

Association Mapping and Significance Estimation via the Coalescent

Gad Kimmel,^{1,2,*} Richard M. Karp,^{1,2} Michael I. Jordan,^{1,3} and Eran Halperin²

The central questions asked in whole-genome association studies are how to locate associated regions in the genome and how to estimate the significance of these findings. Researchers usually do this by testing each SNP separately for association and then applying a suitable correction for multiple-hypothesis testing. However, SNPs are correlated by the unobserved genealogy of the population, and a more powerful statistical methodology would attempt to take this genealogy into account. Leveraging the genealogy in association studies is challenging, however, because the inference of the genealogy from the genotypes is a computationally intensive task, in particular when recombination is modeled, as in ancestral recombination graphs. Furthermore, if large numbers of genealogies are imputed from the genotypes, the power of the study might decrease if these imputed genealogies create an additional multiple-hypothesis testing burden. Indeed, we show in this paper that several existing methods that aim to address this problem suffer either from low power or from a very high false-positive rate; their performance is generally not better than the standard approach of separate testing of SNPs. We suggest a new genealogy-based approach, CAMP (coalescent-based association mapping), that takes into account the trade-off between the complexity of the genealogy and the power lost due to the additional multiple hypotheses. Our experiments show that CAMP yields a significant increase in power relative to that of previous methods and that it can more accurately locate the associated region.

Introduction

Recent advances in genotyping technologies have considerably improved our understanding of common complex diseases through whole-genome association studies. In these studies a population of cases and controls is collected, and hundreds of thousands of single nucleotide polymorphisms (SNPs) are genotyped. These studies search for SNPs that are associated with the studied disease by measuring the difference in the SNP-allele distributions between the cases and the controls (e.g.,^{1,2}).

Because complex diseases are caused by multiple environmental and genetic factors, the differences in allele frequencies between the cases and the controls for any given SNP can be expected to be quite small. Therefore, analyses that achieve high statistical power are essential for these studies. Additionally, although current technology (e.g., the Affymetrix SNP Array 6.0 and the Illumina human1m-duo beadchip) allows measurement of nearly two million genetic variants for each individual, this is still only a fraction of the set of genetic variants, and statistical methods are needed to cope with this partial assessment of genetic variation.

The statistical analysis of a typical association study involves the testing of individual SNPs or genomic regions for association and the evaluation of the significance of the findings. The simplest approach to significance testing is to test each marker separately for association.^{3,4}

Many approaches have attempted to move beyond separate testing by leveraging the unobserved genealogy of the chromosomes (e.g.,^{5,6}). Such proposals aim to increase statistical power by taking into account the dependency

among SNPs. Model-based approaches in particular try to infer aspects of the unobserved genealogy. In practice, however, this is not a trivial task because one must infer the genealogy from the genotypes. As we show in this paper, the loss of information caused by erroneous inference of the genealogy can be detrimental to the association, and thus genealogy-based methods are not always desirable.

Previous methods that have used genealogies in association studies have faced two main challenges. First, the number of possible genealogies is very large, and even more so when recombination events are taken into account; thus, it is infeasible to examine all possible genealogies. Second, an inferred genealogy determines a large set of genealogy-based association tests (these can be expressed as tests of SNP interactions); a major challenge is how to choose a subset of these tests such that the increased number of hypotheses tested will not decrease the power. If the tests are not chosen properly, the statistical power can be reduced considerably as a result of the burden of multiple hypotheses, even when the genealogical modeling is accurate.

In this paper, we suggest a new genealogy-based approach that takes into account the trade-off between the complexity of the genealogy and the power lost because of multiple hypotheses. The approach we present seeks to avoid excessive power loss due to multiple testing while still testing the observed mutations and selected putative unobserved mutations suggested by plausible genealogies. As with previous genealogy-based methods, we test selected SNP interactions. The core of our method is to exploit properties of the coalescent to decide which interactions can be ignored. In a nutshell, we construct a perfect phylogeny graph that represents the genealogy of the haplotypes and restrict

¹Computer Science Division, University of California, Berkeley, Berkeley, CA 94720, USA; ²International Computer Science Institute, 1947 Center Street, Berkeley, CA 94704, USA; ³Department of Statistics, University of California, Berkeley, Berkeley, CA 94720, USA

*Correspondence: kimmel@cs.berkeley.edu

DOI 10.1016/j.ajhg.2008.10.017. ©2008 by The American Society of Human Genetics. All rights reserved.

attention to observed mutations and to unobserved mutations that are consistent with that graph.

