

Maximum Likelihood Haplotyping for General Pedigrees

Ma'ayan Fishelson, Nickolay Dovgolevsky, and Dan Geiger

Computer Science Department
Technion, Haifa, Israel

Running Title: Maximum Likelihood Haplotyping

Corresponding Author: Ma'ayan Fishelson, Computer Science Department, Technion,
Technion City, Haifa 32000, Israel. Tel: +972 4 829 4895, Fax: +972 4 829 3900. Email:
fmaayan@cs.technion.ac.il

Keywords: Haplotyping, Linkage Analysis, Pedigree, SUPERLINK.

Abstract

Haplotype data is valuable in mapping disease-susceptibility genes, especially in the study of complex diseases. We present algorithms for inferring a most likely haplotype configuration for general pedigrees, implemented in the newest version of the genetic linkage analysis system SUPERLINK. In SUPERLINK, genetic linkage analysis problems are represented internally using Bayesian networks. The use of Bayesian networks enables efficient maximum likelihood haplotyping for more complex pedigrees than was previously possible. Furthermore, to support efficient haplotyping in larger pedigrees, we have also incorporated a novel algorithm for determining a better elimination order for the variables of the Bayesian network. The presented optimization algorithm also improves likelihood computations. We present experimental results for the new algorithms on a variety of real and semi-artificial data sets, and use our software to evaluate MCMC approximations for haplotyping.

1 INTRODUCTION

Haplotype data is valuable in mapping disease-related genes, especially in the study of complex diseases (Oehlman et al., 1993; Litt et al., 1994). Such haplotype data defines the nearest flanking recombination events and consequently the smallest interval containing a disease gene. This allows tracing disease genes more easily and cheaply. Other uses of haplotyping include family based statistical tests such as TDT (*Transmission Disequilibrium Test*) which require haplotype data as input (Ewens and Spielman, 1995), or as a means for detecting genotyping errors, which are usually expressed as an excess of recombination events (Lin and Speed, 1997).

The input data for a haplotyping problem can be divided into two categories: pedigree genotype data and population genotype data. The haplotyping problem is to infer the two haplotypes of each individual from the measured unordered genotypes. Haplotype information from population data is often reconstructed using some evolutionary model and is usually applied to data with a dense map of markers (e.g., Stephens et al., 2001; Gusfield, 2002; Greenspan and Geiger, 2003). On the other hand, haplotype information from pedigrees is reconstructed using the information that can be inferred on each individual from his relatives' genotypes, and can be used to reconstruct haplotypes from either dense or widely spaced marker data.

The haplotyping problem can be defined via maximizing a suitable likelihood function or via a combinatorial optimization problem. A common combinatorial approach, called the *Minimum Recombinant Haplotype Configuration (MHRC)* problem, is to seek those haplotype configurations that minimize the total number of recombination events observed in the pedigree. Another common combinatorial approach is to seek those haplotype configurations that show no recombination events. Approaches to solve such combinatorial optimization problems include rule-based systems (Wisjman, 1987; Qian and Beckmann, 2002; Li and Jiang, 2003a), graph-theoretic approaches (Gusfield, 2002), dynamic programming (Li and Jiang, 2003b), or linear programming (Li and Jiang, 2004). These approaches are most appropriate when the expected number of recombination

events is small.

The statistical approach for haplotyping by maximizing a suitable likelihood function has been pursued quite extensively (Sobel et al., 1995; Lin and Speed, 1997) and implemented in programs that perform exact computations, such as GENEHUNTER (Kruglyak et al., 1996; Kruglyak and Lander, 1998), and ALLEGRO (Gudbjartsson et al., 2000), as well as in programs that perform approximate computations, such as SIMWALK2 (Sobel and Lange, 1996), and MERLIN (Abecasis et al., 2002). All these methods take into account intermarker recombination fractions or intermarker genetic distances. The objective of these algorithms is to find one or several haplotype configurations of maximal probability given the observed data on the pedigree.

In this paper we focus on improving exact approaches for generating a maximum likelihood haplotype configuration for larger pedigrees. We present a haplotyping algorithm which we have incorporated into the freely available newest version of SUPERLINK (v1.4), reported herein. SUPERLINK uses Bayesian networks as the internal representation of pedigrees, which allows one to handle a wide variety of linkage problems (Fishelson and Geiger, 2002). In particular, this representation allowed us to naturally implement a maximum likelihood approach for haplotyping. Furthermore, to support efficient haplotyping on larger pedigrees, we have also incorporated a novel algorithm for determining a better elimination order for the variables of the Bayesian network. This algorithm is especially important when solving linkage problems since the Bayesian networks created for such problems are very large. The presented optimization algorithm also improves likelihood computations. In addition, we have adapted the allele recoding algorithm, presented in (O'connell and Weeks, 1995), for the haplotyping task, achieving further reduction in time and space complexity. We present experimental results for the new algorithms on a variety of real and semi-artificial data sets, and use our software to evaluate MCMC approximations for haplotyping via SIMWALK2 (Sobel et al., 1995; Sobel and Lange, 1996).

2 THE HAPLOTYPING PROBLEM

2.1 PROBLEM DEFINITION

The sequence of alleles at different loci inherited by an individual from one parent is called a *haplotype*, and the two haplotypes of an individual constitute this individual's *genotype*. A *recombination* is said to have occurred between two loci, if an haplotype of an individual contains two alleles that resided in different haplotypes of the individual's parent. The *recombination fraction* θ is the probability that a recombination occurs between two loci. For a comprehensive background on human genetic linkage analysis consult (Ott, 1999).

When genotypes are measured by standard measurement procedures, the result is a list of unordered pairs of alleles, one pair for each locus. The *Maximum Likelihood Haplotype Configuration* problem, consists of finding a joint haplotype configuration for all members of the pedigree which maximizes the probability of the data. The haplotyping problem often does not have a unique solution.

2.2 BAYESIAN NETWORKS

Our model for representing pedigree data is a *Bayesian network*. A Bayesian network is a directed acyclic graph, where each vertex $v = 1, \dots, n$ corresponds to a discrete variable X_v and each directed edge represents conditional dependencies between the variables it connects (Pearl, 1988; Lauritzen, 1996). The distribution of each variable X_v is conditional upon the variables in \mathbf{Pa}_v , which is defined as the set of vertices from which there are edges leading into v in the graph. The joint probability of a full assignment x_1, \dots, x_n to variables X_1, \dots, X_n is the product of these conditional probabilities. In other words,

$$Pr(X_1 = x_1, \dots, X_n = x_n) = \prod_v Pr(X_v = x_v \mid \mathbf{Pa}_v = \mathbf{pa}_v),$$

where \mathbf{pa}_v is the joint assignment $\{x_i | X_i \in \mathbf{Pa}_v\}$ to the variables in \mathbf{Pa}_v . From here on, we will use the notation $Pr(\mathbf{y}|\mathbf{z})$ as an abbreviated form of $Pr(\mathbf{Y} = \mathbf{y} | \mathbf{Z} = \mathbf{z})$ for any sets of variables \mathbf{Y} and \mathbf{Z} . For example, the joint probability could be rewritten as $Pr(X_1, \dots, X_n) = \prod_v Pr(X_v | \mathbf{Pa}_v)$. Note also that we use capital letters for variable names and lowercase letters to denote specific values taken by those variables. Sets of variables are denoted by boldface capital letters, and assignments of values to the variables in these sets are denoted by boldface lower case letters.

For haplotyping, we consider the *Most Probable Explanation (MPE)* problem for Bayesian networks, (e.g., Dechter, 1996). That is, finding an assignment $\mathbf{X} = \mathbf{x}^0$ such that

$$Pr(\mathbf{X} = \mathbf{x}^0, \varepsilon) = \max_{\mathbf{x}} Pr(\mathbf{X} = \mathbf{x}, \varepsilon) = \max_{\mathbf{x}} \prod_v Pr(X_v | \mathbf{Pa}_v, \varepsilon)$$

where $\mathbf{X} = \{X_1, \dots, X_n\}$ is the set of variables in the Bayesian network, and ε denotes a particular assignment of values to some of the variables in \mathbf{X} . The assignment ε is called *evidence*. In the case of haplotyping, the evidence is a partial assignment ε of alleles at some or all loci to people in the pedigree under study.

