Design of Robust Global Power and Ground Networks

S. Boyd
E.E. Department
Stanford University

L. Vandenberghe
E.E. Department
UCLA

A. El Gamal
E.E. Department
Stanford University

S. Yun
E.E. Department
Stanford University

ABSTRACT

We consider the problem of determining optimal wire widths for a power or ground network, subject to limits on wire widths, voltage drops, total wire area, current density, and power dissipation. To account for the variation of the current demand, we model it as a random vector with known statistics, possibly including correlation between subsystem currents. Other researchers have shown that when the variation in the current is not taken into account, the optimal network topology is a tree. A tree topology is, however, almost never used in practice, because it is not robust with respect to variations in the block currents. We show that when the current variation is taken into account, the optimal network is usually not a tree.

We formulate a heuristic method based on minimizing a linear combination of total average power and total wire area. We show that this results in designs that obey the reliability constraints, occupy small area, and most importantly are robust against variations in block currents. The problem can be formulated as a nonlinear convex optimization problem that can be globally solved very efficiently.

Keywords

Power and ground network design. Interconnect sizing. Convex optimization.

1. INTRODUCTION

A system-on-chip typically consists of several predesigned hard and soft blocks. The design of such a system involves the hierarchical placement and routing of the blocks to satisfy several delay, power, and area constraints. A crucial part of this process is the design of the power and ground (P&G) distribution networks. These networks are typically designed hierarchically, first at the block level, then at the chip or global level. The P&G networks for the blocks are either completely specified, in the case of hard blocks, or are designed using a conventional standard cell style P&G router. The global P&G network design involves deciding on a topology to supply power to the blocks and selecting appropriate sizes for the wires to satisfy reliability and IR drop constraints. Because of the very large integration and high performance of a system-on-chip, the global P&G networks can occupy a significant fraction of the available wiring capacity. As a result it is important to attempt to minimize the total area occupied by these networks. The global P&G network design problem can be formulated as a constrained optimization problem: we want to minimize the total wire area, subject to bounds on the peak voltages (IR drop), and bounds on the current densities (to limit electromigration). However, solving this optimization problem is very difficult, because the block currents are not exactly known and time-varying.

The existing literature [5, 4, 6, 2, 3] addresses a simplified scenario. It is assumed that each block consumes a constant current and that the network is modeled by a resistive network. Under these assumptions, the minimum area network subject to current density and IR drop constraints is a tree, as shown by Erhard and Johannes [5]. A tree topology is, however, almost never used in practice. The main reason is that it is not robust with respect to variations in the block currents caused either by underestimating the constant current values or, more importantly, by the inevitable variations of the currents in time. Such variations can cause much higher than expected IR drops. Of course it is possible to alleviate this problem by overdesigning the tree for worst case or close to worst case currents, but this would result in significant waste of wiring area.

In this paper we propose a new formulation for the global P&G network design problem, which is both tractable and provides designs that are robust against current variations. We model the network as a resisitive circuit, and take into account current variations by modelling the block currents as random variables. We assume that the first and second moments of the block currents, as well as the correlations between the currents, are known. Such information can be readily obtained from measurements of the block currents, simulations of the blocks, or block level static timing analysis. We minimize the average power dissipated in the P&G network subject to an area constraint. This at first glance appears as an unmotivated objective function, but, as we shall see, results in designs that obey the reliability constraints, occupy small area, and most importantly are robust against variations in block currents. Our formulation yields a convex optimization problem that can be optimally and efficiently solved using interior-point methods. Since the solution can be obtained quickly our approach can be

ISPD'01 April 1-4, 2001, Sonoma, California, USA.

used both for the planning phase of the design as well as for the final detailed phase.

In §2 we introduce the circuit model we use for the P&G network. We state the general design problem and review earlier work on the static and deterministic special cases. We demonstrate via a simple example why the tree topology is not robust against current variations. In §3 we introduce our proposed formulation of the P&G network design problem. We present a numerical implementation and computational results in §4.