Many genealogy-based association tests have been suggested in earlier work. One popular way of using genealogy in association studies is through the use of ancestral recombination graphs⁷ (ARGs). These graphs aim to model the coalescence and the recombination events explicitly. Several studies have proposed performing full-likelihood or Bayesian inference under the ARG model (e.g.,^{8,9}). This is, however, a technically challenging problem, and the proposed solutions are computationally feasible only on relatively small data sets. Zollner and Pritchard⁵ suggested solving this inference problem with an approximation approach in which one tests for association by performing a likelihood ratio test by calculating the probability of the disease mutation given the genotypes and the disease status. The inference is performed by a Markov chain Monte Carlo (MCMC) algorithm. This approach has the advantages of model-based procedures, but it is too expensive computationally to be used in a large-scale whole-genome association study.

A different approximation approach to association mapping was suggested by Durrant et al.¹⁰ Their main idea was to perform a cladistic analysis of SNPs. The cladogram captures the successive partitioning of SNP haplotypes into clusters. At each partition, clusters of haplotypes from the previous partition are merged such that the mean pairwise haplotype diversity is minimized within the new clade. The cladogram is built with a sliding window of SNPs. In each window the best partition of haplotypes is chosen. This procedure incorporates two levels of multiple testing, which are adjusted by a Bonferroni correction.

Minichiello and Durbin⁶ introduced another approximation scheme for the inference of ARGs. There are two stages to their analysis: First, they attempt to infer all plausible ARGs by using a heuristic algorithm. Second, a genealogical tree at each locus is built, and a possible causative mutation at each branch is tested. Because the true ARG is unknown, this analysis is averaged over a set of inferred ARGs.

In general, the genealogy-based methods are meant to improve upon the naive approach to association testing; in such an approach, each SNP is tested separately via a χ^2 test, and the tests are adjusted for multiple-hypothesis testing with a permutation test (we will refer to this approach as *standard χ^2*). To assess the extent to which existing methods have realized this goal, we compared these methods to the naive approach. In our experiments, we found that the naive approach has more power and a lower false-positive rate than any of the tested methods. This surprising result was the motivation for our new genealogy-based method, which we refer to as *CAMP* (coalescent-based association mapping).

Like previous methods, *CAMP* tests for interactions of SNPs or haplotypes with disease. To address the issues of computational complexity and multiple-hypothesis testing, we have emphasized reducing the number of tests.

The core of our method is to exploit properties of the coalescent to decide which interactions can be ignored. In brief, we construct a perfect phylogeny graph which represents the genealogy of the haplotypes and restrict attention to observed mutations and to unobserved mutations that are consistent with that graph, in the sense that each of the unobserved mutations is consistent with a larger graph that retains the perfect-phylogeny property. The larger graph represents a genealogy of the haplotypes with the unobserved mutation.

As opposed to ARGs, our method does not model the recombination events explicitly in detail. Indeed, we begin our presentation by making the simplifying assumption that there are no recombinations across the studied region and that there are no recurrent mutations (this is often referred to as a *perfect phylogeny model*, or a coalescent model with the infinite site assumption). It is well known that in order to satisfy the assertion that a region is consistent with a perfect phylogeny model, the region has to comply with the four-gamete test; put differently, every pair of SNPs has at most three out of the four possible haplotypes. We use this characterization to define a simple version of our method for generating unobserved mutations. We then back off from the simplifying assumption of no recombination and consider a model that allows some deviation from the four-gamete condition. This yields the *CAMP* algorithm, which can be viewed as a procedure for defining tests based on an approximate genealogy. A similar approach has been taken by Eskin et al. in their work on haplotype phasing (Eskin et al., 2003, in RECOMB 03. The Association for Computing Machinery, 104–113).

