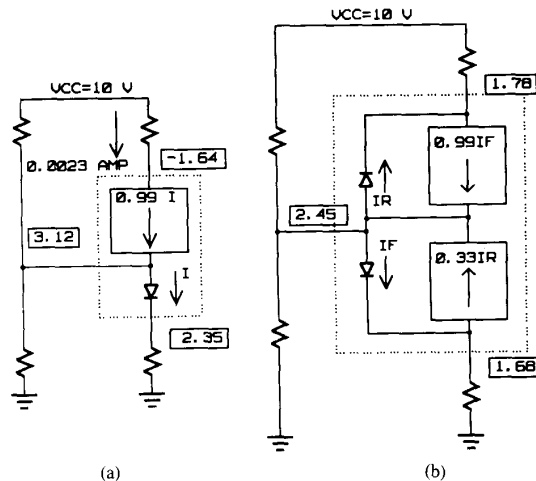

(a)   (b)

Fig. 2. Two models of a transistor.
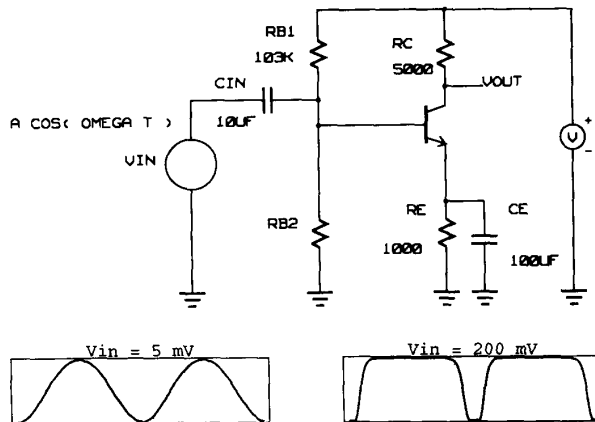
Fig. 3. Amplifier circuit.

A circuit element is *passive* if it does not generate power [24], [60], [63]. This can be stated in mathematical terms by considering the voltage v_k across each element and the current i_k flowing into the element. If the sum $\Sigma i_k v_k$ over all elements is always nonnegative, then the device is *passive*. Passivity is a less restrictive condition than the no-gain property introduced earlier.

Among the electronic devices introduced so far, linear

(positive) resistors, diodes, and transistors are passive. The current amplifier is not passive; however, the particular arrangement of diodes and current amplifiers in the Ebers–Moll transistor model is passive [21]. Any interconnection of passive components is passive. This lets us evaluate the inner product condition for the nodal equations of a nonlinear resistive circuit. The particular values of circuit parameters establish a radius of a ball in $\mathbf{R}^n$ such that for any vector of node voltages x on this ball, the inner product $x^T F(x)$ is a sum of powers that is nonnegative. If we appeal to passivity rather than the no-gain property, then, in general, the radius of this ball will be larger than the sum of the absolute values of the independent voltage sources, and might be difficult to calculate. However, applications of Theorem 1 do not require the *knowledge* of the radius, only its existence. A detailed exposition of this argument appears in [45].

Therefore, a passivity argument can be made for an electronic circuit consisting of independent voltage sources, resistors, diodes, and transistors. This covers all practical cases. Occasionally, designers use voltage amplifiers to model operational amplifiers. A voltage amplifier delivers an output voltage μv_{in}, where μ is a constant and v_{in} is a voltage drop across some branch in the circuit. The graph of the input/output relationship of such a device would be a straight line of slope μ extending to infinity in either direction. Suppose μ is large, say 1000. Then an input voltage of 1 V generates an output voltage of 1000 V. Again, common sense about electronic circuits says that an operational amplifier built using transistors and operated from $+12$ V and -12 V supplies cannot generate an output voltage of 1000 V. Any practical operational amplifier exhibits *limiting* behavior at its output. That is, the output is indeed equal to μv_{in} over some range of v_{in}, but beyond that range, the output voltage is bounded by the positive and negative power supply values. When a voltage amplifier is modified to model this limiting behavior (providing a more accurate model of an operational amplifier), the inner product condition of Theorem 1 can be satisfied at a certain radius r that may depend on the limits set for the voltage amplifiers.

So far, our discussion of passivity has been limited to the case of nonlinear resistive circuits, i.e., circuits with no notion of time. In a later section, we discuss passivity for the time domain response of a circuit.

III. Homotopies for the DC Operating Point Problem

The general homotopy paradigm involves embedding the equations to be solved, $F(x) = 0$, in a system of equations of one higher dimension, $H(x, \lambda) = 0$, with the introduction of one more variable λ, called the *homotopy parameter* or *continuation parameter*. Typically, λ is restricted to the range [0, 1], and the embedding is done so that the augmented system $H(x, 0) = 0$ is easy to solve and reduces to the original system when $\lambda = 1$, i.e., $H(x, 1) = F(x)$.

Starting from the solution to $H(x, 0) = 0$, one follows a connected set of points (x, λ) such that $H(x, \lambda) = 0$ until $\lambda = 1$. In traditional continuation methods [31], sometimes called *monotonic continuation*, a functional relationship is assumed between x and λ, so that there is a unique value $x(\lambda)$ such that $H(x(\lambda), \lambda) = 0$ for $\lambda \in$ [0, 1]. An important extension of the continuation paradigm is the *arc length continuation* [28], in which both x and λ are thought of as functions of arc length s along a connected set of points such that $H(x(s), \lambda(s)) = 0$. Now, s is advanced to a value s^* such that $\lambda(s^*) = 1$.

Continuation methods, either monotonic or arc length, are well known in circuit simulation, either for finding an operating point [9], [10], [13], [14], [38], [49], [50], [59], [64]–[67], [69], or for calculating a dc transfer curve [16]. For example, one approach to the operating point problem is to multiply each voltage source in the circuit by λ. When $\lambda = 0$, the circuit has an obvious solution in which each node voltage is 0 V. A continuation process is used to advance λ to 1, where a point x^* such that $H(x^*, 1) = 0$ represents a solution to the circuit with all voltage sources at their desired values.

Continuation of the supply voltage is a good example of *natural parameter* continuation, as discussed in [52], because the continuation parameter has an obvious physical interpretation. In this section we describe some theoretical work of Chow *et al.* [12], who have given a class of methods that can be proven to have very desirable numerical properties. The methods presented in [12] are called *artificial parameter homotopy* (or *globally convergent probability-one homotopy*) methods to distinguish them from traditional continuation methods, either monotonic or arc length. This distinction is substantial, as discussed at length in [52]. Our work concentrates on the artificial parameter approach.

Consider the scheme for continuation in the values of the voltage sources, using continuation in the arc length parameter s. This process is described by a system of equations

$$H_1(x_1(s), \cdots, x_n(s), \lambda(s)) = 0$$

$$H_2(x_1(s), \cdots, x_n(s), \lambda(s)) = 0$$

$$\vdots$$

$$H_n(x_1(s), \cdots, x_n(s), \lambda(s)) = 0$$

where x_1 through x_n represent the node voltages. The parameter $\lambda \in$ [0, 1] multiplies the value of each independent voltage source. The *zero set* of H is

$$H^{-1}(0) = \{(x(s), \lambda(s)) | \lambda \in [0, 1]$$

$$\text{and} \quad H(x(s), \lambda(s)) = 0\}$$

which is typically a union of curves in $\mathbf{R}^{n+1}$. Given a starting point in the zero set (certainly available at $\lambda = 0$) one can think of moving along a component of the zero set as tracing the trajectory of an initial value problem.

Numerical procedures for carrying out this type of *curve tracking* will be described later.

Unfortunately, this continuation scheme is not always robust. For certain circuits, at a critical point on the zero set, the solution trajectory will attempt to split (bifurcate). If the curve-following procedure is unlucky enough to hit the bifurcation point, or a sufficiently small neighborhood of it, the Jacobian matrix of H will become rank deficient, requiring further analysis. Note that a circuit may have a unique operating point at the *final* value of all the supplies, yet still have multiple operating points at their intermediate values. In such a case, the success or failure of source stepping will be sensitive to the relative rate at which the various sources are turned on. Such behavior is reported in [68].

Consider the flip-flop of Fig. 4. This is a circuit with three dc operating points. However, if the value of the voltage source is set to zero, then the circuit has a unique solution, with all voltages equal to zero. Let $x = (xc_1, xb_1, xc_2, xb_2)^T$ be the vector of node voltages. The following equations represent the voltage source continuation embedding for this circuit:

$$H_1(x, \lambda) = (xc_1 - \lambda VCC)/RC1 + (xc_1 - xb_2)/$$

$$RB1 + ic_1$$

$$H_2(x, \lambda) = (xb_1 - xc_2)/RB2 + ib_1$$

$$H_3(x, \lambda) = (xc_2 - \lambda VCC)/RC2 + (xc_2 - xb_1)/$$

$$RB2 + ic_2$$

$$H_4(x, \lambda) = (xb_2 - xc_1)/RB1 + ib_2$$

where H is a mapping from $\mathbf{R}^{n+1}$ into $\mathbf{R}^n$, with $n = 4$. The Jacobian matrix for H has dimension $n \times (n + 1)$ and can be evaluated for any particular $(x, \lambda) \in \mathbf{R}^{n+1}$. Does this matrix have full rank (i.e., rank 4) for *all* points in the zero set of H? The answer is no; the symmetry in the circuit generates a bifurcation at a critical point in the zero set with $\lambda \approx 0.1135$. Then $\lambda VCC \approx 0.68$ V, just enough to "turn on" the base-emitter junctions of the two transistors. The rank of the Jacobian matrix of H was monitored by performing a singular value decomposition of the Jacobian matrix [43]. At the critical value for λ mentioned earlier, the smallest singular value drops to approximately 2×10^{-14}, indicating a rank of deficiency.