2. PROBLEM FORMULATION

We focus on the design of ground distribution networks, since the results translate immediately to power supply networks. We model the blocks or subsystems on the IC as current sources. We assume that the current variations are slow compared to the time constants of the distribution network, *i.e.*, we do not include dynamic effects (*i.e.*, capacitive or inductive), which can sometimes be significant. We will, however, take into account current variations by modelling the block currents as random variables.

The ground distribution network that connects the current sources to the external ground pins is modeled as a linear resistive circuit. The resistances are the interconnect wire resistances (possibly including nonzero substrate resistance); we can also include the output resistance of the current sources as part of the interconnect network. We assume the resistor network has n branches (which represent the wire segments of the distribution network), m (non-ground) nodes (which represent both the subsystem and internal nodes), and a ground node (which represents all the external ground pin nodes). The external current flowing into the network at node k is denoted I_k (thus I_k is equal to the subsystem current if node k is a subsystem node, and zero if node k is an internal node of the ground network), and the node voltages are denoted V_k. The current in branch k (with some fixed orientation) is denoted i_k, and the voltage across branch k is denoted v_k. The conductance g_k of the kth branch (wire segment) is proportional to its width w_k and inversely proportional to its length l_k: $g_k = w_k/(\rho l_k)$, where ρ is the sheet resistance of the wire. The current density in wire k is (up to a constant) $j_k = i_k/w_k$.

The relation between external currents and node voltages is given by the node equations $I = GV$, where

$$G = A\,\mathbf{diag}(g_1, \ldots, g_n)A^T = \sum_{k=1}^{n} \frac{w_k}{\rho l_k} a_k a_k^T \qquad (1)$$

is the conductance matrix of the circuit, A is the incidence matrix, and a_k is the kth column of A. We will sometimes write the conductance matrix as $G(w)$ to emphasize its dependence on w. The current density in branch k is given by

$$j_k = \frac{1}{\rho l_k} v_k = \frac{1}{\rho l_k} a_k^T G^{-1}(w) I. \qquad (2)$$

The variables in the design problem are the widths w_k of the interconnect wires. We are interested in minimizing the total wire area $\sum_{k=1}^{n} l_k w_k$ subject to the following constraints:

- A limit on the average current density or the RMS current density. High current densities can cause metal migration and lead to failure of the circuit. Some simple and widely used experimental models for metal migration predict that the lifetime is a decreasing function of the average current density or the RMS current density, so we can guarantee a minimum lifetime by imposing an upper bound on the average or RMS current density.

- Limits on the maximum voltage drops V_k from the ground pins on the blocks or subsystems to the external ground pin.

We can express this as a constrained optimization optimization problem:

Problem 1

$$\begin{array}{ll} \text{minimize} & \sum_{k=1}^{n} l_k w_k \\ \text{subject to} & \mathbf{E}\,|j_k| \le J, \quad k = 1, \ldots, n \\ & V_k \le V_{\max}, \quad k = 1, \ldots, n \\ & w_k \ge 0, \quad k = 1, \ldots, n. \end{array}$$

This problem is not yet fully specified because we have not stated the assumptions on the statistical distribution of the ground currents, which determines the distribution of j_k and V_k.

The existing literature on power and ground network design addresses the *deterministic* version of the design problem, *i.e.*, it is assumed that the input currents I are constant and known. If we make this assumption, the design problem can be reformulated as follows.

Problem 2 Deterministic formulation

$$\begin{array}{ll} \text{minimize} & \sum_{k=1}^{n} l_k w_k \\ \text{subject to} & V_k \le V_{\max}, \quad k = 1, \ldots, n \\ & w_k \ge 0, \quad k = 1, \ldots, n \end{array}$$

In general, this is not a convex optimization problem, so finding the global optimum is hard. However, a local minimum can be efficiently computed using standard nonlinear optimization techniques (see Chowdhury and Breuer[2]). Erhard and Johannes [5, 4, 6] also show that the optimal solution of problem 2 must have a tree topology (*i.e.*, if we delete the branches for which the optimal width $w_k = 0$, we obtain a tree network).