In order to evaluate the power achieved by *CAMP*, we have tested *CAMP* on an extensive number of simulated data sets. Our experiments show that *CAMP* yields a significant increase in power relative to that of previous methods. In particular, unlike previous methods, *CAMP* achieves an increase of more than 10% over the standard χ^2 . This advantage was observed with different sampling distances of SNPs and with different numbers of individuals. Thus, by using our method in association studies, we expect that more associated SNPs will be discovered as a result of the increased power.

Material and Methods

The General Framework

We begin by sketching the main idea of our approach. As in previous approaches, our goal is to exploit the unobserved genealogy of the population in order to map and evaluate the significance of associations. We do this by performing additional tests of interactions between SNPs; these tests correspond to unobserved mutations along the genealogical tree.

The basic idea of our approach is to restrict attention to interactions between pairs of SNPs that may represent a plausible mutation along a genealogy. Our approach relies strongly on the theory of *perfect phylogeny* of SNPs and haplotypes. Several studies in the literature have focused on this combinatorial object and yielded

theoretical characterizations that provide the basis of our approach; for background see ^{11,12,13,14}.

Notation and Definitions

Let n be the number of individuals tested and m be the number of markers. The $2n \times m$ haplotype matrix is denoted by H . Hence, $H_{i,j} = s$ if the i^{th} haplotype has type s at the j^{th} marker, where s can be 0 or 1. Let the vector of the disease status be d . The entries of d are 0 (for a healthy individual) or 1 (for an individual who has the disease).

For a pair of discrete vectors x and y , let $\Omega(x, y)$ denote their contingency table; i.e., $\Omega(x, y)$ is a matrix in which $\Omega(x, y)_{i,j} = |\{k | x(k) = i, y(k) = j\}|$ (in our case, the matrix is of size 2×2 because a SNP has two values and there are two disease states). An *association function* A is a function that assigns a positive score to a contingency table. Typical examples of association functions are the Pearson score and the Armitage trend statistic. We have used the Pearson score in our work; however, it is important to point out that our algorithm does not use any specific properties of the association function, apart from the property that the score is a function of the contingency table and the following symmetry property (which holds for the Pearson score):

$$A\left\{\begin{pmatrix} a & b \\ c & d \end{pmatrix}\right\} = A\left\{\begin{pmatrix} a & c \\ b & d \end{pmatrix}\right\}. \quad (1)$$

An *association score* is a function of the haplotype matrix and an arbitrary disease vector e (a binary vector of dimension $2n$) and is defined by $S(H, e) = \max_j A[\Omega(H, j, e)]$, i.e., the value of the association function at the most associated locus.

The goal is to calculate the significance of a pair (H, e) , which is defined as the probability of obtaining an association score at least as large as $S(H, e)$ under a null model. Formally, if e is a random disease vector, the p value is $\Pr[S(H, e) \geq S(H, d)]$. In addition, we want to accurately find the location of the associated SNPs in the genome. In our case, the null model is defined according to the randomization model in which $e = \pi(d)$ is a permutation of the disease vector and all instances $\pi(d)$ are equiprobable.


In addition to testing all the SNPs of H , we also test selected SNP interactions. For a set of SNPs j_1, j_2, \dots, j_k , suppose that there are Ψ different haplotypes induced by these SNPs, and let $h(j_1, j_2, \dots, j_k)$ be the $2n$ -dimensional haplotype vector induced by these SNPs (so that each element of $h(j_1, j_2, \dots, j_k)$ is an integer between zero and $\Psi - 1$). A *combinatorial interaction* of the SNPs j_1, j_2, \dots, j_k is a binary vector v that has dimension $2n$ and corresponds to a partition of the Ψ haplotypes into two sets S_1 and S_2 . Formally, let S_1, S_2 be disjoint sets of integers such that $S_1 \cup S_2 = \{0, 1, 2, \dots, \Psi - 1\}$. Then, $v(i) = 0$ if $h(j_1, j_2, \dots, j_k) \in S_1$, and $v(i) = 1$ if $h(j_1, j_2, \dots, j_k) \in S_2$.