A rather different approach to the formulation of embeddings is taken in [12], resulting in artificial parameter homotopies. Under certain conditions, the zero set of such homotopies can be shown to contain a smooth curve that leads to a solution with bounded arc length. Bifurcations or other types of singularities do not occur. In particular, Chow et al. [12] consider solving a system of equations $F(x) = 0$, where $F: \mathbf{R}^n \to \mathbf{R}^n$ satisfies the inner product condition of Theorem 1. They impose the additional condition that F be C^2, and consider the following homotopy for finding a zero of F:

$$\rho_a(x, \lambda) = \lambda F(x) + (1 - \lambda)(x - a) \qquad (2)$$

where $a = (a_1, \cdots, a_n)$ is a fixed element of $\mathbf{R}^n$. Note that a solution to $\rho_a(x, 0) = 0$ is trivial, and a solution to $\rho_a(x, 1) = 0$ is a zero of F. Following the homotopy paradigm, the easy problem, "solve $(x - a) = 0$," is deformed in a continuous manner into the desired problem, "solve $F(x) = 0$." The zero set for (2) is the set $\Gamma_a = \{(x, \lambda) | \lambda \in [0, 1) \text{ and } \rho_a(x, \lambda) = 0\}$.

The theory developed in [12] shows that with the additional assumption that F is C^2, then for *almost all choices* of a with $\|a\| \leq r$, Γ_a contains a smooth path emanating from $(a, 0)$ and terminating at a point $(x^*, 1)$ where $F(x^*) = 0$. The term "almost all" means that the set of starting points that does not generate successful trajectories has a measure of zero in $\mathbf{R}^n$. This property of the artificial parameter methods is a consequence of a version of Sard's theorem [41]. Thus, the a vector provides an element of random choice in the procedure for finding a zero of F. Each different choice for a provides a different possible path to a zero of F; "most" such paths are successful. The algorithm, then, for finding a zero of F, is to pick such an a at random (which fixes a value of $x_0 = x(0)$), then solve an initial value ordinary differential equation (ODE) problem to generate the solution trajectory emanating from $(x_0, 0) = (a, 0)$.

Watson et al. [52]–[54] have done careful numerical implementation of an algorithm that can follow the solution trajectory of (2). Code is available in a package called HOMPACK. The implementation in HOMPACK requires the user to supply a subroutine to evaluate F and its Jacobian matrix. Given a starting point a, HOMPACK tracks the solution trajectory of (2) for $\lambda \in [0, 1]$. Because the set of bad starting points has a measure of zero, a random choice of a has zero probability of hitting one. In this sense, the algorithm for finding a zero of F is globally convergent.

Of course, one must be careful when applying concepts from real analysis to the computer number system that has only a finite approximation to the real numbers. Nonetheless, our experiments with probability-one homotopy algorithms, as implemented in HOMPACK, indicate exceptionally robust convergence, provided that the device models are smooth enough.

3.1. Formulation of the Operating Point Problem

The *nodal formulation* presents the operating point problem as a system of nonlinear equations $F(x) = 0$, where x is a vector of node voltages. Let there be n nodes, and let x_k denote the voltage at node k. $F_k(x_1, \cdots, x_n)$ will be an equation for the sum of currents flowing into node k through all the branches connected to that node. The system F has been shown to satisfy the inner product condition if all the devices (other than the voltage sources) are passive [45]. Thus, the homotopy (2) can be used to find an operating point for the circuit.

In fact, $\rho_a(x, \lambda) = 0$ has an interesting physical interpretation. For an arbitrary node k, imagine connecting a resistor in series with a voltage source to the ground. Let the value of the voltage source be a_k, and the value of the
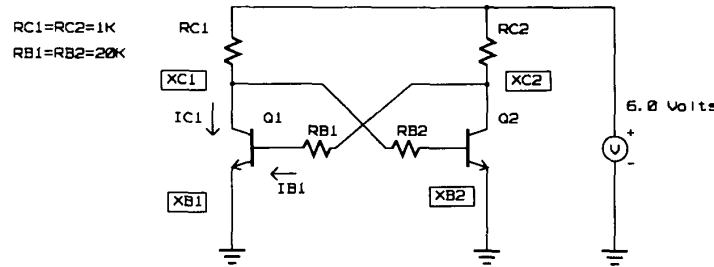
Fig. 4. Flip-flop.

resistor be $\lambda/(1 - \lambda)$. The nodal equation for node k is

$$\sum_{j=1}^{m} I_j = \frac{1 - \lambda}{\lambda} (a_k - x_k) \qquad (3)$$

where there are m branches connected to node k in the original circuit, carrying currents $I_1, \cdots, I_m$. Recognizing that the current summation on the left-hand side of (3) is $F_k(x)$ as previously defined, we see that (3) is equivalent to (2).

Note that the value of the resistor is 0 for $\lambda = 0$. Hence, the node voltage is forced to the value a_k at $\lambda = 0$. Then, as λ approaches 1, the value of the resistor approaches infinity, and the inserted source is decoupled from the circuit.

The physical interpretation of the homotopy also provides intuition about why the homotopy should be bifurcation free for almost all choices of an a vector. Consider the flip-flop of Fig. 4. The dc equations for this circuit have three solutions, corresponding to transistor **Q1** on and transistor **Q2** off, **Q2** on and **Q1** off, and both transistors on. As discussed earlier, the zero curve of the supply continuation equations for this circuit must have a bifurcation. Conversely, imagine connecting a resistor and voltage source combination to the two collector nodes. Let the voltage sources have values a_1 and a_2. Now, sweep the value of the resistors from zero to infinity. The only choice of (a_1, a_2) that will hit the bifurcation is $a_1 = a_2$; any other pair of starting values will steer the circuit to one mode or the other. Given a random choice of (a_1, a_2), the probability of hitting $a_1 = a_2$ is zero. In practice, such pathological cases can be circumvented easily.

Consider again the natural parameter homotopy proposed for the flip-flop in Section III. HOMPACK *was* able to follow the zero curve of this homotopy to a solution at $\lambda = 1$, although there is a bifurcation in the zero curve—the curve tracking procedure was lucky enough to "hop over" the bifurcation point on the zero curve. However, what operating point was found by following the zero curve of the natural parameter homotopy? Answer: The *unstable* operating point at which both transistors conduct equal amounts of current. Although an exact zero of the operating point equations, this operating point is unstable in the sense that the slightest perturbation will cause the circuit to move to one of the two stable states in which one transistor is conducting and the other is not. The un-

stable operating point is of little or no interest to a designer. Conversely, the introduction of the random vector moves the zero curve of the artificial parameter homotopy toward one of the two more interesting stable solutions.

3.2. Alternative Homotopies

Experiments with the homotopy (2) show that it converges robustly, but rather slowly. In some of our examples, more than a thousand steps (thousand evaluations of the Jacobian matrix) were needed to get to the solution. The physical interpretation provides some insight into the problem. Suppose node k has a small signal impedance of 1 MΩ at the operating point of the circuit. The homotopy (2) implies connecting a resistor and a voltage source between this node and the ground. Before the homotopy resistor is "removed from the circuit" it should have a value at least ten times as large as the natural impedance of the node to which it is connected. This means close to 10 MΩ for node k. However $\lambda/(1 - \lambda)$ is equal to 10^7 only for λ *very* close to 1. We experimented briefly with a scheme that attempted to scale the homotopy resistors on a per node basis, but this became cumbersome and required *second* derivatives of device models. Instead, we propose a different approach using the general *curve tracking* option in HOMPACK. This allows a user to construct a customized homotopy. The presentation of [12] and [52] treats a homotopy of the following form:

$$\rho: \mathbf{R}^m \times \mathbf{R}^n \times [0, 1) \to \mathbf{R}^n$$

denoted $\rho(a, x, \lambda)$ where a is an m-vector, x is an n-vector, and $\lambda \in [0, 1)$. The m-vector a is a random parameter. For fixed a, define

$$\rho_a(x, \lambda) = \rho(a, x, \lambda).$$

Now consider the zero set

$$\Gamma_a = \{(x, \lambda) | \rho_a(x, \lambda) = 0\}.$$

If certain conditions on ρ are met, then Γ_a is a union of smooth curves, one component of which connects a zero of $\rho_a(x, 0)$ with a zero of $\rho_a(x, 1)$.