This paper is motivated by an important shortcoming of the deterministic formulation (problem 2). It assumes that the currents I are constant and given. In practice, the values of the currents are not known exactly and vary with time, so in order to apply the techniques of the previous section, we have to consider average or worst-case values of the currents I. This poses two problems:

- *Robustness with respect to current variations.* Using the average values of the supply currents may result in a solution where the peak value of the node voltages exceeds the maximum allowable value $V_{\max}$.

- *Inefficiency with respect to area.* We can guarantee a bound on the peak value of the node voltages by designing the network for the worst-case values of the input currents. However, this is wasteful in terms of area.

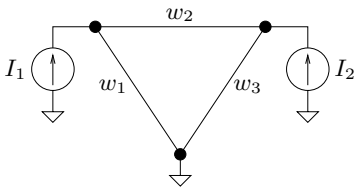

Figure 1: Network with two nodes and three links. The links from nodes 1 and 2 to the ground have length $l_1 = l_3 = 1$, and the link between nodes 1 and 2 has length $l_2 = 0.5$.

Figure 1 shows a small example that illustrates this point. We assume that the input current vector I is random, with two possible values: we either have $I_1 = 1$, $I_2 = 50$, or $I_1 = 50$, $I_2 = 1$, with equal probability. We take $\rho = 1$, so the branch conductances are $g_k = w_k/l_k$. We impose the following limits on the current densities and voltage drops:

$$\mathbf{E}\,|j_k| \leq 1, \quad V_k \leq 1.$$

We will show that solving the deterministic problem 2 with worst case values for I, yields suboptimal solutions for the stochastic problem.

We know from the results in [5, 4, 6] that the optimal solution of problem 2 is a tree (*i.e.*, one of the three widths is zero at the optimum). By symmetry we only have to consider two of the three possible trees in the network.

Let us first consider the tree consisting of links 1 and 3 (*i.e.*, we take $w_2 = 0$). In order to satisfy the voltage drop constraints, we need $w_1 \geq 50$ and $w_3 \geq 50$. With this choice, the resistances of both links are less than or equal to $1/50$, so the voltages never exceed 1. If we choose $w_1 = w_2 = 50$, the average current density in both branches is 0.51, which is less than the maximum allowable value, so $w_1 = w_2 = 50$ is the optimal solution (assuming $w_3 = 0$). The resulting area is equal to 100.

The second tree consists of links 1 and 2, *i.e.*, we set $w_3 = 0$. It is clear that the maximum node voltage will occur at node 2, when $I_1 = 1$ and $I_2 = 50$, so we have to consider the following problem:

$$\begin{array}{ll} \text{minimize} & w_1 + 0.5w_2 \\ \text{subject to} & \max V_2 = 51/w_1 + 25/w_2 \leq 1 \\ & \mathbf{E}\,|j_1| = 51/w_1 \leq 1 \\ & \mathbf{E}\,|j_2| = 0.5(50/w_2 + 1/w_2) \leq 1. \end{array}$$

The optimal values for the widths are $w_1 = 76.2$, $w_2 = 75.5$ (which makes the voltage drop constraint tight at the second node). The average current densities are less than 0.67, and the area is 114.

We conclude that the optimal solution with a tree topology is $w_1 = w_2 = 50$, $w_3 = 0$, with an area equal to 100.

However, this solution is suboptimal for problem 1, and we can achieve a smaller area by using a non-tree topology. For example, choosing $w_1 = w_3 = 31$, $w_2 = 26$ yields a solution with maximum node voltage 1, average current densities $\mathbf{E}\,|j_k(t)| \leq 0.82$, and area 75.

We conclude that if we take current variation into account, the topology of the optimal solution of problem 1 is not necessarily a tree, and the methods described above for problem 2 yield solutions that are either nonrobust against current variations or conservative.

3. DESIGN VIA POWER-AREA TRADEOFF

The contribution of the paper is a heuristic method for obtaining good suboptimal solutions for problem 1. Our method is based on minimizing a weighted sum of average power dissipation and total wire area. Specifically, we consider the following problem.