Another relevant problem for haplotyping is the *Maximum A Posteriori Hypothesis (MAP)* problem (Dechter, 1996), of which *MPE* is a special case. The input to this problem is the same as for the *MPE* problem, with the addition of a set of focus variables $\mathbf{A} = \{A_1, \dots, A_k\}$, $\mathbf{A} \subseteq \mathbf{X}$, for which the most probable assignment, given the evidence, is desired. The MAP problem is to find an assignment $\mathbf{a}^0 = (a_1, \dots, a_k)$, such that

$$Pr(\mathbf{A} = \mathbf{a}^0, \varepsilon) = \max_{\mathbf{a}} \sum_{X_j \in \{\mathbf{X} \setminus \mathbf{A}\}} \prod_v Pr(X_v | \mathbf{Pa}_v, \varepsilon)$$

Note that when $\mathbf{A} = \mathbf{X}$, MPE is identical to MAP. It has been shown that solving the MAP problem is significantly harder than solving the MPE problem or computing the probability of evidence (Park, 2002). Consequently, MPE is often solved instead of MAP, and the most likely assignment of all variables is projected on the focus set of variables \mathbf{A} .

3 METHODS

3.1 HAPLOTYPING IN SUPERLINK

Three types of random variables are used in the representation of pedigrees as Bayesian networks in SUPERLINK: *genetic loci* variables which represent the genotypes of the individuals in the pedigree (two genetic loci variables per individual per locus, one for the paternal allele and one for the maternal allele), *phenotype* variables, and *selector* variables which are auxiliary variables used to represent the gene flow in the pedigree. For example, the paternal selector of individual i at locus j indicates whether the paternal allele of individual i at locus j came from his father's paternal haplotype or from his father's maternal haplotype. Figure 1 presents a fragment of a network that describes parents-child interaction in a simple 3-loci analysis. The genetic loci variables of individual i at locus j are denoted by $G_{i,jp}$ and $G_{i,jm}$. Variables $P_{i,j}$, $S_{i,jp}$, and $S_{i,jm}$ denote the phenotype variable, the paternal selector variable and the maternal selector variable of individual i at locus j , respectively. For more details on the structure of the Bayesian network, consult (Fishelson and Geiger, 2002).

We use the following notation to refer to the different variables: \mathbf{S} for the set of all selector variables, \mathbf{F} for the set of genetic loci variables of individuals with no parents in the pedigree (founders), and \mathbf{N} for the set of genetic loci variables of non-founders. For haplotyping, phenotypes are unordered genotypes of typed individuals, and are included in the evidence ε . The Bayesian network of SUPERLINK represents the joint distribution $Pr(\mathbf{S}, \mathbf{F}, \mathbf{N}, \varepsilon)$ in factored form:

$$Pr(\mathbf{S}, \mathbf{N}, \mathbf{F}, \varepsilon) = Pr(\mathbf{S}, \varepsilon)Pr(\mathbf{F}|\varepsilon)Pr(\mathbf{N}|\mathbf{F}, \mathbf{S}, \varepsilon). \quad (1)$$

A maximum-likelihood haplotype configuration of a pedigree is a maximum-likelihood assignment to all the genetic loci variables, namely a joint assignment $\{\mathbf{N} = \mathbf{n}_0, \mathbf{F} = \mathbf{f}_0\}$ which satisfies:

$$Pr(\mathbf{N} = \mathbf{n}_0, \mathbf{F} = \mathbf{f}_0) = \max_{\mathbf{n}, \mathbf{f}} \sum_{\mathbf{s}} Pr(\mathbf{s}, \mathbf{f}, \mathbf{n}, \varepsilon).$$

Since we are interested in determining the most likely gene flow in addition to the most likely assignment to all the haplotypes, we seek a joint maximum-likelihood assignment to the selector variables and the genetic loci variables of founders, namely a joint assignment $\{\mathbf{S} = \mathbf{s}_0, \mathbf{F} = \mathbf{f}_0\}$ which satisfies:

$$Pr(\mathbf{S} = \mathbf{s}_0, \mathbf{F} = \mathbf{f}_0) = \max_{\mathbf{s}, \mathbf{f}} \sum_{\mathbf{n}} Pr(\mathbf{s}, \mathbf{f}, \mathbf{n}, \varepsilon). \quad (2)$$

The genetic loci variables of non-founders, \mathbf{N} , are a function of the genetic loci variables of founders and the selector variables, which, for every \mathbf{s}, \mathbf{f} , is zero for all values \mathbf{n} except one. Consequently, solving Eq. (2) is equivalent to:

$$Pr(\mathbf{S} = \mathbf{s}_0, \mathbf{F} = \mathbf{f}_0, \mathbf{N} = \mathbf{n}_0) = \max_{\mathbf{s}, \mathbf{f}, \mathbf{n}} Pr(\mathbf{s}, \mathbf{f}, \mathbf{n}, \varepsilon), \quad (3)$$

which is an MPE problem. Thus the MAP problem, defined by Eq. (2), is essentially an MPE problem which can be solved more easily.

Our algorithm for solving Eq. (3) consists of several stages. The first stage is a preprocessing step of value elimination on the graph representation of the pedigree (Fishelson and Geiger, 2002). At this stage inconsistent values for each variable given the evidence are removed, and the values of some of the genetic loci and selector variables can be determined unambiguously from the evidence, namely all values except one are removed. Mendelian inconsistencies are also discovered at this stage. The value elimination performed is based on the well-known observation that the possible genotypes of an individual can be inferred from the genotypes of his relatives (e.g. Lange and Goradia, 1987). When the value of a selector variable is determined by value elimination it implies that the parental origin of the corresponding allele is known, i.e., whether it came from the paternal or maternal haplotype of the parent.

After value elimination, we perform allele recoding. In this stage the genotype lists of untyped individuals are recoded, resulting in a reduction in the number of genotypes that need to be summed over, and hence, in an acceleration of the computations. Our

allele recoding algorithm, which is an adaptation of the ideas presented in (O'connell and Weeks, 1995) to the task of haplotyping, is fully described in the next section.

Finally, haplotyping is done via performing the *Elim-Max* algorithm (Dechter, 1996) on the Bayesian network to determine a maximum-likelihood assignment to the remaining variables. The Elim-Max algorithm, described in the appendix, is a variable elimination algorithm in which variables are eliminated one after another, each time computing the effect of the eliminated variable on the rest of the problem. The order by which variables are eliminated greatly affects both time and space requirements of the computations. In many cases, the memory limitation does not allow solving the problem using variable elimination alone, and hence, variable elimination is combined with conditioning. By conditioning, one means to instantiate some of the variables, perform the rest of the computations for each possible instantiation, and then merge the results. The order of variable elimination and conditioning is determined by a new algorithm described below.

3.2 ALLELE RECODING

When performing likelihood computation or haplotyping, all possible genotype combinations for the individuals in the pedigree need to be iterated over. When using highly polymorphic markers, any person who is untyped at some locus will have a large number of possible genotypes. A possible way to accelerate these computations is to recode alleles and thus reduce the number of possible genotypes that need to be iterated on. Several different methods have been proposed. One method is lumping all alleles that do not appear in the pedigree into a single allele whose population frequency is the sum of frequencies of the lumped alleles (Lange et al., 1988; Schäffer, 1996). A more efficient method, which recodes the paternal and maternal allele lists of each individual separately, has been suggested by O'connell and Weeks (1995), and implemented in VITESSE. The allele recoding algorithm implemented in SUPERLINK is based on the ideas of *set-recoding* and *fuzzy inheritance* defined in VITESSE. These definitions are repeated here for completeness. Our contribution is the adaptation of this algorithm to the task of maximum

likelihood haplotyping for general pedigrees.

The allele-recoding algorithm is based on the observations that alleles have two roles in likelihood computations, and that valid recoding does not alter these roles:

1. **Determine prior probabilities of founders' genotypes.** The genotype frequency of a founder is computed using the population frequencies of the two alleles that constitute the genotype, assuming Hardy-Weinberg equilibrium.
2. **Determine recombination events.** A recombination event is determined by identifying the parental origin of the child's alleles, that is whether the child's alleles came from the paternal or maternal haplotype of his parent. Note that the allele identity does not matter here; only whether the allele matches the parent's paternal or maternal allele.

An allele is defined to be *transmitted* if the following two conditions are fulfilled: (i) the allele appears in the ordered genotype list of a *typed* descendant D of P , as inherited from; (ii) there is some path from P to D containing only *untyped* descendants in the pedigree, namely, D is the nearest typed descendant of P on that path. The remaining alleles are defined to be *non-transmitted*. In terms of determining recombination events, a person's non-transmitted alleles are indistinguishable from one another by data, and can therefore be combined into a single representative allele.