In Section II we gave the homotopy first, then derived its circuit interpretation. In this section we work backward, giving a circuit interpretation first, then deriving from it the homotopy equations. The idea is to start with all nonlinear devices removed from the circuit, then gradually bring them in. A solution trajectory is generated as

the nonlinear devices sink more and more current. Fig. 5 shows a Wilson current source [62].

The operating point equations for this circuit are shown subsequently. The notation is a fragment of an executable code, and requires a bit of explanation. The transistor model Q1 is stimulated by calling Q1.stim with the node voltages for the collector, base, and the emitter, respectively. (Note that the ground node is at 0 V by definition.) Then, a function call such as Q1.ic( ) will deliver the collector current for the given set of stimuli:

```
Q1.stim(x[0],x[1],0.0)
Q2.stim(x[2],x[0],x[1])
Q3.stim(x[1],x[1],0.0)
F[0] = Q2.ib( )+Q1.ic( )+((x[0]
       -(VCC))/RP)
F[1] = Q3.ic( )+Q3.ib( )
       +Q2.ie( )+Q1.ib( )
F[2] = Q2.ic( )+((x[2]
       -(VCC))/RL)
```

To transform these equations into a homotopy, imagine stimulating the nonlinear devices in the following fashion:

```
Q1.stim(lambda*x[0],lambda*
   x[1],0.0)
Q2.stim(lambda*x[2],lambda*
   x[0],lambda*x[1])
Q3.stim(lambda*x[1],lambda*
   x[1],0.0)
```

where lambda ∈ [0,1]. When lambda = 0, all nonlinear devices' terminal currents are zero. Therefore, they are "removed from the circuit." This provides a starting point on a component of Γ_a, as defined earlier, that is tracked until $\lambda = 1$.

There is a problem with this construction for $\lambda = 0$. When $\lambda = 0$, no current flows into transistors, so they might as well be disconnected. This leaves node x_1 floating. In mathematical terms, the Jacobian matrix for the equations will be singular at $\lambda = 0$, so HOMPACK is not even able to start tracking the solution trajectory. We circumvented this problem by attaching extra circuitry to each node. At node k, we attached a conductance of value GLEAK in a series with a voltage source of value a_k. The voltage sources $(a_1, \cdots , a_n)$ provided the random element needed in the probability-one homotopy construction. Finally, $(1 - \lambda)$ GLEAK approaches zero as λ approaches one. The resulting equations for the Wilson current source are as follows:

Fig. 5. Wilson current source.

These equations define a homotopy $\rho(a, x, \lambda)$. Several stipulations are placed on ρ in [52] so that HOMPACK will be able to track the solution trajectory. We address each of these stipulations separately:

- ρ *is* C^2: As in homotopy (2), this depends entirely on the device models. We assume that their characteristics are sufficiently smooth.
- The Jacobian matrix $D\rho(a, x, \lambda)$ has rank n on the set

$$\rho^{-1}(0) = \{(a, x, \lambda)|a \in \mathbf{R}^n, 0 \le \lambda < 1,$$

$$\mathbf{x} \in \mathbf{R}^n, \rho(a, x, \lambda) = 0\}:$$

This property is assured by the terms GLEAK*(x[k]-a[k]). The Jacobian matrix is an $n \times (2n + 1)$ matrix. It can be written as the concatenation of three submatrices:

$$[\partial\rho/\partial a \quad \partial\rho/\partial x \quad \partial\rho/\partial\lambda].$$

The first $n \times n$ submatrix is a diagonal matrix with $-(1 - \lambda)$ GLEAK in each diagonal position. For $\lambda < 1$, this matrix has rank n.

- *For fixed* $a \in \mathbf{R}^n$, $\rho_a(x, \lambda) = \rho(a, x, \lambda)$ *has a unique solution* x_0 *at* $\lambda = 0$: When $\lambda = 0$, the equations for the Wilson current source describe a circuit consisting of resistors and voltage sources only, hence we have a unique solution. Setting all the currents through nonlinear devices to zero and rearranging yield the following equations for node voltages at

```
double lambda_bar = 1.0 - lambda
Q1.stim(lambda*x[0],lambda*x[1],0.0)
Q2.stim(lambda*x[2],lambda*x[0],lambda*x[1])
Q3.stim(x[1],x[1],0.0)
rho[0] = lambda_bar*GLEAK*(x[0]-a[0])+Q2.ib( )+Q1.ic( )+((x[0]-VCC)/
         RP)
rho[1] = lambda_bar*GLEAK*(x[1]-
         a[1])+Q3.ic( )+Q3.ib( )+Q2.ie( )+Q1.ib( )
rho[2] = lambda_bar*GLEAK*(x[2]-a[2])+Q2.ic( )+((x[2]-(VCC))/RL)
```
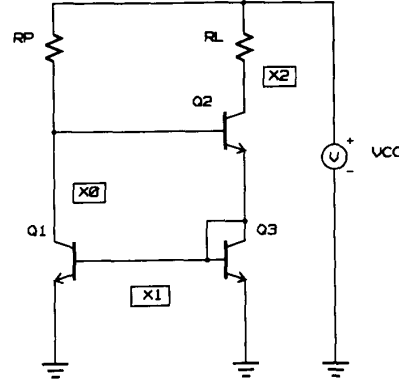
$\lambda = 0$:

```
x[0] = (GLEAK*a[0]+VCC/RP)/
       (GLEAK+1.0/RP)
x[1] = a[1]
x[2] = (GLEAK*a[2]+VCC/RL)/
       (GLEAK+1.0/RL)
```

- $\rho_a(x, 1) = F(x)$: This holds because the "leakage" circuitry is removed and each nonlinear device is stimulated by the nominal voltage.
- *The zero set Γ_a is bounded:* In other words, the solution trajectory stays within some fixed distance from the origin in $\mathbf{R}^n$ for all $\lambda \in [0,1)$. The equations represent operating point equations for a *family* of circuits, one for each value of $\lambda \in [0,1)$. For $\lambda \in (0,1)$ the circuit will comprise some mixture of partially stimulated nonlinear devices and the leakage circuitry from each node to ground. However, a partially stimulated nonlinear devices is still a no-gain element. This means that the entire circuit is a no-gain circuit for intermediate values of λ. Hence, the node voltages are bounded in absolute value by the sum of the absolute values of all the voltage sources in the circuit. Note that the values of the voltage sources do *not* change as λ sweeps from 0 to 1, so that a single fixed bound may be placed on the value of any node voltage in the solution trajectory.

We also make the reasonable engineering assumption that the Jacobian matrix of ρ_a has full rank at an operating point of the circuit, i.e., at a point $(x^*, 1)$ on Γ_a. This simply says that the circuit has one or more well-defined solutions that are isolated points in $\mathbf{R}^n$.

The example of the Wilson current source indicates all the essential features of a homotopy called the *variable stimulus* homotopy. A nodal equation is written for each node that is not connected directly to a voltage source. In addition, for each such node, a series combination of a conductance and a voltage source is connected to the ground. Let the voltage source have fixed value $a[k]$ for the kth node, and let the conductance have the value $(1 - \lambda)\texttt{GLEAK}$, where $\texttt{GLEAK}$ is a fixed parameter. Then, for each nonlinear device, we constructed a variable stimulus model with an additional parameter $\lambda \in [0, 1]$. Each terminal voltage is multiplied by this parameter. Clearly, when $\lambda = 1$, this circuit is equivalent to the circuit whose operating point is desired. When $\lambda = 0$, the circuit is linear, with a unique value for each node voltage. Experiments with this homotopy indicate that the solution trajectories are much smoother than for (2). Moreover, the action is spread out evenly over all values of λ, rather than concentrated near the end of the path. Therefore, HOMPACK is able to take larger steps along the path and use fewer iterations.

A variation on the homotopy theme is particularly suitable to bipolar circuits which often present the greatest challenge to operating point algorithms. Fig. 2(b) shows the conventional Ebers–Moll model for a bipolar transis-

tor, with forward current gain α_F and reverse current gain α_R. Suppose we set the gains of each transistor to $\alpha_F = \alpha_R = 0$. Each transistor transforms into a three-terminal element in which a diode is connected between the base and the collector, and the base and the emitter. This observation suggests a homotopy in which the current gain of each transistor is set to zero at $\lambda = 0$, then gradually increased toward the appropriate value for each transistor. "Floating" nodes present no problem, as with the variable stimulus homotopy; however, we still include the "leakage" circuitry to provide the random element needed to avoid bifurcations. We call this homotopy the *variable gain* homotopy.