Problem 3 Average power-area tradeoff problem

$$\begin{array}{ll} \text{minimize} & \mathbf{E}\, I^T G(w)^{-1} I + \mu \sum_{k=1}^n l_k w_k \\ \text{subject to} & w \geq 0. \end{array}$$

Here μ is a positive constant, and the first term in the objective,

$$\mathbf{E}\, IG(w)^{-1} I = \mathbf{E} \sum_{k=1}^m I_k V_k,$$

is the expected value of the power dissipated in the ground network. The objective is a weighted sum of the expected power dissipation and the total area, with μ controlling the relative importance of both terms. The average power can be expressed in terms of w as

$$\mathbf{E}\, I^T G(w)^{-1} I = \mathbf{Tr}\, G(w)^{-1} \Gamma,$$

where $\mathbf{Tr}\, X = X_{11} + \cdots + X_{nn}$ denotes the trace of the (square) matrix X, and $\Gamma = \mathbf{E}\, II^T$ is the second moment (or correlation matrix) of I. It can be shown that the function $\mathbf{Tr}\, \Gamma G(w)^{-1}$ is a differentiable convex function of w, and therefore problem 3 is a convex optimization problem so it can be globally solved very efficiently. Efficient methods for minimizing this function subject to nonnegativity constraints on w are discussed in §4.

We now show the resulting design has the property that each wire is either zero width (meaning that it is not used) or has a constant RMS current density. This implies that we indirectly solve the problem of minimizing the total area subject to a limit on RMS current density.

3.1 Equal RMS current density property

We first examine the optimality conditions for problem 3. To simplify notation, we write the objective as $\mathbf{E}\, P(w) + \mu A(w)$ where

$$P(w) = I^T G(w)^{-1} I, \quad A(w) = \sum_{k=1}^n l_k w_k.$$

The necessary and sufficient conditions for $w \geq 0$ to be optimal are:

$$\frac{\partial}{\partial w_k} \left(\mathbf{E}\, P(w) + \mu A(w) \right) = 0 \tag{3}$$

for each k with $w_k > 0$, and

$$\frac{\partial}{\partial w_k} \left(\mathbf{E}\, P(w) + \mu A(w) \right) \geq 0 \tag{4}$$

for each k with $w_k = 0$.

The partial derivative of A is simply $\partial A/\partial w_k = l_k$. To find the partial derivatives of $\mathbf{E}\, P$ we use the fact that

$$\frac{\partial}{\partial x} Y^{-1} = -Y^{-1} \frac{\partial Y}{\partial x} Y^{-1}$$

where Y is a symmetric matrix that depends on $x \in \mathbf{R}$. Therefore,

$$\begin{aligned}\frac{\partial \mathbf{E}\, P}{\partial w_k} &= \mathbf{E}\,\frac{\partial}{\partial w_k} I^T G^{-1}(w) I \\ &= -\frac{1}{\rho l_k}\,\mathbf{E}(I^T G^{-1}(w) a_k a_k^T G^{-1}(w) I) \\ &= -\frac{1}{\rho l_k}\,\mathbf{E}(V^T a_k a_k^T V) = -\frac{\mathbf{E}\, v_k^2}{\rho l_k}.\end{aligned}$$

The condition (3) can therefore be expressed as

$$-\frac{\mathbf{E}\, v_k^2}{\rho l_k} + \mu l_k = 0,$$

i.e., $\mathrm{rms}(v_k) = l_k\sqrt{\rho\mu}$, which implies

$$\mathrm{rms}(j_k) = \frac{\mathrm{rms}(i_k)}{w_k} = \frac{g_k\,\mathrm{rms}(v_k)}{w_k} = \frac{\mathrm{rms}(v_k)}{\rho l_k} = \sqrt{\frac{\mu}{\rho}}.$$

Thus, for every wire that has nonzero width, the RMS current density is exactly equal to $\sqrt{\mu/\rho}$.