The allele recoding algorithm (Figure 2) is executed after initial value elimination is performed on the input pedigree as described in the previous section. At this stage, each individual is associated with two allele lists at each locus, a paternal allele list and a maternal allele list. In the first stage of the allele recoding algorithm, all the alleles of typed individuals are marked as transmitted, and all the alleles of untyped individuals are marked as non-transmitted. Next, the pedigree is traversed in a bottom-up manner where each person is updated after his children have been updated. The update is performed as follows: each allele of the father which appears as transmitted in the paternal allele list of the child is marked as transmitted in the father as well. A similar update is done for mothers. After a founder has been updated by all his children, a final processing

is performed on the founder's two lists of non-transmitted alleles. In this stage, only the allele with the highest frequency in each of the two lists is kept. At the end of the algorithm, each set of non-transmitted alleles forms a set (e.g., $\{A, B, C\}$), and each transmitted allele A forms a set including only itself, i.e., $\{A\}$. Recall that if a parent has the ordered genotype $A|B$ and its child has allele C , then C is inherited from the parent if $A = C$ or $B = C$. After allele recoding, however, A , B , and C are now sets of alleles, and hence C is inherited from the parent if $A \subseteq C$ or $B \subseteq C$. This is termed *fuzzy inheritance* in (O'connell and Weeks, 1995).

In the appendix, we prove that the probability of the assignment found for the regular case (without allele recoding) is the same as the one found in the case of allele recoding. This claim proves the correctness of the allele recoding algorithm for maximum likelihood haplotyping. Note that there is often more than one maximum likelihood assignment, but the algorithm described herein produces only one. To produce all possible maximum likelihood assignments, one needs to change the Elim-Max algorithm, described in the appendix, to store all optimizing values of each variable X_v rather than storing a single optimizing value. This addition increases the time and space complexity of the computations.

3.3 COMPUTATION ORDER

The problem of determining a good combined order of variable elimination and conditioning is important for both likelihood computations and haplotyping. This problem has been addressed quite extensively in the context of genetic linkage analysis. The two main approaches for performing likelihood computations on pedigrees are the Elston-Stewart algorithm (Elston and Stewart, 1971) which peels one nuclear family after another, and the Lander-Green algorithm (Lander and Green, 1987) which peels one locus after another. These two approaches are both variants of variable elimination methods which use different fixed elimination orders. Finding a good elimination order is also essential in a variety of combinatorial problems, such as: constraint satisfaction, independent set, dom-

inating set, graph K-colorability and Hamiltonian circuit (Arnborg, 1985; Dechter, 1996) as well as in other applications of Bayesian networks.

The problem of determining a good combined order of variable elimination and conditioning can be reduced to a graph-theoretic problem, namely, all elimination and conditioning operations are performed on the undirected graph representation of the Bayesian network. The undirected graph representation is obtained from the Bayesian network by connecting each pair of vertices that have edges leading into a common vertex, and removing the directionality of the edges (Pearl, 1988).

When a vertex is eliminated from the graph, its set of neighbors are connected to form a clique. The cost of eliminating vertex v from graph G_i is $c_{G_i}(v) = \prod_{u \in \bar{N}_{G_i}(v)} w(u)$, where $\bar{N}_{G_i}(v)$ represents the set of neighbors of v including v itself, and $w(v)$ is the weight of v , namely, the number of possible values of variable X_v . In the case when there is no memory limitation, we aim to find an elimination order \hat{X}_α which satisfies $\hat{X}_\alpha = \arg \min_\alpha C(X_\alpha)$, where

$$C(X_\alpha) = \sum_{i=1}^n c_{G_i}(X_{\alpha(i)}). \quad (4)$$

and α denotes a permutation on $\{1, \dots, n\}$. In Eq. (4), G_i , $i = 2, \dots, n$ denotes the sequence of residual graphs obtained from a given graph $G_1 = G$ by eliminating its vertices in the order $X_{\alpha(1)}, \dots, X_{\alpha(i-1)}$.

This cost function, which is often referred to as the *total state space* (Kjærulff, 1990), is an approximated measure of the time and space complexity of the computations, provided that the heaviest clique created fits into the RAM size of the working environment. If this is not the case, then conditioning is needed and a more elaborate cost function, described in (Fishelson and Geiger, 2003), is required. In this case we obtain a constrained elimination order $X_{\alpha,\beta} = ((X_{\alpha(1)}, \dots, X_{\alpha(n)}), \beta)$ which is a sequence of vertices along with a binary vector β such that vertex $X_{\alpha(i)}$ is eliminated if $\beta_i = 0$ and conditioned on if $\beta_i = 1$.

If we replace the summation in Eq. (4) with maximization, then the problem is reduced

to finding the *weighted treewidth* of the graph. If, in addition, the weight of all vertices is constant, then this problem is reduced to finding the *treewidth* of the graph, which is NP-complete (Arnborg et al., 1987).

We devised a new algorithm for finding a combined order of variable elimination and conditioning and applied it for both haplotyping and likelihood computations. The algorithm is composed of two stages. First, a set of reduction rules are applied on the graph as a preprocessing step. Second, several stochastic-greedy algorithms are applied sequentially to determine an elimination order for the residual graph.

3.3.1 Preprocessing Rules

Eijkhof et al. (2002) present a set of safe reduction rules for the *weighted treewidth* problem. Application of these rules can significantly reduce the size of the graph, without increasing the weighted treewidth of the graph. We tested these rules for our optimization problem and found that two of these reduction rules, the *simplicial* and the *almost simplicial* rules, are worthy to incorporate. The run time of these rules is negligible compared to the total run time for finding an elimination order, and by reducing the size of the graph, each iteration of the stochastic-greedy algorithms applied later is shorter. Throughout the application of the reduction rules, a variable *low* which represents the largest lower bound known for the weighted treewidth of the original graph is maintained. We denote by $nw(v)$ the product $\prod_{u \in \bar{N}_{G_i}(v)} w(u)$.

Simplicial rule. Let v be a *simplicial* vertex in G_i , namely its set of neighbors form a clique. The simplicial rule removes v from the graph, and updates the variable *low*: $low = \max(low, nw(v))$.

Almost simplicial rule. A vertex v is called an *almost simplicial* vertex in G_i if all its neighbors, except one u , form a clique. Vertex v is removed if $low \geq nw(v)$ and $w(v) \geq w(u)$.

3.3.2 Stochastic-Greedy Algorithms

Three Stochastic-Greedy algorithms for finding a combined order of variable elimination and conditioning, all based on the same common procedure SG (Figure 3), have been incorporated. The input to this procedure is a weighted undirected graph $G(V, E, w)$ resulting from the application of the reduction rules, a threshold T which represents the memory limitation, and two cost functions, C_1 and C_2 , that vary between the three algorithms. The threshold T is determined dynamically according to the memory available at run time. According to cost function C_1 the next vertex to eliminate is chosen, and according to C_2 a vertex to condition on is chosen. In each iteration, three vertices with a minimal cost (according to C_1) are selected, and a coin, biased according to the costs of the vertices, is flipped to choose between them. If in iteration i the weight of the clique created by the elimination of each vertex is above the given threshold, a vertex is chosen (according to C_2) to condition on rather than eliminate. Procedure SG is run many times, each time finding a new elimination order, and comparing it to the best order found so far. If the cost of the new elimination order is smaller than that of the best previously found order, then the new order and cost are recorded.

The stochastic-greedy algorithms used are: *Min-Weight* (Min-W), *Min-Fill*, and *Weighted Min-Fill* (WMin-Fill). The first two algorithms are based on known cost functions (Kjærulff, 1990), whereas the cost function of the *Weighted Min-Fill* algorithm is new and shows superior performance in many cases, as demonstrated by Experiment B. We now describe the different cost functions. The cost C_1 of eliminating a vertex according to the *Min-Weight* heuristic is the product of weights of its neighbors, whereas the cost of eliminating a vertex according to the *Min-Fill* heuristic is the number of edges that need to be added to the graph due to its elimination. The *Weighted Min-Fill* heuristic is a novel modification of the *Min-Fill* heuristic to a weighted graph. If we define the *weight of an edge* to be the product of weights of its constituent vertices, then the cost of eliminating a vertex according to the *Weighted Min-Fill* heuristic is the sum of weights of the edges that need to be added due to its elimination. The cost function C_2 for the

Min-Fill and Weighted Min-Fill algorithms, is the same as the first option described in (Fishelson and Geiger, 2003), i.e.,

$$C_2(X) = \sqrt{|N_{G_i}(X)|}C_1(X),$$

where $N_{G_i}(X)$ represents the set of neighbors of X in G_i . The cost function C_2 for the Min-Weight algorithm is the same as the second option described in (Fishelson and Geiger, 2003), i.e.,

$$C_2(X) = \sqrt{|f_{G_i}(X)|}C_1(X),$$

where $f_{G_i}(X)$ are the factors in G_i that include X .

