Homotopy Techniques for Obtaining a DC Solution of Large-Scale MOS Circuits

J.S. Roychowdhury R.C. Melville

AT&T Bell Laboratories, Murray Hill, NJ

Abstract

A new technique for obtaining a DC operating point of large, hard-to-solve MOS circuits is reported in this paper. Based on homotopy, the technique relies on the provable global convergence of arclength continuation and uses a novel method for embedding the continuation parameter into MOS devices. The new embedding circumvents inefficiencies and numerical failures that limit the practical applicability of previous simpler embeddings. Use of the technique in a production environment has led to the routine solution of large, previously hard-to-solve circuits.

1 Introduction

A problem of considerable practical importance in circuit design is that of finding the DC operating point of a nonlinear circuit. The operating point, apart from being important in itself, is a prerequisite for subsequent tasks such as small-signal, transient, and noise analyses that are repeatedly invoked in the design of a circuit. From a simulation standpoint, finding an operating point corresponds to solving a system of nonlinear equations describing the circuit. The Newton-Raphson method, widely used in circuit simulators for solving nonlinear equations, often fails to converge to a solution. Despite the advent of other approaches (e.g., stepping, pseudo-transient), no technique has emerged that solves the operating point problem reliably and efficiently. As a result, DC convergence problems often create a significant bottleneck in the design process, especially for large circuits.

Homotopy or continuation methods [1, 2] are a relatively recent numerical technique for solving systems of nonlinear equations. These methods are appealing because they can be proven to be globally convergent, a property that eludes other nonlinear solution techniques. Previous applications of homotopy to solving the operating point problem include parameter switching [3], piecewise-linear (PL) methods [4, 5] and arclength continuation [6, 7]. Despite the guarantee of convergence offered by the theory of homotopy, however, it has not so far had a significant impact on practical circuit simulation and design.

In this work, a technique for the successful practical application of arclength continuation to the operating point problem for large circuits, typically consisting primarily of MOS devices, is described. The novelty of the method lies in a new way of embedding the continuation parameter λ into the device equations. The new embedding circumvents inefficiencies and numerical failures that limit the applicability of previous simpler embeddings. Existing device models, which represent considerable investment especially for in-house simulators, can be used with this method. Implemented in a proprietary circuit simulator, the technique is now in production use within AT&T and has been found to have a near 100% success rate.

An overview of arclength continuation is provided in Section 2. In Section 3, the new embedding is described. Results on large, difficult-to-solve industrial circuits are presented in Section 4.

2 Arclength continuation and a simple embedding

The principle of continuation is similar to that of source or GMIN stepping[1], familiar to users of circuit simulators such as SPICE. The circuit equations are modified by introducing a *continuation parameter* λ. The parameter is first set to a value ($\lambda = 0$ by convention) at which the circuit becomes easy to solve or its solution becomes known. The parameter is then slowly changed back to a value at which the original circuit is retrieved ($\lambda = 1$) and simultaneously, the solution of the changing circuit is followed. The underlying hypothesis is that small changes in the parameter cause small changes to the circuit and its solution, hence the new solution is easy to obtain using numerical techniques with local convergence properties (e.g., the Newton-Raphson method). It appears natural to expect this hypothesis to hold for circuits described by equations that are smooth (i.e., continuous and several times differentiable).

For many practical circuits, however, this hypothesis is not true. A familiar example is the Schmitt trigger circuit, where stepping can fail at critical values of the continuation parameter because the state of the circuit can change abruptly from low to high (and vice-versa) for even the slightest monotonic change in the parameter. The phenomenon is illustrated in Figure 1. As the supply voltage V_{cc} is stepped upward from 0V, the output of the circuit suffers a large jump at $V_{cc} \approx 4.5$V. Points where monotonic increase or decrease of the continuation parameter leads to an abrupt jump in the solution are termed *turning points* or *folds*. Many practical feedback systems composed of smoothly-behaved components exhibit turning points that can cause stepping algorithms to fail.

It is in the treatment of turning points that homotopy differs from stepping. By appropriately incrementing or decrementing the continuation parameter, discontinuities in the solution are avoided when turning points are encountered. For the Schmitt trigger characteristic of Figure 1,

[1] Stepping is also known as *monotonic continuation*.