The condition (4) can be expressed as

$$\mathrm{rms}(v_k) \le l_k\sqrt{\mu\rho}.$$

This condition is satisfied by every wire whose optimal width is zero.

As a second consequence of the optimality conditions, we note that the solution of problem 3 changes with μ in a very simple way: if $\lambda > 0$ and w solve problem 3 for μ, then λw solves the problem for μ/λ^2. In particular, the topology of the optimal solution is independent of μ.

3.2 Design method

We propose solving problem 3 as a technique for designing a ground network in the presence of current variations, metal migration, and voltage drop constraints, for the following reasons.

- As we have seen, at the optimum the RMS values of the current densities in all branches are equal, and

  $$\mathrm{rms}(j_k) = \sqrt{\frac{\mu}{\rho}} \text{ if } w_k > 0.$$

  By choosing $\mu = \rho J^2$, we can guarantee that the RMS current densities will be equal to J.

- By penalizing the average dissipated power, we indirectly reduce voltage drop. We will illustrate this effect with an example in §3.3.

- The optimal topology of the solution of problem 3 is independent of μ. Once we know the solution for a given maximum RMS current density J, we can obtain the solution for a different value of J very easily, by simply scaling all widths.

  Suppose for example that we design the network for a certain RMS current density, and then find out by simulation that the maximum voltage drop is too high. By scaling all the widths by the same factor, we can reduce the voltage drop to an acceptable level, while maintaining the property that the RMS current densities in all links are equal.

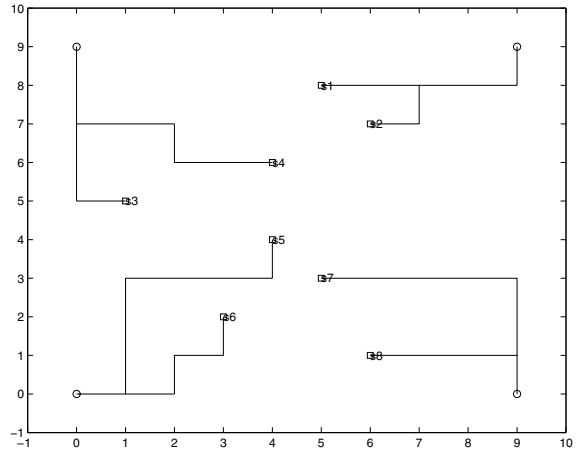

Figure 2: 10×10 mesh with 4 ground pins (circles) and 10 sources (squares), and the optimal topology if the currents are constant.

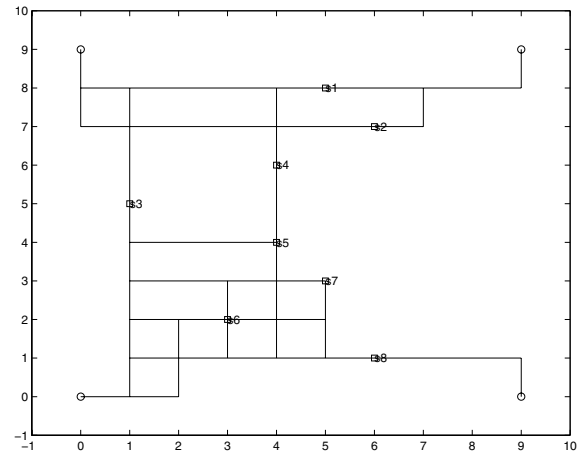

Figure 3: Optimal topology for a stochastic current model.

If we apply the method to the example of §2, (with $\mu = 1$), we find optimal widths $w_1 = w_3 = 26.3$, $w_2 = 17.9$, for which we have $\max V_1 = \max V_2 = 1.22$. To satisfy the maximum voltage drop constraint of 1, we scale the widths by 1.22. This yields $w_1 = w_3 = 32$, $w_2 = 21.8$. This solution satisfies $\mathbf{E}\,|j_k| \le 0.82$, $V_k \le 1.0$ and its area is 75.1. Compared to the tree solutions of §2 we see that this method results in a feasible solution with a smaller area.