The incorporation of these three algorithms and not others that were tried, such as Maximum Cardinality Search MCS (Tarjan and Yannakakis, 1984) or a weighted version of it (WMCS), is based on the fact that the other algorithms are superior in only a few cases. Neither of the three algorithms that were incorporated is better than the others in all cases, and therefore each of the algorithms is run a certain percentage of the total optimization time (Figure 4). We denote by %MW, %MF, and %WMF the percentage of iterations spent on running the Min-Weight, Min-Fill, and Weighted Min-Fill algorithms respectively. These percentages have been determined experimentally based on the relative performance of each algorithm. The total number of iterations N is determined according to the complexity of the problem at hand, which is estimated according to the cost of the elimination order found by the *deterministic-greedy* Min-Weight algorithm (Fishelson and Geiger, 2003). The only difference between the deterministic algorithm and the stochastic algorithm is that in each iteration, the deterministic algorithm chooses to eliminate a vertex with a minimal elimination cost according to the Min-Weight cost function, rather than flip a coin. If the cost of the elimination order found by this deterministic algorithm is lower than some threshold C_{min} then no optimization is performed.

The superior performance of the Min-Fill algorithms, as observed from Experiment B, may stem from the following observation. Assume there is a set of vertices that almost forms a clique. Despite the fact that the elimination of one of these vertices would add

only a few edges to the graph, this vertex would not be a preferred vertex to eliminate by the Min-Weight heuristic if the clique that is created is heavy. Hence, the Min-Weight Heuristic may complicate the given graph by choosing in some early iterations vertices whose elimination creates light cliques but possibly adds many edges to the graph.

4 RESULTS

We performed several experiments to test the new algorithms. The experiments can be divided into two classes. The first class compares the performance of our haplotyping algorithm to existing haplotyping algorithms. The second class is designed to test the performance of the new optimization algorithm. In all experiments in which pedigree data was simulated, the assumptions underlying the simulation were: Hardy-Weinberg and linkage equilibrium, no mutation and no interference. All input files and results of the experiments are readily available online.

4.1 Evaluation of the Haplotyping Algorithm

Experiment A (Simulation Study). We tested our haplotyping algorithm on a complex pedigree of moderate size (Figure 2 in (Lin, 1996)). So far, only an approximate haplotype analysis was possible for this pedigree. We simulated a random haplotype configuration for this pedigree using the simulation guidelines described by Lin and Speed (1997), and obtained a maximum likelihood haplotype configuration in several minutes using SUPERLINK. This pedigree consists of 27 individuals and is highly inbred. All individuals, except for those in the first two generations, were typed at 10 polymorphic markers, each with 5 alleles of equal frequencies. The recombination fraction between each pair of consecutive markers was set to 0.05. The progress made in resolving this pedigree can be appreciated by citing Lin and Speed (1997):

This is a very complex though moderate sized pedigree, with 10 polymorphic markers, and it does not seem to us to be possible to carry out a haplotype analysis for it with existing non-simulation-based statistical methods.

We note that GENEHUNTER removes 12 individuals from the pedigree in order to perform the computations. To the best of our knowledge, no previous exact algorithm can produce the maximum likelihood haplotype configuration for this pedigree.

Experiment B (Testing Accuracy). Existing approximate methods for haplotyping provide no guarantee on the accuracy of the output. Using our haplotyping algorithm,

approximated haplotyping can be compared with the optimal solution on larger pedigrees than was previously possible. This experiment tested the accuracy of a state of the art program that uses MCMC, called SIMWALK2 (Sobel and Lange, 1996). We tested 75 random data sets consisting of 15 to 50 individuals and up to 10 markers. SIMWALK2 found a maximal likelihood assignment in 45 out of the 75 data sets. In the other 30 data sets, the average difference in the log-likelihood of the assignment found by SIMWALK2 compared to the maximal likelihood assignment was merely 1%.

The MRH software (Qian and Beckmann, 2002) could not run on any of the data sets used in this experiment. We also tested the Block-Extension option of PEDPHASE (Li and Jiang, 2003b), running it 200 times on each data set. Coherent output was obtained in %37.64 of the runs. In those cases where coherent output was produced, SUPERLINK produced haplotype configurations with less or equal number of recombination events compared to PEDPHASE. This is interesting to note because the goal of PEDPHASE is to minimize the number of recombination events while SUPERINK maximizes the likelihood of data.

Experiment C (Published Disease Data). We analyzed two published data sets from a study of the Krabbe disease by Oehlman et al. (1993), and from a study on Episodic Ataxia (EA) by Litt et al. (1994). The first data set consists of 9 individuals typed at 8 polymorphic markers. The second data set consists of 29 individuals, which are all typed at 9 polymorphic markers except for the first two generation founders. For the Krabbe data set, the most likely haplotype configuration obtained by SUPERLINK is identical to the one obtained by MCMC via SIMWALK2 (Sobel et al., 1995; Sobel and Lange, 1996), by Lin and Speed (1997), and by PEDPHASE (Li and Jiang, 2003b).

For the Episodic Ataxia data set, the most probable configuration found by SUPERLINK is configuration D in Figure 2 of Qian and Beckmann (2002), which differs from the one obtained by SIMWALK2 (Sobel et al., 1995; Sobel and Lange, 1996) in the position of one recombination event. The only difference is the genotype phase in the fourth marker of individuals 1007 and 113. This configuration is also very similar to the one found by

Lin and Speed (1997).

4.2 Evaluation of the Optimization Algorithm

Experiment D (Stochastic Algorithms). This experiment compared the performance of different stochastic-greedy algorithms (Table 1). Each of the stochastic algorithms is run for 1000 iterations, after the reduction rules have been applied. We used graphs created from simulated pedigree data. Note that this experiment compares the elimination costs found by the algorithms for the case where $T = \infty$, namely no conditioning is performed. The results presented in Table 1 are a representative sample from the experiment that was performed on 100 data sets. In 100 data sets, the distribution of algorithms that found the lowest cost was as follows: Min-Weight - 4%, MCS - 9%, WMCS - 7%, Min-Fill - 25%, and Weighted Min-Fill - 76 %. Note that these percentages do not sum to 100% since, for some data sets, several algorithms found a minimal cost elimination order. As can be seen, the Min-Fill and Weighted Min-Fill are superior to the other heuristics. However, since the Min-Weight heuristic is the fastest and it works well when conditioning is needed (results not shown), it is profitable to first run it and then run the Min-Fill and Weighted Min-Fill heuristics. The MCS and Weighted-MCS heuristics have been found to hardly contribute when applied after the Min-Weight heuristic and are therefore not incorporated. To summarize, using an algorithm which combines the three algorithms, Min-Weight, Min-Fill, and Weighted Min-Fill, is superior to running only one of them, provided the optimization time is small enough compared to the total run time, as is the case for sufficiently large pedigrees.

Experiment E (Total Run Time). This experiment compared the run time of likelihood computation with the new optimization algorithm for determining an elimination order presented herein, to the run time with the previous optimization algorithm. The total run time includes both optimization time and inference time. We demonstrate the performance of the optimization algorithm using likelihood computations rather than haplotyping, since haplotyping was not implemented in the previous version of SUPER-

LINK. We tested 50 randomly simulated data sets chosen so that the run time in the new version is above 10 seconds and below 10 hours. The graph in Figure 5 shows that timing was improved in 47 out of 50 data sets, often by an order of magnitude.

Experiment F (Benchmarks). This experiment (Table 2) tested the performance of the three stochastic-greedy algorithms, on eight known benchmarks for Bayesian networks inference. The stochastic algorithms are run after the reduction rules have been applied. We present the best elimination cost found by the algorithms after 100 iterations and after 1000 iterations. Also presented for each benchmark are the best known elimination cost and the elimination cost found by HUGIN6.1 (Andersen et al., 1989; Hugin, 2002), which is a leading software for Bayesian networks. As can be seen, in most cases, running a few iterations of the stochastic-greedy algorithms is superior to the algorithm of HUGIN6.1. Another conclusion that can be drawn from this experiment is that the results of the stochastic-greedy algorithms are comparable to *simulated annealing* (Kjærulff, 1990) which is a very time consuming algorithm, and hence not a viable option for minimizing the total run time for genetic linkage analysis. As can be seen from Table 2, after 100 iterations the best known cost for the Water problem is matched, and after 1000 iterations the best known cost for the Barley problem is also matched.