33rd Design Automation Conference ®

DAC 96 - 06/96 Las Vegas, NV, USA 0-89791-833-9/96/0006..$3.50

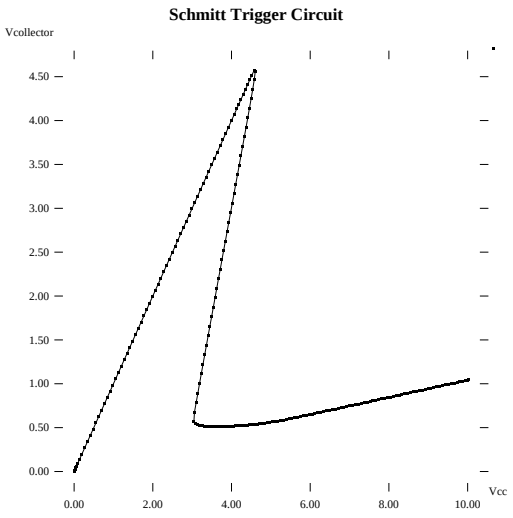

Figure 1: Schmitt trigger V_{out} vs V_{cc}

this corresponds to reducing V_{cc} after the turning point at $V_{cc} \approx 4.5\text{V}$ is reached, taking care to follow the central section of the characteristic and not backtrack onto the initial section already covered. Another turning point is reached at $V_{cc} \approx 3\text{V}$, after which V_{cc} is increased again and the lower right section of the characteristic followed.

Some versions of continuation (e.g., parameter switching) rely on the detection and special treatment of turning points. Arclength continuation, on the other hand, can negotiate turning points automatically without their explicit detection. Its power stems from that it does not treat the continuation parameter λ differently from the unknowns of the circuit being solved for, but treats it as another unknown whose next value on the curve it determines. More precisely, the technique solves a special differential equation (the *defining ODE*) that produces as output a sequence of values of λ (in general not monotonically increasing) together with solutions of the circuit at these values of λ. The key property of this sequence is that a point where $\lambda = 1$ is always reached; therefore that point is the desired solution of the original circuit.

Any nonlinear circuit's equations can be put in the general form (barred variables denote vectors):

$$\bar{g}(\bar{x}) = \bar{0} \tag{1}$$

In order to apply homotopy, the parameter λ is embedded into this system and a related system

$$\bar{f}(\bar{x}, \lambda) = \bar{0} \tag{2}$$

is obtained. The embedding of λ is designed so as to reduce $\bar{f}$ to the original system at $\lambda = 1$, i.e., $\bar{f}(\bar{x}, 1) \equiv \bar{g}(\bar{x})$. In addition, the *start system* $\bar{f}(\bar{x}, 0) = \bar{0}$ is constructed so as be easy to solve by traditional methods. By solving the defining ODE for the problem, the arclength continuation algorithm generates samples $\{(\bar{x}^i, \lambda^i)\}$ of the *continuation track*. Every point on the track satisfies Equation 2, hence when $\lambda^i = 1$ is reached, $\bar{x}^i$ is a solution of $\bar{g}(\bar{x}) = \bar{0}$.

Different types of continuation track are possible, as shown in Figure 2. The horizontal axis depicts the progress

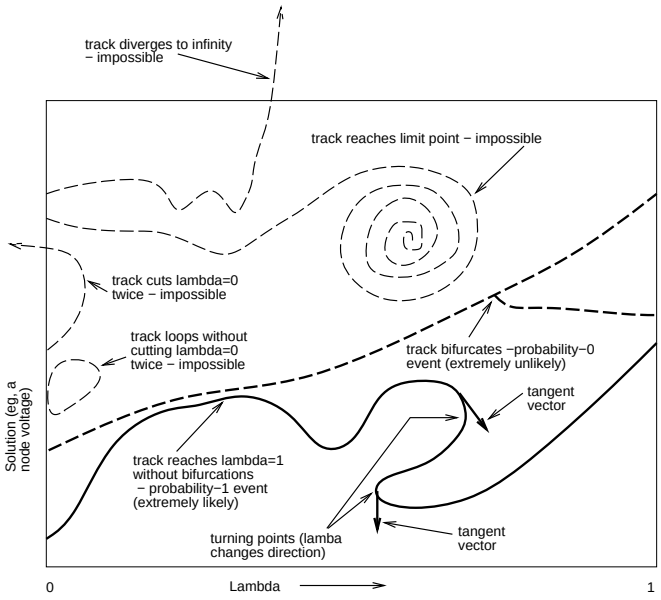

Figure 2: Different kinds of homotopy track

of λ, which varies between 0 and 1. The vertical axis represents the solution of the circuit at a given value of λ. The algorithm starts at $\lambda = 0$ and generates points on the track until $\lambda = 1$ is reached.