In summary, the variable gain homotopy for bipolar circuits is a two-stage procedure. Set the forward and reverse current gains of each transistor to zero. Attach a series combination of a conductance and a voltage source from each node to ground. Choose a random value for each of these voltage sources. The only nonlinear components in the resulting circuit are diodes. This circuit has a unique operating point [11], [47]. For phase 1, use the variable stimulus homotopy to find the operating point of this circuit. Call this operating point x^{**}. For phase 2, define the homotopy $\rho_a(x, \lambda)$ to represent the circuit with all transistor current gains multiplied by λ and all leakage conductances multiplied by $(1 - \lambda)$. The point $(x^{**}, 0)$ is clearly on a component of the zero set of the homotopy for phase 2. Now track this component as the gains of each transistor change from zero to the appropriate final value. Simultaneously, let the leakage conductance become zero. The endpoint of the second phase is a point $(x^*, 1)$, such that x^* is an operating point.

As with the variable stimulus homotopy, we must address five criteria required by the probability-one theory:

- ρ *is* C^2: As in the previous cases, this depends entirely on the device models.
- *The Jacobian matrix $D\rho(a, x, \lambda)$ has rank n along the zero curve of ρ:* As in the variable stimulus homotopy, the leakage circuitry assures that the Jacobian matrix has rank n. During the second phase, in which the gain elements are brought in, the leakage conductances are allowed to go to zero, but still remain nonzero for $\lambda < 1$. Again, $D\rho$ contains a submatrix of rank n.
- *For fixed $a \in \mathbf{R}^n$, $\rho_a(a, x, \lambda) = \rho(a, x, \lambda)$ has a unique solution x_0 at $\lambda = 0$:* The starting point for the second phase is the solution x^{**} to the diode-only circuit obtained at the end of the first phase. Is this solution unique? Duffin [17] proved that a circuit consisting of voltage sources, resistors, and diodes has a unique operating point.
- $\rho_a(x, 1) = F(x)$: This is clear from the form of the equations, because the "leakage" circuitry is removed, and each transistor gain is at its appropriate final value.
- *The zero set of ρ_a is bounded:* The theory presented in [61] shows that a bipolar transistor remains no-

gain as long as the absolute values of the current gains remain less than or equal to one. This means that intermediate circuits for $\lambda < 1$ are no-gain, hence all node voltages are limited to a fixed value set by the independent sources in the circuit.

Experimental results with the variable gain homotopy are presented in the last section. For the case of bipolar circuits, a hybrid approach provides a fast solution method. Our experiments with the variable gain homotopy indicate that most of the work is spent in phase 1— i.e., in finding the solution to the diode-only circuit. Duffin [17] proved that a circuit consisting of diodes as the only nonlinear components has a unique solution. (We assume the diode model includes the reverse breakdown. All practical diodes display such behavior.) Duffin's result suggests that norm reducing variations of Newton's method ("damped" Newton), such as proposed by Bank and Rose [4], will converge robustly for such circuits. The hybrid technique, then, is to use the damped Newton code to find the operating point of the circuit with all the current gains set to zero, then employ homotopy to track the solution as the gains are restored to their desired full values. This indeed works quite well, and is two to three times faster than using homotopy alone. Performance statistics are presented in the last section.

IV. STEADY-STATE COMPUTATION IN THE TIME DOMAIN

Time domain simulations were described briefly in Section I. A system of first-order nonlinear differential equations is formulated for a circuit that contains capacitors and inductors, then standard numerical methods are employed to integrate the equations from the initial state computed using a dc operating point algorithm.

For an important class of circuits, the solution to the system of differential equations approaches a periodic steady state asymptotically over time as initial transients die out. Unfortunately, the amount of integration steps needed to get to the steady state can be prohibitively large. Various techniques that result in savings in computer time and an improvement in the accuracy of the solution have been proposed [3], [30], [54]–[57] to calculate the steady-state response directly. One such approach is to replace the nonlinear differential equations with nonlinear difference equations, employing a numerical approximation for the derivatives [30]. Each solution waveform, assumed to be periodic, is represented as a vector of sampled values on a fixed partition of time points over one period. Thus, if there are n equations in the system of nonlinear differential equations, and m samples per period, the discretization process results in nm equations in nm unknowns. The resulting system can be large, but is very sparse.

Getting a solution to this system of nonlinear algebraic equations presents the same difficulties as finding dc operating points of a circuit. Can the artificial parameter homotopy machinery that we have developed for the operating point problem be applied to the systems of equations that arise from the steady-state problem? In this section, we answer this question affirmatively. The key idea is to prove the inner product condition for the equations that result from discretizing the differential equations.

Consider the simple one-transistor amplifier circuit of Fig. 3. If the time constant established by C i n is large compared to the period of the stimulus, finding the time domain response of this circuit by treating it as an initial value problem is expensive. Suppose, instead, we treat this as a two-point boundary value problem. The following system of nonlinear differential equations describes the time domain behavior of the circuit. The unknowns are periodic functions of time, e.g., $x_C(t)$ is the waveform of the collector voltage of Q1. The notation $i_C(x_C, x_B, x_E)$ denotes the instantaneous current flowing into the collector of Q1, given instantaneous voltages (x_C, x_B, x_E) of the collector, base, and emitter, respectively.

$$(x_B - VCC)/RB1 + x_B/RB2 + i_B(x_C, x_B, x_E)$$
$$+ C i n \frac{d}{dt}(x_B - \cos(\omega t)) = 0$$
$$x_E/RE + CE\frac{d}{dt}x_E + i_E(x_C, x_B, x_E) = 0$$
$$(x_C - VCC)/RC + i_C(x_C, x_B, x_E) = 0.$$

The following additional conditions enforce periodicity, where T is the (known) period of the stimulus:

$$x_C(0) = x_C(T)$$
$$x_B(0) = x_B(T)$$
$$x_E(0) = x_E(T).$$

Given some convenient number of sample points, say m, the preceding system can be discretized on $[0,T)$ using the mesh points $0 = t_0 < t_1 < t_2 < \cdots < t_{m-1} < t_m = T$. The derivative operator dx/dt is replaced by a forward difference approximation $(x_{k+1} - x_k)/(t_{k+1} - t_k)$, where the subscripts are taken modulo m, and $x_k \approx x(t_k)$.

The result is a system of $3m$ equations in $3m$ unknowns. The equations are still nonlinear because of terms such as $i_E(x_C(t_k), x_B(t_k), x_E(t_k))$. We argue that this system of nonlinear algebraic equations is simply the dc operating point equations for a passive circuit. Therefore, all the global convergence arguments we have developed earlier apply. The original circuit can be copied m times, and the copies labeled as t_0, t_1, etc. Consider a capacitor between nodes P and Q in the original circuit. This capacitor can be replaced by a resistor and a voltage-controlled current source in the t_{k-1} and t_k copies, as shown in Fig. 6.

In the original circuit, the nodal equation for node P is

$$A + Cd(v_P(t) - v_Q(t))/dt$$

to account for the current flowing into the capacitor. The term A depends on other connections to node P. Writing dt as an abbreviation for $t_k - t_{k-1}$, the nodal equation for
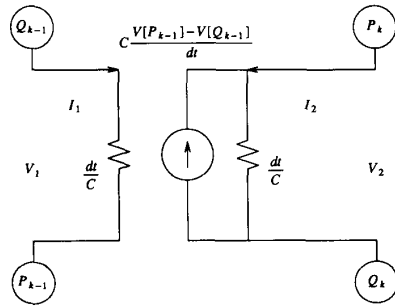
Fig. 6. Coupled circuits.

node P in the t_k circuit is

$$A_k + \frac{C}{dt} (v_P(t_k) - v_Q(t_k)) - \frac{C}{dt} (v_P(t_{k-1}) - v_Q(t_{k-1}))$$

where the voltage-controlled current source introduces coupling from the previous time point t_{k-1}. This equation is the discrete approximation to the differential equation given earlier. The equation for node Q is similar with a sign change for the derivative term. Finally, there is a connection from copy t_0 to copy t_{n-1} to enforce periodicity.

The "unrolled" circuit of Fig. 6 is a nonlinear resistive circuit, but does it satisfy the inner product condition of Theorem 1 on some r ball? A voltage-controlled current source is a nonpassive element; therefore, we have to prove that the passivity condition is satisfied.

The circuitry in the center of Fig. 6 may be treated as a component with four terminals. In circuit terminology, it is a *linear two-port*. We argue that it is a *passive* linear two-port. This implies that the entire circuit is passive, since the other components are resistors, diodes, and transistors. Let a current I_1 flow in and out of the terminals on the left side of the component, and let there be a voltage drop of V_1 across these two terminals. Define I_2 and V_2 similarly for the right-hand side. The total power dissipated by the component is $I_1 V_1 + I_2 V_2$. We show that this power is always nonnegative. Using Ohm's law, and the definition of the voltage-controlled current source, we may write

$$I_1 V_1 = \frac{C}{dt} V_1^2$$

$$I_2 V_2 = \frac{C}{dt} V_2^2 - \frac{C}{dt} V_1 V_2$$

which yields

$$I_1 V_1 + I_2 V_2 = \frac{C}{dt} V_1^2 + \frac{C}{dt} V_2^2 - \frac{C}{dt} V_1 V_2$$

$$= \frac{C}{2dt} (V_1^2 + V_2^2) + \frac{C}{2dt} (V_1 - V_2)^2.$$

The last quantity is clearly nonnegative for all v_1 and v_2. Thus, the discretization of the original system of differ-

ential equations produces a system of nonlinear algebraic equations that satisfies the inner product condition of Theorem 1.