Experiment G (Reduction Rules). This experiment tested the gain due to the reduction rules presented in Eijkhof et al. (2002). The data sets used are 20 graphs, out of which 12 are created from simulated pedigree data and 8 are known benchmarks for Bayesian networks inference. We have found that two of the reduction rules, the *simplicial* rule and the *almost simplicial* rule, are worthy to apply in almost every problem. Applying these rules usually eliminated between 50% and 60% of the vertices in the data sets (Table 3). The other two reduction rules, the *buddies* rule and the *cube* rule, are more time consuming, and applying them does not yield sufficient improvement; therefore we do not use these rules.

5 DISCUSSION

The use of Bayesian networks enables efficient maximum likelihood haplotyping for more complex pedigrees than was previously possible. We adapted and combined several algorithms, allele recoding, reduction rules, elimination algorithms, and the Elim-Max inference algorithm, into a working system for haplotyping readily available for use by geneticists. This advancement can also be utilized for checking approximate haplotyping algorithms such as MCMC.

Several new features have been incorporated in the algorithm for optimizing the elimination order presented herein. First, reduction rules are applied to speed the computations. Second, several greedy algorithms are run, rather than one. Their allocated relative run time is based on experiments on many Bayesian networks. Third, among the greedy algorithms that are applied, a new greedy algorithm is introduced (Weighted Min-Fill). Fourth, the threshold which controls the time-space tradeoff is determined at run time according to the memory of the computer, rather than fixed a priori. Our new optimization algorithm is applied whenever the elimination cost found by a quick greedy algorithm is above a certain threshold ($\log_{10} C \geq 9$ in the current implementation). The result is an optimization algorithm which is superior in total run time to the algorithm in (Fishelson and Geiger, 2003) in 47 of the 50 instances we tried, often by an order of magnitude (Experiment D). The presented optimization algorithm is used in SUPERLINK for likelihood computations as well as for haplotyping.

Acknowledgments

This research was supported by the Israel Science Foundation. We thank Gil Rubin and Shaul Karni for ideas regarding efficient implementation of our haplotyping algorithm. We thank Uffe Kjærulff and the Research Unit of Decision Support Systems Aalborg University for providing us the datasets and support to conduct experiment C. We also thank Anna Tzemach for computer assistance.

Appendix A

Algorithm Elim-Max (Dechter, 1996)

Input: A Bayesian network $\langle G, P \rangle$; an ordering d of the variables; evidence ε .

Output: The most probable assignment to the variables of the Bayesian network, and its probability.

1. **Initialize:** Generate an ordered partition of the conditional probability tables $\{P_i\}$ into buckets, where bucket B_i contains all the probability tables and evidence whose largest-index variable is X_i .

2. **Backward phase (compilation phase):**

For $i = n$ to 1, process bucket B_i as follows:

Let h_1, \dots, h_k be all the probability tables (new and old) in B_i at the time it is processed, and let S_1, \dots, S_j be the subset of variables in B_i on which probability tables (new and old) are defined.

- **If** B_i contains $X_i = a_i$ (X_i is observed), assign $X_i = a_i$ to each h_l , and place it in the bucket of the largest-index variable that appears in its scope.
- **Else**, $\mathbf{U}_i = \bigcup_{j=1}^k S_j - \{X_i\}$. Generate function $h_i = \max_{x_i} \prod_{l=1}^k h_l$, and place it in the bucket of the largest-index variable in \mathbf{U}_i . Store the optimizing value of X_i for each tuple of \mathbf{U}_i , $X_i^{opt}(\mathbf{U}_i) = \arg \max_{x_i} \prod_{l=1}^k h_l$.

3. **Store:** the constant computed in B_1 . This is the probability of the most probable assignment to the variables of the network.

4. **Forward phase (process X_i after finding an assignment to X_1, \dots, X_{i-1}):**

For $i = 1$ to n , process bucket B_i as follows:

- Given the assignment $\mathbf{U}_i = \mathbf{u}_i$, choose $x_i = X_i^{opt}(\mathbf{u}_i)$.

5. **Return:** the assignment selected for the variables of the network and its probability.

Appendix B

Proposition: The probability of the haplotype configuration found using allele recoding is the same as the probability of the haplotype configuration found without allele recoding.

Proof: Recall that a haplotype configuration is an assignment of values to all genetic loci variables of individuals in the pedigree. For the proof, we use the word *assignment* rather than haplotype configuration. We divide the set of all possible assignments to two sets: A_1 and A_2 . The set A_1 consists of all assignments which fulfill the following conditions: (i) at least one allele assigned to a founder belongs to the set of non-transmitted alleles of this founder, and (ii) this allele is not the one with the highest frequency among the set of non-transmitted alleles of this founder. The set A_2 consists of all other consistent assignments given the data, namely, all assignments which fulfill the following condition: each allele assigned to a founder either belongs to the set of transmitted alleles of this founder, or it is the allele of maximum frequency among the list of non-transmitted alleles of this founder. The proof follows from the following two claims.

Claim 1: The maximum likelihood assignment is in the set A_2 .

Proof: Assume to the contrary of the claim that the maximum-likelihood assignment is in the set A_1 and denote it by a_0 . Suppose, without loss of generality, that for some founder P in locus i , there are two non-transmitted alleles, N_1 and N_2 . Assume that allele N_1 is chosen in assignment a_0 , and that $Pr(N_2) > Pr(N_1)$. Since N_1 is non-transmitted, no typed descendant of P inherited this allele from P . Hence, we can replace allele N_1 with allele N_2 in assignment a_0 in all individuals that inherited this allele from P according to a_0 , obtaining a consistent assignment denoted by a_1 . The only difference in the likelihood of these two assignments is in the frequency of allele N_2 compared to the frequency of allele N_1 . Since $Pr(N_2) > Pr(N_1)$, it follows that $Pr(a_1) > Pr(a_0)$ in contrary to the assumption that a_0 is the maximum likelihood assignment.

Claim 2: The probability of each assignment $a \in A_2$ using allele-recoding is the same as without using allele recoding.

Proof: There are three types of probability functions:

1. $Pr(\mathbf{S}, \varepsilon)$ - this function does not change under allele recoding because selector variables do not change.
2. $Pr(\mathbf{F}, \varepsilon)$ - this function is the product of frequencies of all alleles in the genotypes of founders determined by the given assignment. These functions do not change under allele recoding.
3. $Pr(\mathbf{N}|\mathbf{F}, \mathbf{S}, \varepsilon)$ - this probability equals 1 if the assignment is consistent and 0 otherwise. When using allele recoding, it is defined via fuzzy inheritance. If the assignment is consistent without allele recoding, it is consistent also with allele-recoding, since all founders' alleles exist also after recoding by definition of the set A_2 . In the case of non-founders, no allele is erased. The alleles are merely divided into groups of transmitted and non-transmitted alleles.

To summarize, all probability functions, whose product constitutes the probability of an assignment $a \in A_2$, remain the same with allele recoding and without allele recoding, and hence, the probability of assignment a is maintained under allele recoding.

6 Electronic-Database Information

The SUPERLINK program is available as an executable for Linux, Windows and Unix operating systems, at <http://bioinfo.cs.technion.ac.il/superlink/temp/V1.4.html>, along with user documentation. Also available are the input files and results of all the experiments performed.

References

- Abecasis, G. R., Cherny, S. S., Cookson, W. O. and Cardon, L. R. (2002). Merlin-rapid analysis of dense genetic maps using sparse gene flow trees., *Nat Genet* **58**: 97–101.
- Andersen, S. K., Olesen, K. G., Jensen, F. V. and Jensen, F. (1989). HUGIN - a shell for building Bayesian belief universes for expert systems., *Proc. of the 11th International Joint Conference on Artificial Intelligence (IJCAI)*, Vol. 2, pp. 1080–1085.
- Arnborg, S. (1985). Efficient algorithms for combinatorial problems on graphs with bounded decomposibility., *BIT* **25**: 2–23.
- Arnborg, S., Corneil, D. G. and Proskurowski, A. (1987). Complexity of finding embeddings in a k-tree., *SIAM J. Alg. Disc. Meth.* **8**: 277–284.
- Dechter, R. (1996). Bucket elimination: A unifying framework for probabilistic inference., *Proc. Twelfth Conference on Uncertainty in Artificial Intelligence (UAI)*, pp. 211–219.
- Eijkhof, F., Bodlaender, H. and Koster, A. (2002). Safe reduction rules for weighted treewidth., *Technical Report 02-49*, ZIB, Berlin, Germany.
- Elston, R. C. and Stewart, J. (1971). A general model for the analysis of pedigree data., *Hum. Hered.* **21**: 523–542.
- Ewens, W. J. and Spielman, R. S. (1995). The transmission/disequilibrium test: History, subdivision, and admixture., *Am. J. Hum. Genet.* **57**: 455–465.
- Fishelson, M. and Geiger, D. (2002). Exact genetic linkage computations for general pedigrees., *Bioinformatics* **18**(Suppl. 1): S189–S198.
- Fishelson, M. and Geiger, D. (2003). Optimizing exact genetic linkage computations., *Proc. 7th Conf. on Computational Molecular Biology (RECOMB)*., pp. 114–121.
- Friedman, N., Geiger, D. and Lotner, N. (2000). Likelihood computation with value abstraction., *Proc. 16th Conf. on Uncertainty in Artificial Intelligence (UAI)*.