3.3 Example

Figure 2 shows the geometry of the problem. The network is a 10×10-mesh, with a node at each point of the mesh, and a link between each node and its nearest neighbors. Each link has unit length, plus a small randomly generated perturbation to break the symmetry. There are 8 input current sources, at the positions indicated by squares. The external ground nodes are at the four corners.

Figure 2 shows the optimal topology for deterministic currents. The solution is a tree (if we identify the four corners

with the ground node), which can be constructed in a very simple way: each current source is connected to the nearest ground node.

The solution in figure 3 is the solution for stochastic currents, and was obtained by solving problem 3. We assume that the current vector I can take three values

$$\begin{aligned} I^{(1)} &= (10,\ 10,\ 0,\ 0,\ 10,\ 10,\ 0,\ 0) \\ I^{(2)} &= (0,\ 0,\ 0,\ 0,\ 0,\ 0,\ 10,\ 10) \\ I^{(3)} &= (0,\ 0,\ 10,\ 10,\ 0,\ 0,\ 0,\ 0) \end{aligned}$$

each with probability 1/3.

Let us compare the two topologies assuming the currents switch periodically between $y^{(1)}$, $y^{(2)}$ and $y^{(3)}$. We use a value $\rho = 1$ for the sheet resistance, and assume the maximum allowable RMS current density is one.

- If we size the links in figure 2 so that the RMS current density in each link is one, we obtain a network with total area 448. If we apply current $y^{(1)}$, then the maximum node voltage is equal to 7.7 (at source no. 5). If we apply current $y^{(2)}$, the maximum node voltage is 6.7 (at source no. 7). For current $y^{(3)}$, the maximum node voltage is 6.8 (at source no. 4).
- Figure 3 is optimized for random currents. The RMS current density in all branches is one. The area is 347. If we apply current $y^{(1)}$, the maximum node voltage is equal to 5.3 (at source no. 5). If we apply current $y^{(2)}$ the maximum node voltage is 4.2 (at source no. 7). For current $y^{(3)}$ the maximum node voltage is 4.1 (at source no. 4).

We see that by using a non-tree topology, we obtain a solution with a smaller area and smaller peak voltage drops, for the same values of the RMS current densities.

4. COMPUTATIONAL RESULTS

In this section we discuss numerical methods for solving problem 3. The most straightforward method is to cast the problem in terms of a more general standard convex optimization problem for which software is readily available. This can be done in (at least) two ways: using semidefinite programming (SDP), or second-order cone programming (SOCP) [1]. The drawback of these methods is that general purpose software does not exploit the structure of the problem (*e.g.*, the sparsity of $G(w)$ and the matrices $a_k a_k^T$). The method is therefore limited to small problems (*e.g.*, up to several hundred nodes and links).

We now describe a method that does exploit the problem structure, and is effective for large problems, even with several thousand nodes or branches.

4.1 Barrier method with pruning

The objective function in problem 3 is a smooth convex function; the only constraints are the inequalities $w_k \geq 0$. These constraints can be handled by augmenting the objective function with a logarithmic barrier function. We form the function

$$\phi(w) = \mathbf{Tr}\,\Gamma G(w)^{-1} + \mu l^T w - \beta \sum_{i=1}^{n} \log w_i, \qquad (5)$$

where $\beta > 0$ is a parameter [7]. The function (5) is defined for all $w > 0$, and it is smooth and convex on its domain. It can be shown that the minimizer of (5) is suboptimal for (3) with an accuracy of at least $n\beta$ (see, *e.g.*, [1]). In order to solve (3) to an accuracy ϵ, one can therefore set $\beta = \epsilon/n$, and solve the unconstrained convex optimization problem (5). An alternative, which is often more efficient, is to solve (5) repeatedly for a sequence of decreasing β, until the desired accuracy is reached.