It can be shown that the top four kinds of track in Figure 2 cannot occur with arclength continuation, and that the situation depicted by the fifth (dashed bold) track is extremely unlikely to occur (a probability-0 event). Another possibility (not depicted), a bounded space-filling curve that does not reach $\lambda = 1$, can also be shown to be impossible. The lowermost track illustrates the normal, extremely likely (or probability-1) case of tracks that reach $\lambda = 1$ without bifurcations.

An important concept in arclength continuation is that of the *tangent vector*, which has a simple interpretation: it is the tangent to the track at any point. Two instances of the tangent vector are shown on the lowermost track. Roughly speaking, the algorithm proceeds by calculating the tangent vector from which it determines the next point on the curve by extrapolation. Turning points correspond to the λ-component of the tangent vector becoming 0; two turning points are shown on the lowermost curve in the figure.

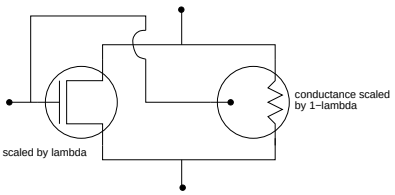

Figure 3: Simple homotopy for MOSFETs

The choice of the embedding $\bar{f}$ is crucial for continuation to work robustly and efficiently in practice. A simple embedding that satisfies the requirements of arclength con-

tinuation is the following:

$$\bar{f}(\bar{x},\lambda) \equiv \lambda \bar{g}(\bar{x}) + (1-\lambda)(\bar{x}-\bar{a}) \qquad (3)$$

where $\bar{a}$ is a constant vector. The above equation has a simple circuit interpretation: the current through each device is multiplied by λ and new resistors and current sources are added from each node to ground, of conductance $1-\lambda$ and current $(1-\lambda)a_i$ respectively. A variation is to consider only the nonlinear elements (e.g., MOS devices) when adding the conductances, depicted in Figure 3. The currents through the resistor are multiplied by $1-\lambda$ while the currents through the MOSFET are multiplied by λ, as indicated by Equation 3. This straightforward embedding is however ineffective for most large circuits. Excessively long tracks or outright failure of the numerical method for solving the defining ODE result. Solving for an operating point using arclength continuation with this embedding can be inferior to even traditional Newton-Raphson methods.

3 The BLHOM MOS homotopy model

In this section, an embedding of λ that is effective for large MOS circuits is described. A key feature of this MOS homotopy model is that it is constructed with *two* λ parameters, λ_1 and λ_2. In Section 3.1, the dependence of the model on these parameters is described; in Section 3.2, the use of a single-λ-based arclength continuation solver with a coupling between λ_1 and λ_2 is presented.

3.1 Dependence on λ_1 and λ_2

The BLHOM model is symmetric and bulk-referenced [8, 9], taking the electrical inputs $V_{gb} = V_g - V_b$, $V_{sb} = V_s - V_b$ and $V_{db} = V_d - V_b$. V_s, V_b, V_g and V_d represent the voltages at the source, bulk, gate and drain nodes respectively. In addition, the model uses two homotopy parameters λ_1 and λ_2 which take values in $[0,1]$. λ_1 influences the drain-source driving-point characteristic whereas λ_2 controls the transfer characteristic, i.e., the influence of the gate on the drain current.

The form of the drain-source current I_{ds} for the BLHOM homotopy is:

$$I_{ds} = \beta\left[V'_{gs}(V_{gb}, V_{db}, V_{sb}, \lambda_2)\right]^2 h(V_{db} - V_{sb}, \lambda_1) \qquad (4)$$

Equation 4 is a single-piece model, qualitatively resembling the Schichman-Hodges (SH) model in that it contains a quadratic term in V_{gs} multiplying a term determined by V_{ds}. An appreciation of how varying λ_1 and λ_2 affects the characteristics of the model can be gained from Figure 4. Each small three-dimensional plot represents the variation of the drain-source current (plotted on the vertical axis) as a function of the gate-source and drain-source voltages (represented on the horizontal axes) at fixed values of λ_1 and λ_2. λ_1 and λ_2 vary on the large vertical and horizontal axes. The bottom left corner depicts the $(\lambda_1, \lambda_2) = (0,0)$ case and the top right the $(1,1)$ case. Moving vertically from bottom to top, λ_1 increases from 0 to 1; likewise, λ_2 increases from 0 to 1 horizontally from left to right.