- Greenspan, G. and Geiger, D. (2003). Model-based inference of haplotype block variation., *Proc. 7th Conf. on Computational Molecular Biology (RECOMB)*., pp. 131–137.
- Gudbjartsson, F., Jonasson, K., Frigge, M. L. and Kong, A. (2000). Allegro, a new computer program for multipoint linkage analysis., *Nat. Genet.* **25**: 12 – 13.
- Gusfield, D. (2002). Haplotyping as perfect phylogeny: conceptual framework and efficient solutions., *Proc. 6th Conf. on Computational Molecular Biology (RECOMB)*., pp. 166 – 175.
- Hugin (2002). The API reference manual version 5.4.
- Kjærulff, U. (1990). Triangulation of graph - algorithms giving small total state space, *Technical Report R90-09*, Department of Computer Science, Aalborg University, Denmark.
- Kruglyak, L., Daly, M. J., Reeve-Daly, M. P. and Lander, E. S. (1996). Parametric and nonparametric linkage analysis: A unified multipoint approach., *Am. J. Hum. Genet.* **58**(6): 1347–1363.
- Kruglyak, L. and Lander, E. S. (1998). Faster multipoint linkage analysis using fourier transform., *J. of Comp. Biol.* **5**: 1–7.
- Lander, E. S. and Green, P. (1987). Construction of multilocus genetic maps in humans., *Proc. Natl. Acad. Sci.*, Vol. 84, pp. 2363–2367.
- Lange, K. and Goradia, T. M. (1987). An algorithm for automatic genotype elimination, *Am. J. Hum. Genet.* **40**: 250–256.
- Lange, K., Weeks, D. and Boehnke, M. (1988). Programs for pedigree analysis: Mendel, fisher, and dgene., *Genet. Epidemiol.* **5**: 471–473.
- Lauritzen, S. L. (1996). *Graphical Models*., Oxford University Press.

- Li, J. and Jiang, T. (2003a). Efficient rule-based haplotyping algorithms for pedigree data., *Proc. 7th Conf. on Computational Molecular Biology (RECOMB)*., pp. 197 – 206.
- Li, J. and Jiang, T. (2003b). Pedphase: haplotype inference for pedigree data., *Submitted*.
- Li, J. and Jiang, T. (2004). An exact solution for finding minimum recombinant haplotype configurations on pedigrees with missing data by integer linear programming., *Proc. 8th Conf. on Computational Molecular Biology (RECOMB)*.
- Lin, S. (1996). Multipoint linkage analysis via metropolis jumping kernels., *Biometrics* **52**: 299–309.
- Lin, S. and Speed, T. P. (1997). An algorithm for haplotype analysis., *J. Comput. Biol.* **4**: 535–546.
- Litt, M., Kramer, P., Browne, D., Gancher, S., Brunt, E., Root, D., Phromchotikul, T., Dubay, C. and Nutt, J. (1994). A gene for episodic ataxia/myokymia maps to chromosome., *Am. J. Hum. Genet.* **55**: 702–709.
- O’connell, J. R. and Weeks, D. E. (1995). The vitesse algorithm for rapid exact multi-locus linkage analysis via genotype set-recoding and fuzzy inheritance., *Nat. Genet.* **11**: 402–408.
- Oehlman, R., Zlotogora, J., Wenger, D. and Knowlton, R. (1993). Localization of the krabbe disease gene (galc) on chromosome 14 by multipoint linkage analysis., *Am. J. Hum. Genet.* **53**: 1250–1255.
- Ott, J. (1999). *Analysis of Human Genetic Linkage*., Johns Hopkins University Press, Baltimore, Maryland.
- Park, J. (2002). Map complexity results and approximation methods., *Proc. 18th Conf. on Uncertainty in Artificial Intelligence (UAI)*, pp. 388–396.

- Pearl, J. (1988). *Probabilistic Reasoning in Intelligent Systems.*, Morgan Kaufmann, San Francisco, CA.
- Qian, D. and Beckmann, L. (2002). Minimum-recombinant haplotyping in pedigrees., *Am J Hum Genet* **70**: 1434–1445.
- Schäffer, A. A. (1996). Faster linkage analysis computations for pedigrees with loops or unused alleles., *Hum. Hered.* **46**: 226–235.
- Sobel, E. and Lange, K. (1996). Descent graphs in pedigree analysis: applications to haplotyping, location scores, and marker sharing statistics., *Am J Hum Genet* **58**: 1323–1337.
- Sobel, E., Lange, K., O’connell, J. R. and Weeks, D. E. (1995). Haplotyping algorithms., *IMA volumes in mathematics and its applications* **81**: 89–110.
- Stephens, M., Smith, N. J. and Donnelly, P. (2001). A new statistical method for haplotype reconstruction from population data., *Am. J. Hum. Genet.* **68**: 978–989.
- Tarjan, R. E. and Yannakakis, M. (1984). Simple linear time algorithm to test chordality of graphs, test acyclicity of hypergraphs, and selectively reduce acyclic hypergraphs., *SIAM J. Comput.* **13**: 566–579.
- Wisjman, P. (1987). A deductive method of haplotype analysis in pedigrees., *Am. J. Hum. Genet.* **41**: 356–373.

Table 1: Comparing Different Stochastic-Greedy Algorithms.

#Loci ^a	#People ^b	%Typed ^c	Min-W ^d	MCS ^e	WMCS ^f	Min-Fill ^g	WMin-Fill ^h
12	25	80	3.97	5.15	5.36	3.94	3.93
14	27	63	17.00	15.54	22.27	14.17	14.22
16	22	89	5.30	8.08	8.61	5.22	5.22
20	31	71	10.99	12.62	15.67	10.38	10.03
17	29	56	12.77	11.41	20.15	11.14	10.56
19	33	95	6.19	7.03	7.55	5.70	5.91
18	31	47	8.36	10.27	13.01	8.23	7.90
13	30	30	2.79	2.81	2.81	2.80	2.80
22	38	86	6.72	9.41	13.85	6.46	6.37
20	48	91	9.28	14.73	13.45	7.84	7.61

Note - The cost reported is the \log_{10} of the sum of weights of cliques created during elimination. Best results are in bold.

^aNumber of loci being analyzed.

^bNumber of people in the pedigree.

^cPercentage of typed people in the pedigree.

^dElimination cost obtained by the Min-Weight algorithm.

^eElimination cost obtained by the Maximum Cardinality Search algorithm.

^fElimination cost obtained by the Weighted Maximum Cardinality Search algorithm.

^gElimination cost obtained by the Min-Fill algorithm.

^hElimination cost obtained by the Weighted Min-Fill algorithm.

Table 2: Testing some Stochastic-Greedy Algorithms on eight known benchmarks.

Problem	Known Cost ^a	Hugin Cost ^b	Cost after 100 Iterations			Cost after 1000 Iterations		
			Min-W ^c	Min-Fill ^d	WMin-Fill ^e	Min-W	Min-Fill	WMin-Fill
Barley	1.71 E07	1.73 E07	1.95 E07	1.73 E07	1.82 E07	No change	1.71 E07	1.8 E07
Diabetes	9.83 E06	1.04 E07	1.03 E07	1.56 E07	1.24 E07	1.01 E07	1.38 E07	1.21 E07
Link	2.4 E07	2.62 E07	4.31 E07	2.76 E07	3.78 E07	4.15 E07	2.66 E07	3.14 E07
Mumin1	8.69 E07	1.88 E08	1.84 E08	1.54 E08	1.39 E08	No change	8.76 E07	No change
Mumin2	2.05 E06	2.76 E06	4.05 E06	3.53 E06	4.36 E06	3.79 E06	No change	4.35 E06
Mumin3	3.08 E06	3.24 E06	3.28 E06	3.26 E06	3.11 E06	3.27 E06	3.2 E06	3.11 E06
Mumin4	9.84 E06	1.64 E07	1.72 E07	2.0 E07	1.37 E07	1.58 E07	No change	1.34 E07
Water	3.03 E06	8.04 E06	3.66 E06	3.03 E06	3.47 E06	3.62 E06	No change	No change

^aBest known elimination cost.