Any unconstrained minimization method can be applied to (5), *e.g.*, Newton's method, conjugate gradients, or coordinate descent. The choice involves a tradeoff between speed of convergence and amount of work per iteration. For example, Newton's method converges in few iterations, but each iteration involves the evaluation of the gradient and Hessian of (5), *i.e.*,

$$\frac{\partial \phi(w)}{\partial w_k} = -\frac{1}{\rho l_k} a_k^T G(w)^{-1} \Gamma G(w)^{-1} a_k + \mu l_k - \beta / w_k$$

and

$$\frac{\partial^2 \phi(w)}{\partial w_k \partial w_j} = \frac{2}{\rho^2 l_k l_j} a_k^T G(w)^{-1} a_j a_j^T G(w)^{-1} \Gamma G(w)^{-1} a_k,$$

$$\frac{\partial^2 \phi(w)}{\partial w_k^2} = \frac{2}{\rho^2 l_k^2} a_k^T G(w)^{-1} a_k a_k^T G(w)^{-1} \Gamma G(w)^{-1} a_k + \beta / w_k^2$$

for $k \neq j$. The conjugate gradient method requires more iterations, but each iteration is cheap, and involves only gradient evaluations. In our implementation (described below, and used to carry out the numerical examples) we used Newton's method, but for extremely large problems, say with ten or hundreds of thousands of wires and nodes, one could use a conjugate-gradient or coordinate descent method to minimize (5).

Recall that without loss of generality we can set the parameter μ to have any particular value, since all other points of the optimal trade-off curve can be found by scaling the solution found for the particular value. As our initial guess of wire widths, we take all wires to have unit width: we assign $w = \mathbf{1}$, where $\mathbf{1}$ is the vector with all components one. We can then choose the value of μ that makes the two terms in the objective, *i.e.*, the average power $\mathbf{Tr}\,\Gamma G(\mathbf{1})^{-1}$ and the scaled area $\mu l^T \mathbf{1}$, equal. This is achieved with $\mu = \mathbf{Tr}\,\Gamma G(\mathbf{1})^{-1} / l^T \mathbf{1}$.

Several general guidelines for choosing the initial value of β have been suggested, *e.g.*, in [7]. We have found that the value $\beta = 0.05(\mathbf{Tr}\,\Gamma G(\mathbf{1})^{-1})/n$ works well.

The barrier method consists of the following: Newton's method is used to minimize the function (5) for a fixed value of β. Then, the value of β is decreased by a factor of approximately 10–50, and Newton's method is used again to minimize (5) for the new value of β. This continues until $n\beta$ is smaller than the required tolerance. We used a relative tolerance of 1% as the stopping criterion.

As suggested by the small example described in §3.3, the optimal solution often has many wire widths equal to zero, and many isolated nodes. This can be exploited to greatly improve the efficiency of the algorithm. After only a few iterations, one can often clearly identify a large number of wires whose widths are converging to zero. As a heuristic to speed up subsequent iterations, we set these wire widths to zero, remove any nodes that become isolated, and then continue the barrier method. This periodic pruning of the wires and nodes greatly improves the efficiency of the algorithm, without affecting the quality of the solution. In

Name	Wires	Nodes	CPU time (s)
mesh15	420	221	4
mesh25	1200	621	105
mesh35	2380	1221	816
mesh45	3960	2021	3800
rnng100	242	96	1
rnng200	494	196	8
rnng500	1216	496	90
rnng1000	2415	996	860

Table 1: Summary of computational results.

our implementation, we pruned wires and nodes after each Newton step in the barrier method.

4.2 Computational results

We tested our implementation on two randomly generated problem classes. The results are summarized in Table 1. In the first set of examples (labeled as mesh$k$) we use a mesh topology of $k \times k$ nodes, where each node is connected to its nearest neighbors. The ground nodes are at the four corners of the mesh. Current sources are placed at k randomly selected nodes. For the second class (labeled as rnng$k$) we place k nodes at random positions in a square, and connect each node to its four nearest neighbors in Manhattan distance. Four nodes are randomly selected as the ground nodes. Current sources are placed at $k/2$ randomly selected nodes. In all problems, we use a randomly generated current correlation matrix Γ of rank 3.