In this section we drew a connection between the dc operating point problem and the steady-state time domain problem. In fact, passivity considerations apply to the latter, allowing us to apply the theory of globally convergent homotopy methods to the steady-state analysis.

V. EXAMPLES AND PRACTICAL RESULTS

In this section we give some practical results on some small, but difficult circuits. All of the examples are bipolar circuits. We used the Ebers–Moll model, without modeling the reverse breakdown. Our model should be close to the simple model used in the SPICE 2G6 circuit simulator [2]. These operating point benchmarks are available from *netlib* with the request "send all from benchmark/bipopt."

5.1. Example Circuits

The first circuit in our benchmark suite is a cascade of three Schmitt Trigger circuits [26]. This circuit is not particularly useful, but is a tough challenge for operating point algorithms. The output voltages of a Schmitt Trigger will switch abruptly between two extremes for a small excursion of the input voltage, called the "trigger." The two complementary outputs of the center circuit are fed through buffering resistors to the trigger inputs of copies of the same circuit, thus the number of switching possibilities multiply. If V x is set below the lower trigger point, then one output of the middle circuit goes low to about 0.9 V and the complementary output goes high to about 12 V. The voltage on the low output of the center circuit is below the lower trip point of the trigger to which it is connected, and the voltage on the high output of the center circuit is above the upper trip point of the trigger to which it is connected. Thus, the circuit has a unique operating point. However, for a supply voltage below 12 V, the circuit may have multiple solutions.

Newton's method, as implemented in a commercial circuit simulator, is unable to get an operating point for this circuit. Moreover, continuation in the supply voltage is confused by the multiple solutions at intermediate values of the supply. Homotopy methods computed an operating point for this circuit accurate to approximately 10^{-14} A at each node.

An important class of practical circuits that also exhibit multiple operating points are bandgap voltage references (see, for example, [7]). Our second circuit is a bipolar implementation of such a reference. The convergence of a proprietary simulator depends on the initial voltage for the output node of this circuit. For certain settings, the simulator will converge to one of two operating points, at which the output is about 1.23 V (the desired value) or about 14 V. For other initial settings, the simulator does not converge. Homotopy techniques applied to this example, and starting from a variety of widely spaced start-

ing points, exhibited robust convergence. Of the three dc solutions for this circuit, two are stable and one is unstable. In other words, if the actual circuit were to be forced into the unstable state, it will immediately move to one of the two stable states. Thus, the two stable solutions are of most interest to a circuit designer. In our experiments, the homotopy method converged to one of the two stable solutions for any choice of the random vector. This is in contrast to Newton's method, which sometimes converges to the unstable solution.

Our third circuit is the input stage of the famous 741 operational amplifier [22] shown in Fig. 7. The operation of the circuit is sensitive to the Early voltage parameter V a for transistors Q 5 and Q 6. The transistor model used in our experiments does not model the Early effect; however, we can mimic the effect of a large Early voltage on these two transistors by connecting large resistors in parallel between the collector and the emitter. With R V A 1 = R V A 2 = 1 e 1 2 the circuit has a very large dc gain. One "figure of merit" for a numerical implementation of an operating point algorithm is to see how large a value can be used for these resistors while still analyzing the operation of the circuit correctly.

As the input terminals are slightly unbalanced (V P O S ≠ V N E G), the output voltage switches abruptly from approximately 0 V to −12 V. In the middle of this switching region, the circuit operates as a linear amplifier with very high dc gain. The dc gain in the linear region can be estimated by forming a divided difference of two operating points in this region. Our simulation gave an estimate of 5.711×10^9 for the gain, compared to a hand-calculated value of 5.713×10^9 using a linearized model of the transistors. The same perturbation analysis using a commercial simulator produced an estimate of 4.950×10^8—off by a factor of ten. The commercial simulator attaches artificial conductances from each node to ground as a convergence aid; the default value of these conductances is set large enough so that they should not disturb the operation of the circuit. However, the gain of this circuit is so large that the presence of these artificial conductances is significant.

Fig. 8 shows the HOMPACK *convergence trace* for this circuit. The computed trajectory is a path in ten-dimensional space. This figure plots the average of these ten paths against λ:

$$x_{\mathrm{AVG}}(\lambda) = \frac{1}{10} \sum_{k=1}^{10} x_k.$$

This trace can be displayed as the algorithm is working toward the solution, and can provide useful visual feedback to the designer.

Fig. 9 shows a particularly difficult circuit. This is a feedback circuit used in an operational amplifier [5] with a very good common mode rejection ratio. The independent source V 8 is inserted to exercise the feedback loop. For values of V 8 in a small hysteresis range, the circuit has three operating points (one is unstable). A homotopy

method allowed us to determine robustly the two desired (i.e., stable) solutions in the hysteresis region. In practice, this circuit behaves well, but both SPICE [35] and a proprietary simulator failed to analyze the dc behavior of this circuit correctly.

Fig. 10 is of particular interest because our techniques helped improve the design of the circuit. The operation of this circuit is described in [6]. If R 5 and R 6 are given slightly different values, the performance of the circuit will improve, but the possibility of multiple dc solutions arises. This is undesirable in practice, because the circuit will function only as required at one of these dc solutions. Diode-connected transistor Q 1 0 is included as a so-called "start-up" device to force the circuit to have only the desired dc solution. Our methods were able to identify *two* stable dc solutions, even with the start-up device, when the values of R 5 and R 6 differed by about 5%. This information enabled the designer to estimate the manufacturing tolerances necessary for two resistors to ensure that the circuit would behave as anticipated.

5.2. Performance Data and Comparisons

Convergence and timing data were generated for the following circuits from our benchmark suite:

Cascade	Schmitt Trigger cascade
DcNine	Circuit with nine distinct solutions from [48]
Brokaw	Brokaw voltage reference circuit
Hybrid	Hybrid voltage reference of Fig. 10
CmMode(1)	Common mode circuit of Fig. 9 at edge of hysteresis region
CmMode(2)	Common mode circuit of Fig. 9 in hysteresis region

Each circuit was analyzed fifty times using values of the *a* parameter vector uniformly distributed between the most negative and most positive supply voltages in each particular circuit. Numerical parameters were set so that at $\lambda = 0$, the kth node was forced to a value of a_k. Thus, the circuit was exercised over a large space of starting points. Table I shows convergence statistics for the preceding examples, using the routine FIXPNS from HOMPACK with a tracking tolerance of 1.0×10^{-7} [52]. The minimum, average, and maximum number of Jacobian matrix evaluations is reported for each circuit over the fifty trials. Convergence to a valid solution was obtained for every circuit from every starting point. For those circuits with multiple operating points, the homotopy method always converged to one of the two stable solutions.

The preceding results support the theoretical claims of global convergence. In practice, it is possible to set numerical parameters a bit differently so that convergence is achieved more quickly (in case of failure, the circuit is reanalyzed with more stringent parameter settings). Table II shows one sample run for each circuit with parameter settings that seem adequate for general use. Timing results were obtained on a SUN 4 workstation running Berkeley Unix.
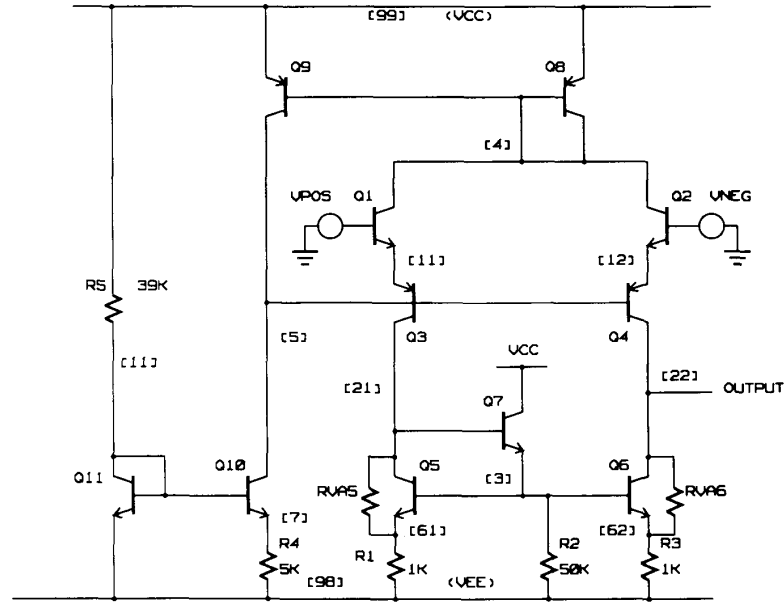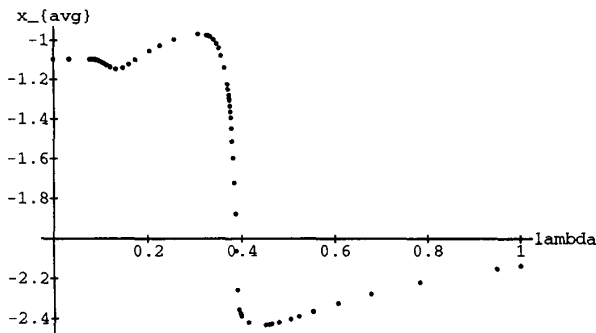
Fig. 7. Input stage of 741 op amp.