At $(\lambda_1, \lambda_2) = (1,1)$ (the top right), the model characteristics are similar to that of the SH model, exhibiting a quadratic dependence on V_{gs} and linear and saturation regions as a function of V_{ds}. At $(\lambda_1, \lambda_2) = (0,0)$ (the bottom left), it can be seen that there is no transfer characteristic (varying V_{gs} does not alter I_{ds}), and that the driving point characteristic is much less sharp than for the original MOSFET. The start system corresponds to $(\lambda_1, \lambda_2) = (0,0)$,

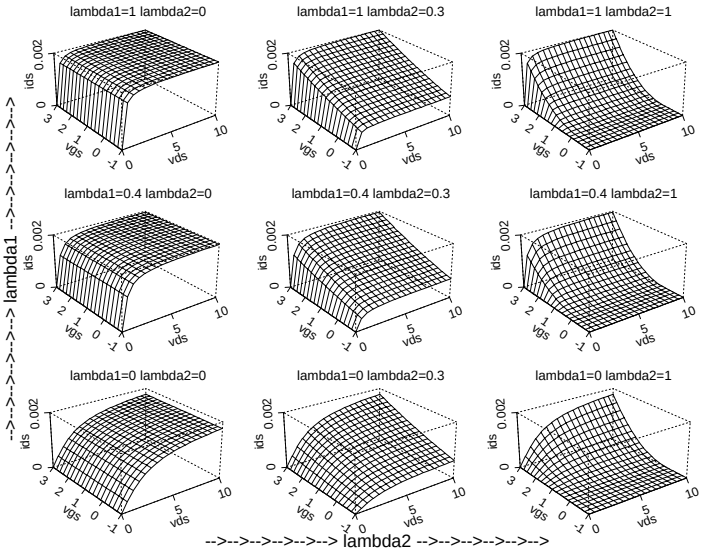

Figure 4: **BLHOM**: Model characteristics as a function of λ_1, λ_2

at which each MOS device becomes a two-terminal almost-linear resistor; hence the circuit becomes easy to solve using the Newton-Raphson method[2]. The effect of varying λ_1 and λ_2 is also apparent from the figure: λ_1 sharpens the driving point characteristic without affecting the gain whereas λ_2 ramps the gain without sharpening the driving-point characteristic.

3.2 Homotopy using two λ parameters

Practical arc-length continuation algorithms [2] are based on a single continuation parameter λ, leading to a system of n equations in $n+1$ variables. Since BLHOM has two continuation parameters, a system of n equations in $n+2$ variables results. One approach to converting this into a one-parameter homotopy is to add an extra equation to obtain a system of $n+1$ equations in $n+2$ variables to which a conventional homotopy solver can be applied.

It is necessary for the extra equation to be specified such that the solution of the original circuit is respected and that the requirements for arclength continuation continue to hold. Any smooth curve relating only λ_1 and λ_2 and passing through $(\lambda_1, \lambda_2) = (0,0)$ and $(\lambda_1, \lambda_2) = (1,1)$ satisfies the above conditions. An infinite number of such curves is possible; one such family $\phi_m(\lambda_1, \lambda_2) = 0$ is shown in Figure 5. As $m \to 0$, $\psi_m(\lambda_1, \lambda_2) \to \lambda_1 - \lambda_2$; as m increases from 0, $\psi_m(\lambda_1, \lambda_2) = 0$ is shown by the upper curves in the figure; likewise, as m decreases, $\psi_m(\lambda_1, \lambda_2) = 0$ is shown by the lower curves. Of interest are the limiting curves obtained as $m \to \pm\infty$, given by the left and upper boundaries of Figure 5, and by its lower and right boundaries, respectively. Corresponding to these limit curves are the first column and top row of Figure 4, and the bottom row and third column, respectively.

While these limit curves are not smooth (violating smoothness requirements for arclength continuation methods),

[2] Typically, the start system takes fewer than 10 iterations to solve using Newton-Raphson.