^bElimination cost found by HUGIN6.1 (Andersen et al., 1989; Hugin, 2002).

^cElimination cost obtained by the Min-Weight algorithm.

^dElimination cost obtained by the Min-Fill algorithm.

^eElimination cost obtained by the Weighted Min-Fill algorithm.

Table 3: Reduction Rules Experiments

#Vertices	#Edges	Using Two Reduction Rules			Using Four Reduction Rules		
		% of vertices eliminated	run time ^a	% of run time ^b	% of vertices eliminated	run time	% of run time
48	126	39.6	0.00	0.00	39.6	0.00	0.00
413	819	19.6	0.00	0.00	19.6	0.01	4.76
724	1738	53.2	0.00	0.00	53.2	0.02	7.13
189	366	52.2	0.00	0.00	52.2	0.00	0.00
1003	1662	68.4	0.00	0.00	68.4	0.03	9.56
1044	1745	82.4	0.00	0.00	83.5	0.03	9.18
1041	1843	74.0	0.00	0.00	74.0	0.02	6.83
32	123	31.3	0.00	0.00	31.3	0.00	0.00
1676	2987	58.0	0.03	5.54	58.0	0.10	18.7
2449	4320	57.7	0.06	6.74	57.7	0.34	33.2
3224	5641	54.5	0.08	7.14	54.7	0.30	16.7
3254	5754	59.0	0.06	4.41	59.0	0.30	21.9
4720	8211	57.8	0.15	7.58	57.8	0.75	29.8
6242	10814	54.7	0.16	4.98	55.2	0.63	18.3
4859	8517	59.6	0.12	6.85	59.6	0.56	23.6
7040	12104	60.0	0.28	9.14	60.0	1.32	31.7
9276	15926	55.3	0.30	7.47	55.6	1.56	24.3
6372	11198	57.8	0.17	5.87	57.8	1.56	24.3
9292	16096	57.9	0.37	7.73	57.9	2.23	38.9
12278	21215	54.3	0.54	7.78	54.5	6.02	73.2

Note - Results in bold indicate problem instances where application of the four reduction rules eliminated more vertices than application of only two rules.

^aRun time is specified in seconds.

^bThe percentage of time spent on reduction rules out of the time required for finding an elimination order, when merely one iteration of the stochastic-greedy algorithm is performed.

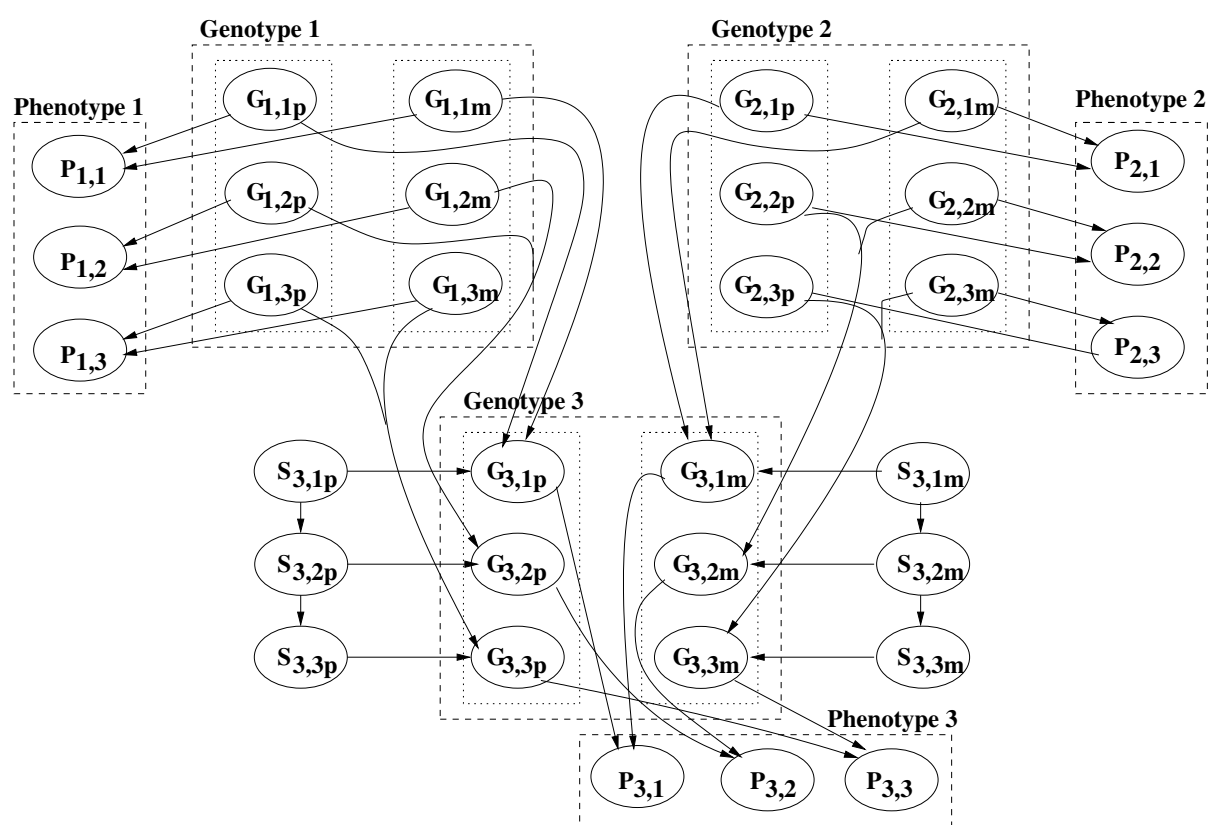


Figure 1: A fragment of a Bayesian network representation of parents-child interaction in a 3-loci analysis. Adapted from Friedman et al. (2000).

Algorithm Allele-Recoding

Input: A pedigree P of size n , where each individual is associated with a list of possible paternal alleles and a list of possible maternal alleles at a given locus l .

Output: Each of the two allele lists of each individual is replaced by a list of sets of alleles.

1. **For** $i \leftarrow 1$ to n **do** {Initialize}
 - **If** individual i is typed at locus l , **then**

Mark all the alleles in both his allele lists as *transmitted*.
 - Else**

Mark all the alleles in both his allele lists as *non-transmitted*.
2. **Traverse** P in a bottom-up manner. Update individual i as follows: {Mark}
 - **For** each child j of i **do**
 - If** i is a male, **then**

Mark each allele of i which appears as transmitted in j 's *paternal* allele list as transmitted in i as well.
 - Else**

Mark each allele of i which appears as transmitted in j 's *maternal* allele list as transmitted in i as well.
 - **If** i is a founder, **then** {Only for haplotyping}

Let A_{nt} be the set of non-transmitted alleles of i at locus l , and let $a_k \in A_{nt}$ be the allele with the highest population frequency in A_{nt} .

Remove all alleles $a_m \in A_{nt}$, $a_m \neq a_k$.
3. **For** $i \leftarrow 1$ to n **do** {Recode}
 - Replace each transmitted allele T by the set $\{T\}$.
 - Replace all non-transmitted paternal alleles by one set P_n , which consists of these alleles.
 - Replace all non-transmitted maternal alleles by one set M_n , which consists of these alleles.
4. **Return.**

Figure 2: The allele recoding algorithm for the haplotyping problem. For likelihood computations, the second stage of step 2 is removed.