Fig. 8. Convergence trace.

Table III shows convergence statistics for the variable gain homotopy of Section 3.2. The data are presented in the form "phase 1 + phase 2." In phase 1, the operating point of the diode-only circuit is calculated using a damped Newton code. In phase 2, the current gains of the transistor models are swept from zero to their final values. It is interesting that more work is spent in phase 1 than in phase 2. Except for the Schmitt Trigger cascade, these iteration counts are within a factor of three of the number of iterations for a damped Newton code alone. Yet the damped Newton method does not always converge, or may converge to a zero of the equations that is physically unstable. Thus, at least in the case of bipolar transistors, we can enjoy the robust convergence of the homotopy method with quite reasonable computing cost. Additional timing results are reported in [46].

The circuits discussed so far are difficult, but small. Does the encouraging performance of the variable gain homotopy scale up to larger circuits typical of those used in full-scale commercial designs? Such circuits will have at least 100 unknowns. To answer this question, we prepared a custom version of HOMPACK, which incorporates a direct sparse solver for nonsymmetric matrices. To demonstrate that our techniques are really industrial strength, we implemented the most detailed bipolar transistor model discussed in [20]. Thus, our timing and convergence results reflect the use of models with the level of complexity necessary for detailed designs. Table IV shows timing and iteration counts for four circuits; all are analog circuits to be incorporated in a bipolar chip. All times are for a SUN SPARCstation 1 + .

Circuits rli13b and is7 exhibit a turning point in the homotopy trajectory. Previous to the application of homotopy methods, an operating point for is7 could be obtained only by a tedious and time-consuming interactive procedure in which a designer would estimate initial values for crucial node voltages. Moreover, even a small change to the second circuit invalidates the previously computed set of initial voltages.

5.3. Example of a Steady-State Computation

Consider a system of linear first-order ODE's driven by a sinusoidal forcing function. Theory shows that the solution to such a system must be sinusoidal with the same period as the driving function. However, for nonlinear ODE's, the situation is much more interesting. Among other possibilities, the solution to the equations may be periodic with the frequency equal to a harmonic of the driving frequency. This effect is used in electronic circuits to generate integer multiples of a reference frequency. Fig. 11 shows the input and output from such a frequency multiplier circuit [15] driven by a sinusoidal input.
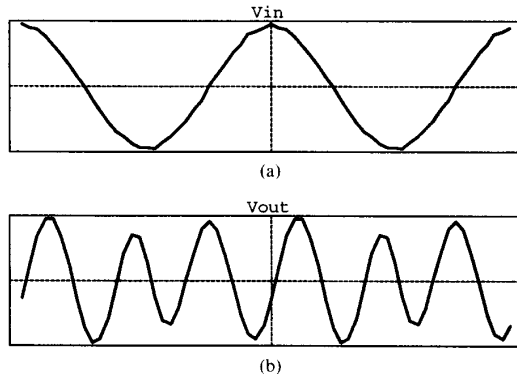
Fig. 9. Difficult common mode bias circuit.

Fig. 10. Hybrid voltage reference.

We applied the discretization technique of Section IV with thirty-two sample points per period to produce a system of algebraic equations that describe the steady-state solution of this circuit. The input and output waveforms are shown annotating the figure. The collector voltage is indeed periodic, with the period equal to one-third of the input period.

5.4. Implementation of a Test Platform

HOMPACK is a general curve tracking package that has no specific knowledge of electronic circuit equations. For our numerical experiments with circuit equations we wrote a driver in the C++ language [44]. The class mechanism and operator overloading facility of C++ allows us to formulate equations in familiar notation. These

Fig. 11. Input and output for frequency tripler circuit.

TABLE I
PERFORMANCE OF VARIABLE STIMULUS HOMOTOPY

Variable Stimulus Homotopy Using **F I X P N S**

Circuit	Number of Iterations		
	min	avg	max
Cascade	351	621	1299
DcNine	163	290	463
Brokaw	299	355	406
Hybrid	346	435	604
CmMode(1)	310	494	787
CmMode(2)	254	341	451

TABLE II
PERFORMANCE WITH NORMAL PARAMETER SETTINGS

Variable Stimulus Homotopy Using **F I X P N S**

Circuit	Number of Iterations	Time(s)
Cascade	496	31.70
DcNine	127	7.60
Brokaw	193	8.70
Hybrid	184	7.25
CmMode(1)	277	11.80
CmMode(2)	195	7.75

TABLE III
PERFORMANCE OF VARIABLE GAIN HOMOTOPY

Variable Gain Homotopy

Circuit	Number of Iterations (Newton + homotopy)
Cascade	12 + 184
DcNine	9 + 23
Brokaw	16 + 27
Hybrid	23 + 16
CmMode(1)	16 + 46
CmMode(2)	18 + 24

expressions do not evaluate to numbers, rather each expression generates a tree data structure in memory. The leaf vertices of this tree are constants or variables introduced by the user, such as unknown node voltages. The

TABLE IV
PERFORMANCE ON LARGER INDUSTRIAL CIRCUITS

Variable Gain Homotopy

Circuit	n	Number of Iterations (Newton + homotopy)	Time(s)
rli13b	44	22 + 137	14.23
ups01a	58	10 + 90	7.70
bgatt	124	23 + 36	11.00
is7	470	45 + 358	410.00

homotopy parameter λ is a special leaf vertex. Internal vertices are labeled with arithmetic operations to be applied to the subtrees of a vertex, or labeled with the address of a subroutine that computes an arbitrary function, using the values of subtrees as arguments to the function. This arrangement allows us to form the Jacobian matrix of the homotopy map using *automatic differentiation* [23]. This is *not* a difference approximation to the derivative, nor is it symbolic differentiation such as would be computed by a symbolic manipulation program. Instead, the chain rule is applied systematically throughout the tree to produce a numerical value for partial derivatives at the same time a numerical value is being computed for the root of each tree.

Here are the equations for the one-transistor bias circuit of Fig. 1, using the variable stimulus homotopy of Section 3.2. This is the complete input required of the user; computation of the Jacobian matrix is automatic.

```
! lambda*(xc-ac)  +  q1.ic( )

  +  (xc-VCC)/RC  ==  0

! lambda*(xb-ab)  +  q1.ib( )

  +  xb/RB2  +  (xb-VCC)/RB1  ==  0

! lambda*(xe-ae)  +  q1.ie( )

  +  xe/RE  ==  0.
```

The operator "!" is defined specially for the homotopy parameter (a derived class) so that ! lambda means (1.0- lambda). The function q1 computes the collector current for an npn transistor, given voltages on the collector, base, and emitter.

Formulating the equations in this manner, let us experiment quickly with a variety of different homotopies, without the tedious and error-prone hand computation of expressions for derivatives. Some simple timing studies have shown that for each Jacobian matrix evaluated at a point of the zero curve, considerably more time is spent in HOMPACK performing linear algebra computations on the matrix than is spent in our driver constructing the derivative. Therefore, the automatic computation of derivatives is not a bottleneck. The need for highly accurate Jacobian matrices varies with the problem domain; high accuracy has been demonstrated to be essential for some fluid mechanics problems [56], [57], and the same is true for circuit analyses here.

VI. Related Work

In addition to the references cited in Section I, we would like to mention some other work related to the material presented herein.

Yamamura *et al.* [64]–[67] have demonstrated the application of homotopy methods for finding the solutions of difficult systems of circuit equations. Their implementation concentrates on *simplicial* methods that are rather different than the technique used in HOMPACK. In particular, [66] reports impressive timing results. It is difficult, however, to make a direct comparison with our benchmark timings, because their implementation assumes that the equations are in the form of a "separable" homotopy. The separability assumption is needed to allow for solving a large system of equations (more than 100 transistors), but is difficult to satisfy for sophisticated bipolar models. Our equation formulation does not have such an assumption.

In [45], we prove that the inner product condition holds in the case of nodal equations with grounded voltage sources. Vandenberghe and Vandewalle [49], [50] presented a similar result for a more general class of equations. Their implementation results are also based on simplicial techniques. Therefore, direct timing comparisons are difficult.

An algorithm along the lines of HOMPACK that was developed specifically to find the dc operating points of difficult nonlinear electronic circuits appears in [13] and [48]. They use a homotopy similar to (2), but do not include the random element. For a certain restricted class of circuits, they prove that their method is able to find *all* the dc operating points.