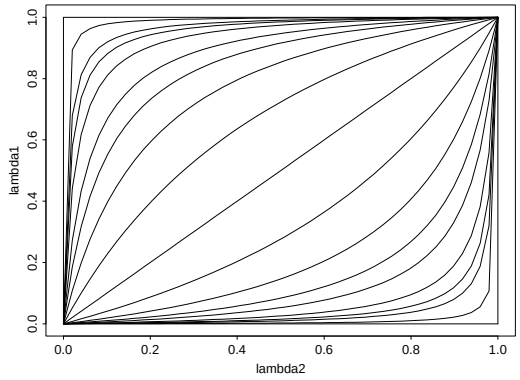

Figure 5: $\phi_m(\lambda_1, \lambda_2) = 0$

they have the property of decoupling the homotopy into two independent parts, one controlled by λ_1 keeping λ_2 fixed, the other by λ_2 keeping λ_1 fixed. For the $m \to \infty$ limit curve, λ_1 is ramped first, whereas for the $m \to -\infty$ limit curve, λ_2 is ramped first. In this work, the lower curve ($m \to -\infty$) is used; the horizontal and vertical segments of this path are referred to henceforth as *phase 1* and *phase 2* of the homotopy, respectively. It has been observed that using the $m \to -\infty$ limit leads to a robust and efficient DC solution technique, while the $m \to \infty$ curve causes failures due to inefficiency and numerical problems. An intuitive understanding of this behaviour is provided by Figure 4, where it can be seen that the latter path is "smoother" than the former, which reaches a highly nonlinear characteristic at $(\lambda_1, \lambda_2) = (1, 0)$ before becoming smoother again at $(\lambda_1, \lambda_2) = (1, 1)$.

For practical design, it is necessary to obtain the operating point of the circuit using existing in-house MOS models that have been characterized to model fabricated devices very accurately [8]. The utility of BLHOM lies in that the operating point obtained with it is very similar to that with more accurate models – hence this operating point can be used as a starting guess to solve the circuit with standard models using, for example, the Newton-Raphson method, relying on its local convergence properties. This approach works very well for most circuits. It is possible, however, to use continuation for smoothly substituting the standard model for the BLHOM model as well. Each MOSFET is replaced by a composite weighted combination of BLHOM and the accurate model with the weights depending on a third continuation parameter λ_3. Using continuation of λ_3 (*phase 3*), the composite is changed smoothly from BLHOM at $\lambda_3 = 0$ to the accurate model at $\lambda_3 = 1$.

From a theoretical standpoint, it is preferable to perform all three phases (ramping λ_2, λ_1 and λ_3) as part of a single smooth homotopy, since it restores smoothness conditions that are violated by the approach outlined in the previous paragraphs[3]. This can be achieved by the straightforward extension of the construction of Figure 5 to three continuation parameters. Our experience however has been that in practice, very few circuits fail as a result of the sharp corners in the limit curves of Figure 5 and its three-dimensional extension; only one has in fact been identified, out of a conservative estimate of a few thousand conventionally hard-to-solve circuits on which the three-phase technique has been effective. The three-phase technique is preferred over the single unified homotopy because implementation becomes significantly simpler due to the decoupling of the λ_2, λ_1 and λ_3 homotopies. Further, a saving in computation is also achieved during the first and second phases because BLHOM is several times less expensive to compute than accurate in-house MOS models.

[3]It is assumed throughout that all device models are smooth, a condition that BLHOM and most AT&T MOS models satisfy, being C^∞.

4 Results

The BLHOM homotopy described in the previous section is in production use by the AT&T MOS design community. In this section, examples demonstrating its global convergence property are presented.

Circuit	Type	Size	Homotopy (CPU secs)	ADVICE CPU secs for no convergence
dlopata1	analog	127	13	4331
heideh	analog	192	49	244
test9	A/D	1380	599	3209
vf_test	A/D	1621	565	2101
rabb-xare	A/D	1877	1035	2340
addas.com	A/D	3413	1195	4395
s1423	digital	3736	678	4207
dctl.t	A/D	7199	10385	10150
goh	digital	8489	3339	11700