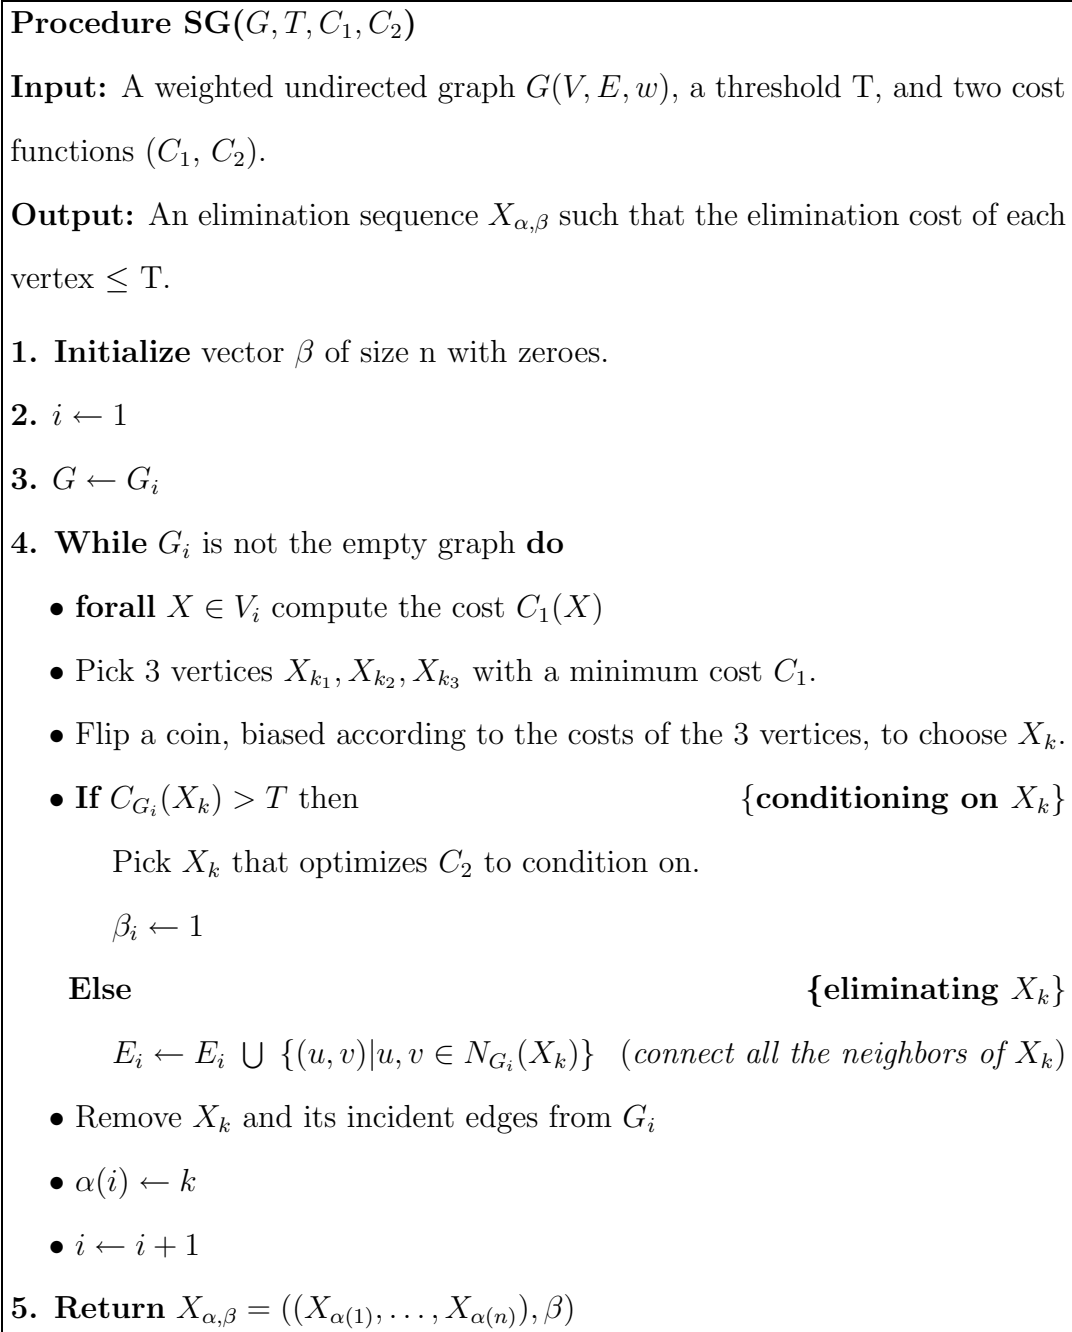


Figure 3: Procedure SG

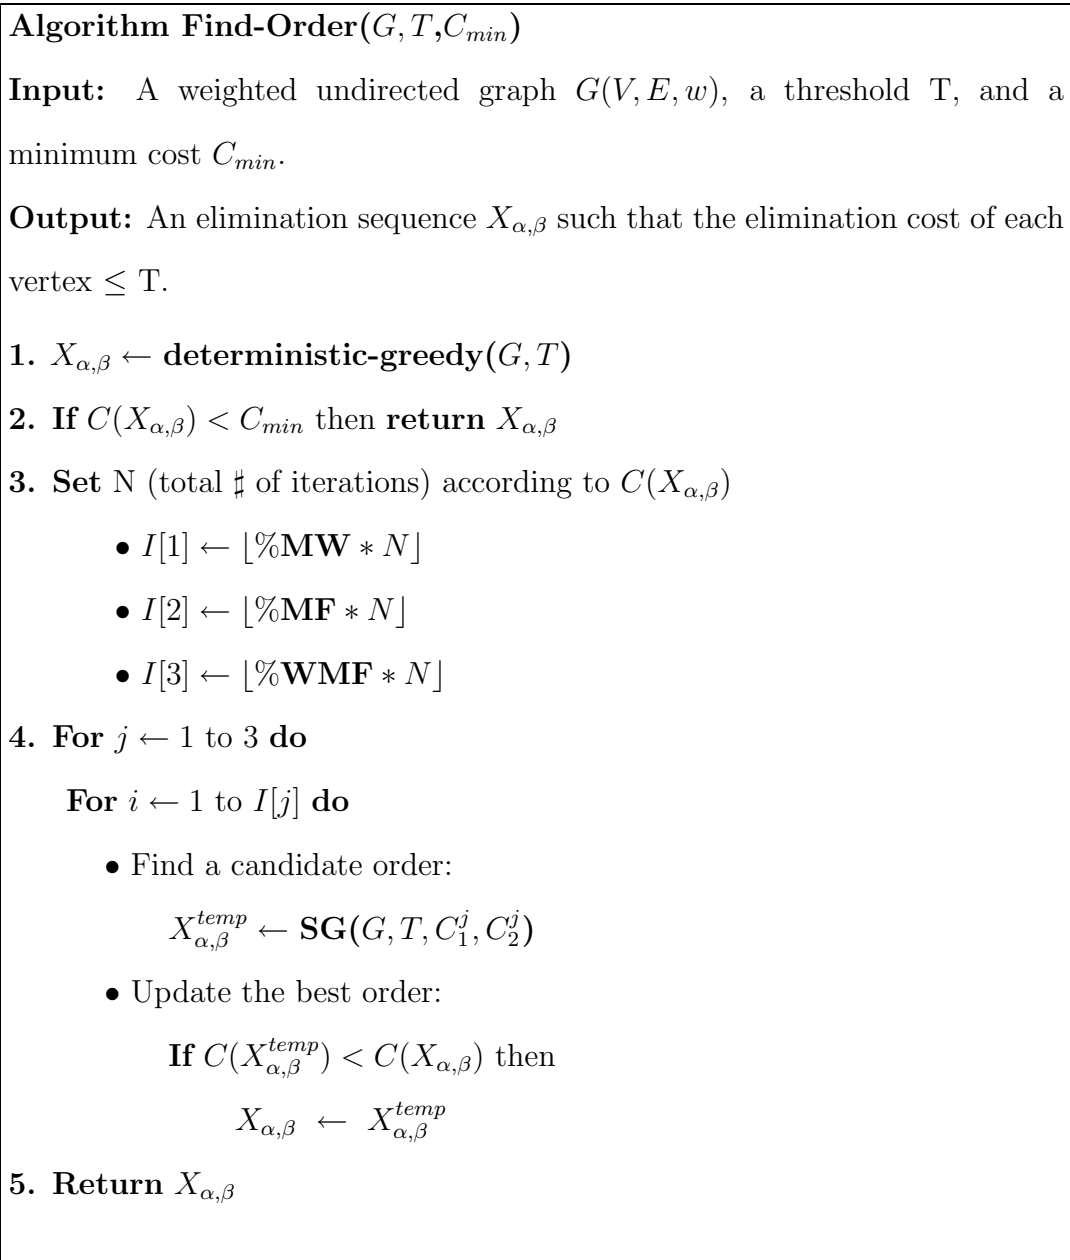


Figure 4: Algorithm Find-Order.




Figure 5: Run time comparison of likelihood computation using the new optimization algorithm versus the previous optimization algorithm. On each dataset, the old version was run for up to a couple of hours, and the total run time was estimated according to the percentage of the computation that was completed at that time. For those datasets where the estimated run time of the old version was well over 300 hours, the run time appears as 10^6 seconds.