VII. Summary

Robust computation of a dc operating point (or points) for a nonlinear network is an essential first step for a circuit simulator. Design experience shows that this problem is most acute for analog circuits using bipolar transistors. Even small networks (ten or fewer transistors) can exhibit qualitative behavior that makes the operating point problem difficult. Such behavior includes 1) multiple solutions (among which some are physically stable and some unstable), 2) bifurcations in solution manifolds, and 3) turning points in solution manifolds. Larger networks (100–1000 transistors) present the same qualitative difficulties with additional numerical difficulties due to size.

The theory of globally convergent homotopy methods addresses the qualitative issues in solving operating point equations. Assuming sufficiently smooth modeling functions, and a so-called "coercivity condition," the theory demonstrates the existence of a smooth solution manifold leading to at least one solution to the operating point equations. We have demonstrated that the required coercivity condition is quite natural for the equations that arise in circuit simulation.

Various homotopies have been presented, all meeting the conditions for global convergence, but differing in terms of performance. Using insight gained from studying these different homotopies, we have formulated the "variable gain" homotopy for bipolar networks that has excellent performance. The time required to solve a system of operating point equations with this homotopy is not more than two to three times slower than the time required to solve the same equations by less widely convergent methods, yet the homotopy method works even on examples that cause competing methods to fail. Thus at the cost of a little extra computer time, on the average, a user is freed from the tedious task of "nursing" a simulator into convergence. In extensive trials on difficult circuits generated with AT&T, the homotopy method has solved every circuit. The example circuits have ranged in size from ten to more than 1800 unknowns. The homotopy code used was a version of HOMPACK modified to use automatic differentiation, arbitrary sparse matrix data structures, and direct instead of iterative solvers for the sparse linear equations.

Finally, we have considered the solution of a two-point boundary value problem that represents the operation of a circuit in the time domain. A finite difference formulation of the circuit's differential equations has been shown to be equivalent to an operating point problem on a larger dc network. The coercivity condition required in the hypothesis for global convergence of a homotopy method has been established for the equations of this larger network. Thus, robust convergence is enjoyed for the time domain case as well.

References

[1] P. Antognetti, D. O. Pederson, and H. De Man, *Computer Design Aids for VLSI Circuits*. Boston, MA: Martinus Nijhoff, 1986, pp. 69–73.

[2] P. Antognetti and G. Massobrio, *Semiconductor Device Modeling with SPICE*. New York: McGraw-Hill, 1988, pp. 13–22, 40–55.

[3] T. J. Aprille and T. N. Trick, "Steady state analysis of nonlinear circuits with periodic inputs," *Proc. IEEE*, vol. 60, pp. 108–114, 1972.

[4] R. E. Bank and D. J. Rose, "Global approximate Newton methods," *Numer. Math*, vol. 37, pp. 279–295, 1981.

[5] M. Banu, J. M. Khoury and Y. Tsividis, "Fully differential operational amplifier with accurate output balancing," *IEEE J. Solid-State Circuits*, vol. 23, pp. 1410–1414, 1988.

[6] T. C. Banwell, "An all-NPN hybrid voltage reference," *IEEE J. Solid-State Circuits*, vol. 26, pp. 77–80, 1991.

[7] A. P. Brokaw, "A simple three-terminal IC bandgap reference," *IEEE J. Solid-State Circuits*, vol. 9, pp. 388–393, 1974.

[8] L. E. J. Brouwer, "Uber abbildung von mannigfältigkeiten," *Math. Ann.*, vol. 71, pp. 97–115, 1912.

[9] K. S. Chao, D. K. Liu, and C. T. Pan, "A systematic search method for obtaining multiple solutions of simultaneous nonlinear equations," *IEEE Trans. Circuits Sys.*, vol. 22, pp. 748–753, 1975.

[10] K. S. Chao and R. Saeks, "Continuation methods in circuit analysis," *Proc. IEEE*, vol. 65, pp. 1187–1194, 1977.

[11] P. Chauffoureaux and M. Hasler, "Monotonicity in nonlinear resistive circuits," in *Proc. IEEE Int. Symp. Circuits and Systems*, May 1990, pp. 395–398.

[12] S. Chow, J. Mallet-Paret, and J. A. Yorke, "Finding zeroes of maps: homotopy methods that are constructive with probability one," *Math. Comp.*, vol. 32, pp. 887–899, 1978.

[13] L. O. Chua and A. Ushida, "A switching-parameter algorithm for finding multiple solutions of nonlinear resistive circuits," *Int. J. Circuit Theory Appl.*, vol. 4, pp. 215–239, 1976.

[14] L. O. Chua and N. N. Wang, "On the application of degree theory to the analysis of resistive nonlinear networks," *Int. J. Circuit Theory Appl.*, vol. 5, pp. 35–68, 1977.

[15] K. K. Clarke and D. T. Hess, *Communication Circuits: Analysis and Design*. Reading, MA: Addison-Wesley, 1978, pp. 6–8.

[16] W. M. Coughran, Jr., E. H. Grosse, and D. J. Rose, "CAzM: a circuit analyzer with macromodeling," *IEEE Trans. Elect. Dev.*, vol. 30, pp. 1207–1213, 1983.

[17] R. J. Duffin, "Nonlinear networks IIa," *Bull. Amer. Math. Soc.*, vol. 53, pp. 963–971, 1947.

[18] ECAP, "1620 Electronic Circuit Analysis Program [ECAP] [1620-EE-02X]," IBM application program file H20-0170-1, 1965.

[19] C. B. Garcia and W. I. Zangwill, *Pathways to Solutions, Fixed Points, and Equilibria*. Englewood Cliffs, NJ: Prentice-Hall, 1981, pp. 1–23.

[20] I. Getreu, *Modeling the Bipolar Transistor*. Beaverton, OR: Tektronix, 1976, pp. 9–23.

[21] B. Gopinath and D. Mitra, "When are transistors passive?," *Bell Sys. Tech. J.*, vol. 50, pp. 2835–2847, 1971.

[22] P. R. Gray and R. G. Mayer, *Analysis and Design of Analog Integrated Circuits*, 2nd ed. New York: Wiley, 1984, pp. 363–373.

[23] A. Griewank, "On automatic differentiation," *Mathematical Programming 88*. Boston, MA: Kluwer, 1989.

[24] M. Hasler and J. Neirynck, *Nonlinear Circuits*. Norwood, MA: Artech House, 1986.

[25] D. A. Hodges, Special issue on computer aided circuit analysis and device modeling, *IEEE J. Solid-State Circuits*, vol. 6, 1971.

[26] P. Horowitz and W. Hill, *The Art of Electronics*. Cambridge, MA: Cambridge University Press, 1983, pp. 124–127.

[27] C. W. Ho, A. E. Ruehli, and P. A. Brennan, "The modified nodal approach to network analysis," *IEEE Trans. Circuits Sys.*, vol. 22, pp. 504–509, 1975.

[28] H. B. Keller, "Numerical solution of bifurcation and nonlinear eigenvalue problems," in *Applications of Bifurcation Theory*, P. Rabinowitz. New York: Academic Press, 1977, pp. 359–384.

[29] S. Kesavan, *Topics in Functional Analysis and Applications*. New York: Wiley, 1989, p. 219.

[30] K. S. Kundert, J. K. White, and A. Sangiovani-Vincentelli, *Steady-State Methods for Simulating Analog and Microwave Circuits*. Boston, MA: Kluwer, 1990, pp. 55–60.

[31] E. Lahaye, "Une méthode de résolution d'une catégorie d'équation transcendantes," *C. R. Acad. Sci.*, pp. 1840–1842, 1932.

[32] W. J. McCalla, *Fundamentals of Computer-Aided Circuit Simulation*. Boston, MA: Kluwer, 1988, pp. 53–85.

[33] L. Miller, Private communications, 1990.

[34] L. D. Milliman, W. A. Massena, and R. H. Dickhaut, "CIRCUIT—a digital computer program for transient analysis of electronic circuits—user's guide," Boeing Co., Harry Diamond Lab., Rep. AD-346-1, Jan. 1967.

[35] L. Nagel, "SPICE2: a computer program to simulate semiconductor circuits," ERL Memo. ERL-M520, Univ. of California, Berkeley, May 1975.

[36] D. Nelson, Private communications, 1990.

[37] J. M. Ortega and W. C. Rheinboldt, *Iterative Solutions of Nonlinear Equations in Several Variables*. New York: Academic Press, 1969, pp. 161–165.

[38] S. L. Richter and R. A. DeCarlo, "Continuation methods: theory and applications," *IEEE Trans. Circuits Sys.*, vol. 30, pp. 347–352, 1983.

[39] A. E. Ruehli, *Circuit Analysis, Simulation and Design, Part 1*. New York: North-Holland, 1986, pp. 207–234.

[40] F. M. Salam, L. Ni, X. Sun, and S. Guo, "Parallel processing for the steady state solutions of large-scale non-linear models of power systems," in *Proc. IEEE Int. Symp. on Circuits and Systems*, May 1989, pp. 1851–1854.