Table 1: Homotopy vs conventional algorithms

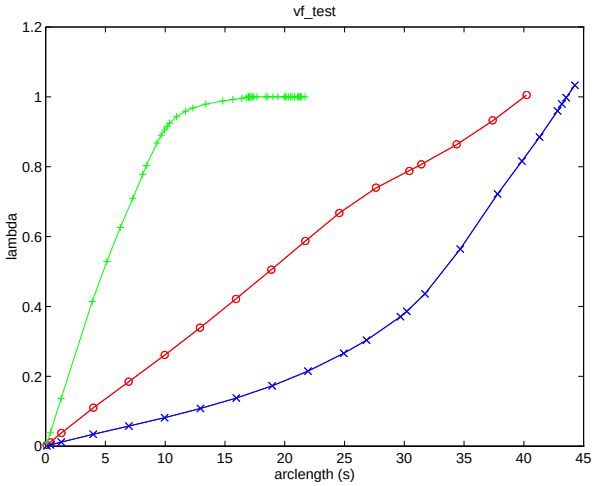

Figure 6: vf_test: λ vs s

The first and second columns of Table 1 list the names and types of a sampling of circuits that exhibit problems

with conventional methods. The circuits range from active filters (dlopata1, heideh), mixed analog-digital circuits involving sigma-delta ADCs, filters, phase mixers, control and division circuitry (test9, vf_test, rabb-xare, addas, dctl_t) to digital blocks and SRAMs (s1423, goh). All circuits except s1423 were obtained from AT&T Microelectronics; s1423 is an ISCAS benchmark circuit which exhibited convergence difficulties with one of AT&T's in-house MOS models.

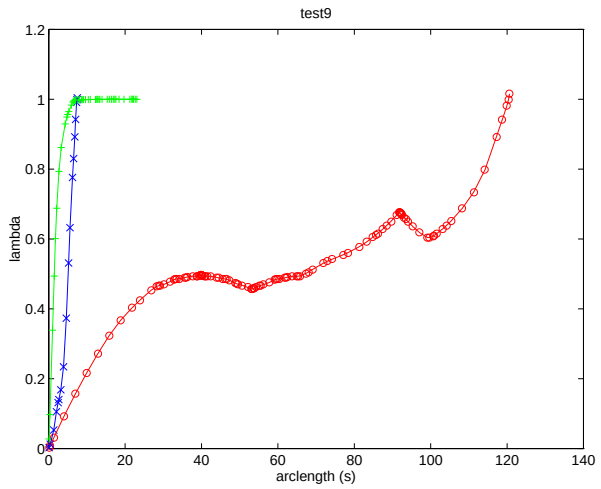

Figure 7: test9: λ vs s

The third column lists the number of MOS devices in the circuits, which range from small (127 MOSFETs) to relatively large (8489 MOSFETs) in size. The fourth column lists the CPU time (on a Sun SPARCstation 2 with 96MB of memory) required by the BLHOM homotopy to obtain an operating point of the circuit. The fifth column lists the CPU time for conventional techniques to announce *failure* – this is helpful as a lower bound on the time wasted by a designer trying to obtain a solution of the circuit.

It can be seen that in most cases, it takes BLHOM considerably less time to obtain the DC operating point of the circuit than it takes for conventional methods to give up. It should be noted however that for circuits on which the Newton-Raphson method[4] succeeds, it is a factor of 2–3 faster than the BLHOM homotopy on the average. The main impact of BLHOM stems from its ability to solve many circuits that cannot be solved by other methods. This has led to large savings in design time, for it was not unusual for several days to be spent in obtaining operating points of "tough" circuits[5].

Figures 6 and 7 provide a graphical representation of the progress of the BLHOM homotopy for three of the above circuits. The horizontal axis represents the arclength s[6] of the continuation of the $n+1$-dimensional solution curve generated by the homotopy solver. Roughly speaking, it is a measure of computation time for a given circuit[7]. On the vertical axis, the value of the continuation parameter λ is plotted. This is a measure of the progress the algorithm has made; success is indicated by the track's reaching $\lambda = 1$. The λ axis represents λ_2, λ_1 or λ_3 depending on the marker on the plot. The plot marked with $\circ$ corresponds to the first phase of BLHOM where (λ_1, λ_2) changes from $(0,0)$ to $(0,1)$, i.e., $\lambda = \lambda_2$ is varied by the continuation algorithm while λ_1 is kept constant at 0. The second phase, where $\lambda = \lambda_1$ is varied while λ_2 is kept constant at 1, is depicted by the plot with the $\times$ markers. The $+$ plot depicts the final phase, the transition from the BLHOM model to the accurate model controlled by $\lambda = \lambda_3$. The solution of the circuit with the accurate in-house model is found when this track reaches 1 on the λ axis.