In all of these examples, the final solution had a sparsity of about 80%, *i.e.*, 80% of the wires were pruned. The CPU times reported are for a C++ implementation, on a 500MHz Pentium II running Linux. For problems with around 1200 variables (wires) the optimal solutions were found in a few minutes; for the largest problem, which had around 4000 variables, the CPU time was around an hour.

We also tested our implementation on an example taken from a real circuit design. The circuit consists of 45 subsystems, each of which had a fixed internal ground network and a number of nodes to connect to the external, chip level ground network. We used (potential) wires between all adjacent nodes, for a total of 1670 (potential) wires and 745 nodes. Eight ground pins were fixed around the perimeter of the chip.

The correlation matrix of the subsystem supply currents was estimated from circuit simulation traces. First, we estimated Γ by the average,

$$\Gamma = \frac{1}{T}\int_0^T i(t)i(t)^T\,dt,$$

where $i(t)$, for $0 \leq t \leq T$, was obtained from a simulation of the chip in typical operation. We then formed a rank 11 approximation via an eigenvalue decomposition.

The resulting topology is shown in figure 4. Of the original 1670 wires, only around 150 remain. This problem was solved in around 400 seconds.

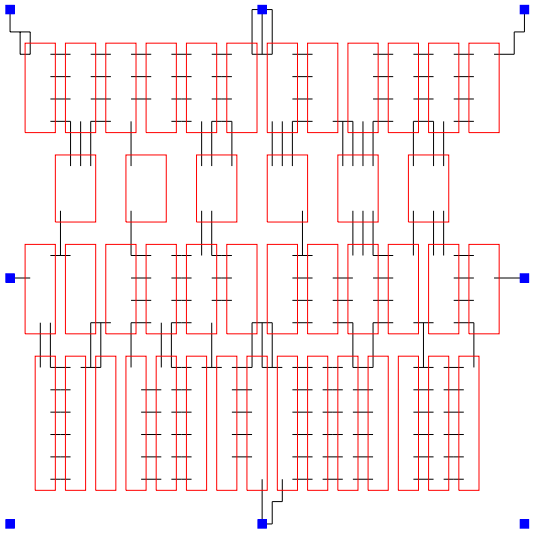

Figure 4: Circuit example. The optimal ground network wires are shown as lines, the subsystems are shown as lighter rectangles, and the dark squares are the external ground pins. Each subsystem has its own (fixed) ground network connecting its ground pins.

Acknowledgments

We thank Maryam Fazel and Nejat Cingi for help in coding an early implementation of the method. We are grateful to Res Saleh from Simplex Solutions for providing the data for the circuit example. This research was partially supported by NSF under grant CCR-9714002.

5. REFERENCES

[1] S. Boyd and L. Vandenberghe. Introduction to convex optimization with engineering applications. Course Notes, 1997. http://www.stanford.edu/class/ee364/.

[2] S. Chowdhury and M. Breuer. Minimal-area design of power/ground nets having graph topologies. *IEEE Transactions on Circuits and Systems*, 34(12):1441–1451, Dec. 1987.

[3] S. Chowdhury and M. A. Breuer. Optimum design of IC power/ground nets subject to reliability constraints. *IEEE Transactions on Computer-Aided Design*, 7:787–796, 1988.

[4] K. H. Erhard and F. M. Johannes. Area minimization of IC power/ground nets by topology optimization. In A. Halaas and P. B. Denyer, editors, *VLSI 91*, pages 119–126, Edinburgh, 1992.

[5] K. H. Erhard and F. M. Johannes. Power/ground networks in VLSI: are general graphs better than trees? *INTEGRATION, the VLSI journal*, 14:91–109, 1992.

[6] K.-H. Erhard, F. M. Johannes, and R. Dachauer. Topology optimization techniques for power/ground networks in VLSI. In *Proceedings of the European Design Conference*, pages 362–367, Hamburg, Germany, 1992.

[7] A. Fiacco and G. McCormick. *Nonlinear programming: sequential unconstrained minimization techniques*. Wiley, 1968. Reprinted 1990 in the SIAM Classics in Applied Mathematics series.