[41] A. Sard, "The measure of the critical values of differential maps," *Bull. Amer. Math. Soc.*, vol. 48, pp. 883–890, 1942.

[42] M. Spivak, *Calculus on Manifolds*. Boston, MA: Benjamin/Cummings, 1968.

[43] G. Strang, *Linear Algebra and its Applications*. Orlando, FL: Harcourt Brace Jovanovich, 1988, App. A.

[44] B. Stroustrup, *The C++ Programming Language*. Reading, MA: Addison-Wesley, 1986.

[45] LJ. Trajković, R. C. Melville, and S. C. Fang, "Passivity and no-gain properties establish global convergence of a homotopy method for DC operating points," *Proc. IEEE Int. Symp. Circuits and Systems*, May 1990, pp. 914–917.

[46] LJ. Trajković, R. C. Melville, and S. C. Fang, "Improving DC convergence in a circuit simulator using a homotopy method," in *Proc. IEEE Custom Integrated Circuits Conf.*, 1991, pp. 8.1.1–8.1.4.

[47] LJ. Trajković, R. C. Melville, and S. C. Fang, "Finding DC operating points of transistor circuits using homotopy methods," *Proc IEEE Int. Symp. Circuits and Systems*, 1991, pp. 758–761.

[48] A. Ushida and L. O. Chua, "Tracing solution curves of nonlinear equations with sharp turning points," *Int. J. Circuit Theory Appl.*, vol. 12, pp. 1–21, 1984.

[49] L. Vandenberghe and J. Vandewalle, "A globally convergent algorithm for solving a broad class of nonlinear resistive circuits," *Proc. IEEE Int. Symp. Circuits and Systems*, May 1990, pp. 403–406.

[50] L. Vandenberghe and J. Vandewalle, "Variable dimension algorithms for solving nonlinear resistive circuits," in *Proc. Eur. Conf. on Circuit Theory and Design*, Sept. 1989, pp. 385–389.

[51] A. Wallace, *Differential Topology: First Steps*. Reading, MA: Benjamin/Cummings, 1968.

[52] L. T. Watson, S. C. Billups, and A. P. Morgan, "Algorithm 652: HOMPACK: a suite of codes for globally convergent homotopy algorithms," *ACM Trans. Math. Software*, vol. 13, pp. 281–310, 1987.

[53] L. T. Watson, "Numerical linear algebra aspects of globally convergent homotopy methods," *SIAM Rev.*, vol. 28, pp. 529–545, 1986.

[54] L. T. Watson, "An algorithm that is globally convergent with probability one for a class of nonlinear two-point boundary value problems," *SIAM J. Numer. Anal.*, vol. 16, pp. 394–401, 1979.

[55] L. T. Watson, "Solving finite difference approximations to nonlinear two-point boundary value problems by a homotopy method," *SIAM J. Sci. Stat. Comput.*, vol. 1, pp. 467–480, 1980.

[56] L. T. Watson and M. R. Scott, "Solving spline-collocation approximations to nonlinear two-point boundary-value problems by a homotopy method," *Appl. Math. Comput.*, vol. 24, pp. 333–357, 1987.

[57] L. T. Watson and M. R. Scott, "Solving Galerkin approximations to nonlinear two-point boundary value problems with a globally convergent homotopy method," *SIAM J. Sci. Stat. Comput.*, vol. 8, pp. 768–789, 1987.

[58] W. T. Weeks, A. J. Jiminez, G. W. Mahoney, D. Mehta, H. Qassemzadeh, and T. R. Scott, "Algorithms for ASTAP—a network-analysis program," *IEEE Trans. Circuits Sys.*, vol. 20, pp. 628–634, 1973.

[59] G. Wettlaufer and W. Mathis, "Finding all DC-equilibrium-points of nonlinear circuits," in *Proc. 32nd Midwest Symp. on Circuits and System*, 1989, pp. 462–465.

[60] A. N. Willson, Jr., *Nonlinear Networks: Theory and Analysis*. New York: IEEE Press, 1975.

[61] A. N. Willson, Jr., "The no-gain property for networks containing three-terminal elements," *IEEE Trans. Circuits Sys.*, vol. 22, pp. 678–687, 1975.

[62] G. R. Wilson, "A monolithic junction FET-npn operational amplifier," *IEEE J. Solid-State Circuits*, vol. 3, pp. 341–348, 1968.

[63] J. L. Wyatt, Jr., L. O. Chua, J. W. Gannett, I. C. Gonkar, and D. N. Green, "Energy concepts in the state-space theory of nonlinear n-ports: part I—passivity," *IEEE Trans. Circuits Sys.*, vol. 28, pp. 48–61, 1981.

[64] K. Yamamura and K. Horiuchi, "A globally and quadratically convergent algorithm for solving nonlinear resistive networks," *IEEE Trans. Computer-Aided Design*, vol. 9, pp. 487–499, 1990.

[65] K. Yamamura and M. Kiyoi, "A piecewise linear homotopy method with the use of the Newton homotopy and polyhedral subdivision," *Trans. IEICE*, vol. 73, pp. 140–148, 1990.

[66] K. Yamamura and K. Horiuchi, "Solving, nonlinear resistive networks by a homotopy method using a rectangular subdivision," *Proc. IEEE Int. Symp. Circuits and Systems*, June 1988, pp. 1225–1228.

[67] K. Yamamura and M. Ochiai, "An efficient algorithm for finding all solutions of piecewise-linear resistive circuits," in *Proc. IEEE Int. Symp. Circuits and Systems*, 1991, pp. 3039–3042.

[68] P. Yang, "Circuit simulation and modeling," *IEEE Circuits Devices Mag.*, vol. 5, pp. 48–50, 1989; vol. 6, pp. 8–10, 1990.

[69] C. W. Yun and K. S. Chao, "Simple solution curves of non-linear resistive networks," *Int. J. Circuit Theory Appl.*, vol. 11, pp. 47–55, 1983.

Robert C. Melville received the B.S. degree from the University of Delaware in 1976 and the Ph.D. degree from Cornell University in 1981, both in computer science.

He was a junior faculty member at The Johns Hopkins University before joining AT&T Bell Laboratories in 1985. His work at the labs concentrates on CAD tools for simulation and design of analog circuits, with particular emphasis on the qualitative analysis of nonlinear systems.

He is a member of SIAM, and has served as a referee for the IEEE TRANSACTIONS ON COMPUTER-AIDED DESIGN, the *SIAM Journal of Computing*, and numerous professional conferences.
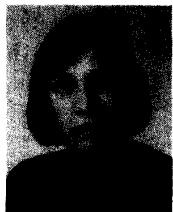
Ljiljana Trajković received the Dipl. Ing. degree from the University of Priština, Yugoslavia, in 1974, the M.S. degree in electrical engineering and the M.S. degree in computer engineering from Syracuse University, Syracuse, NY, in 1979 and 1981, respectively, and the Ph.D. degree in electrical engineering from UCLA, in 1986.

From 1986 to 1988 she was a visiting lecturer at the Department of Electrical Engineering at UCLA. She was a member of technical staff at the AT&T Bell Laboratories, Murray Hill, NJ, from 1988 to 1990. Since 1990 she has been a member of the technical staff in the Computer Networking Research Department at Bell Communications Research, Morristown, NJ. Her research interests include network management and congestion control of high-speed data networks, computer-aided circuit analysis and design, and the theory of nonlinear networks and dynamical systems.

She was a Fulbright Scholar in 1978 and a recipient of the Amelia Earhart Fellowship in 1985 and 1986. Dr. Trajković is a member of ACM and Sigma Xi.

San-Chin Fang received the B.S. degree from Princeton University in 1974, the M.S. degree from Stanford University in 1977, and the Ph.D. degree from Columbia University in 1983.

In 1982, he joined AT&T Bell Laboratories in Murray Hill, NJ, as a member of the technical staff to work on analog CAD tools and high-speed analog circuit design. Currently, he is a supervisor in the VLSI Technology Laboratory, responsible for circuit design research. From June 1985 to May 1986 he was on leave from Bell Labs at the Center for Telecommunications Research, Columbia University, as an Associate Research Scientist. He has also taught several courses in analog IC design at Columbia University from 1982 to 1984.

Layne T. Watson received the B.A. degree (magna cum laude) in psychology and mathematics from the university of Evansville, Evansville, IN, in 1969, and the Ph.D. degree in mathematics from the University of Michigan, Ann Arbor, in 1974.

He was worked for USNAD Crane, Sandia National Laboratories, and General Motors Research laboratories, and has served on the faculties of the University of Michigan and Michigan State University before coming to Virginia Polytechnic Institute and State University, Blacksburg, where he is currently a professor of computer science and mathematics. His current research interests include fluid dynamics, structural mechanics, homotopy algorithms, parallel computation, mathematical software, and image processing.

Dr. Watson is a member of Phi Kappa Phi, Blue Key, Psi Chi, Kappa Mu Epsilon, Phi Beta Chi, ACM, and SIAM. He is an Associate Editor for the *SIAM Journal on Optimization* and the *ORSA Journal on Computing*.