The three tracks in Figure 6 are for the vf_test circuit. Both $\circ$ and $\times$ tracks (phases 1 and 2) proceed monotonically and with relatively few points from $\lambda = 0$ and $\lambda = 1$, indicating that the circuit is not particularly challenging for the BLHOM homotopy. The $+$ track (phase 3) shows fast progress initially, indicating very little change from the solution obtained with the BLHOM model; the progress slows as it approaches $\lambda_3 = 1$, indicating that the solution is changing at the last stages of the substitution of BLHOM by the accurate in-house model. This is typical of circuits in which some node voltages depend strongly on the second-order details of the MOS model being used – for example near-floating nodes whose voltages are primarily determined by the g_{ds} of MOS devices connected to them.

More interesting behaviour is observed in Figure 7 for the test9 circuit. The first phase marked by $\circ$ is seen to be non-monotonic; it displays two pairs of turning points at which λ_2 changes from increasing to decreasing or vice-versa. Circuits that display such turning points often fail with conventional methods.

Despite its success, BLHOM has not been successful in solving *every* circuit it has encountered efficiently. A few circuits (fewer than 5 out of thousands) have been found that are very slow to solve using the technique. We have developed a probabilistic technique that makes large perturbations to the homotopy to circumvent the inefficiency, which is caused by circuit topologies involving parallel bistable structures. Occurences of this phenomenon are rare, however; the examples shown above are representative of the typical performance of BLHOM on circuits encountered in industry. Testimony to the reliability of the technique are methodology changes that have occurred within some design groups [10] within AT&T – time-consuming worst-case testing of large designs are now relegated to automated scripts without user intervention.

5 Conclusion

A homotopy technique for finding a DC operating point of large-scale MOS circuits has been presented. The method uses arclength continuation together with a new two-phase embedding of λ into the circuit equations. The technique converges reliably on virtually all practical circuits and has led to significant savings in design time and effort within AT&T.

[4] One of a number of conventional techniques that are traditionally used.

[5] Using manual initialization, splitting circuits into smaller parts and other ad-hoc techniques.

[6] s is the distance travelled from $(\bar{x}_0, \lambda = 0)$ to $(\bar{x}, \lambda)$ along the continuation track.

[7] Note: similar values of s do not correspond to similar computation times across different circuits.

Acknowledgments

The authors are indebted to Bruce McNeill and Dave Rich for detailed feedback on initial versions of BLHOM and their close co-operation during the course of this work. Doug Lopata, Bob Walden, Scott Fetterman, Jeff Sonntag, Bernie Morris, Haideh Khorramabadi and Goh Komoriya contributed circuits that were instrumental in testing and refining the new homotopy. Sani Nassif provided useful insights into the properties of different embeddings and helped incorporate the algorithm into a production circuit simulator. Colin McAndrew provided invaluable help in integrating BLHOM with AT&T's proprietary MOS models. Karti Mayaram, David Lee, Jeff Hantgan, Bijan Bhattacharyya, Richard Booth and Mike Toth made important contributions to this work. Discussions with Layne Watson, Peter Feldmann, Roland Freund, Ljiljana Trajković, Mike Green and Jeff Lagarias led to valuable insights about the operation of homotopy methods. The assistance and encouragement of Kishore Singhal throughout the progress of this work is greatly appreciated.

References

[1] E.L. Allgower and K. Georg. *Numerical Continuation Methods.* Springer-Verlag, New York, 1990.

[2] 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*, 13:281–310, 1987.

[3] L.O. Chua and A. Ushida. A switching-parameter algorithm for finding multiple solutions of nonlinear resistive circuits. *Int. J. Ckt.Th. Appl.*, 4:215–239, 1976.

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

[5] L. Vandenberghe and J. Vandewalle. A globally convergent algorithm for solving a broad class of nonlinear resistive circuits. *Proc. IEEE Int. Symp. Ckts. Sys.*, pages 403–406, 1990.

[6] R.C. Melville, L. Trajković, S.-C. Fang and L.T. Watson. Artificial Parameter Homotopy Methods for the DC Operating Point Problem. *IEEE Trans. CAD*, 12(6):861–877, June 1993.

[7] L. 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. Ckts. Sys.*, pages 914–917, May 1990.

[8] C.C. McAndrew and B.K. Bhattacharyya. A single-piece C^{∞}-continuous MOSFET model including subthreshold conduction. Technical Report 52864-900901-01TM, AT&T Bell Laboratories.

[9] Y.P. Tsividis. *Operation and modeling of the MOS transistor.* McGraw-Hill, New York, 1987.

[10] B. McNeill, D. Lopata, D. Rich. Personal communications, 1995.