The Perfect Phylogeny Tree

Our method is based on the construction of a perfect phylogeny tree and a specific choice of interactions among the SNPs based on the tree. A perfect phylogeny tree is a genealogical tree in which every node corresponds to a haplotype and every edge corresponds to a mutation (Figure 1). In a perfect phylogeny tree, we assume no recombination events and no recurrent mutations. Thus, such a genealogy is equivalent to the coalescent tree with the infinite-site assumption.

In a perfect-phylogeny model, every pair of SNPs satisfies the four-gamete test. Formally, for two SNP vectors $H_{\cdot, i}, H_{\cdot, j}$ we consider the haplotype counts $C_{a,b}(i, j) = |\{H_{x,i}, H_{x,j} | H_{x,i} = a, H_{x,j} = b\}|$. For example, $C_{0,0}(i, j)$ is the number of haplotypes in which SNPs i and j both equal 0. We say that the pair of SNPs

A



B




Figure 1. An Example of a Perfect Phylogeny Tree

Each node corresponds to a haplotype. The mutations appear on the edges. (A) A perfect phylogeny with five SNPs. (B) An additional sixth SNP that was mutated and can be expressed as the interaction between SNPs 4 and 5.

(i, j) satisfies the four-gamete test if there exists at least one pair (a, b) for which $C_{a,b}(i, j) = 0$.

The Algorithm

The intuition for our method is based on the case where the data are consistent with the perfect-phylogeny model. Our algorithm can be applied also to cases where there are deviations from the perfect phylogeny, as discussed in the section [Handling Recombination Events](#) below. The algorithm can be outlined as follows:

1. Build a perfect-phylogeny tree.
2. Select all pairs of SNPs that correspond to adjacent edges in the tree (edges that share a common vertex).
3. For each selected pair of SNPs, add a combinatorial interaction vector to the haplotype matrix H as a column.
4. Perform an association test by using the augmented haplotype matrix H .

The newly added columns represent putative unobserved SNPs that are plausible given the observed SNPs. In the algorithm described above, we added all pairwise interactions of SNPs but no higher-order interactions of SNPs. It is straightforward to extend this algorithm to also test higher orders of interactions, i.e., haplotypes with more than two SNPs. We restrict the discussion in this paper to pairwise interactions because we observed experimentally that higher-order interactions did not attain statistically significant improvement of the power (data not shown). We note that in the case of pairwise interactions, the association tests are practically done on haplotypes; the extension of this method to higher-order interactions cannot, however, be expressed as a haplotype test.

The algorithm finds the value of the association function for each SNP and each interaction, along with the corresponding association score. We use a permutation test to determine the significance of this score (corrected for multiple hypotheses). Because permutation tests can be quite inefficient, we use an importance sampling method for efficient calculation of the permutation test.⁴ Note that the algorithm we use is generic and that we could use any test for association for each of the interactions (e.g., a two-by-two χ^2 test or a three-by-two trend test).

Even though the above algorithm is quite simple, it is not immediately clear where the gain in power comes from. In the remainder of this section, we will describe the rationale for the algorithm. In order to do so, we will begin with the case where the perfect-phylogeny model is consistent with the data. We will explain later how we deal with deviations from the perfect-phylogeny model. We begin by describing in detail the process for adding combinatorial interaction vectors and the interpretation of this process as imputed unobserved SNPs.

Selecting the SNP Interactions

Each edge of a perfect phylogeny corresponds to a mutation in some SNP. Contracting an edge in the tree corresponds to the removal of the SNP associated with the edge from the dataset. Thus, every unobserved SNP corresponds to a contracted edge. More generally, we can view a perfect phylogeny on a set of observed SNPs as the result of a series of edge contractions on a larger perfect phylogeny determined by both observed and unobserved SNPs. It follows that the effect of adding an unobserved SNP to a perfect phylogeny must be to reverse an edge contraction; i.e., to split a node into two copies and insert an edge joining the two copies.

Every putative unobserved SNP that our algorithm constructs corresponds to such an edge insertion. Here, we limit ourselves to the simplest kind of edge insertions: those resulting from the interaction between a pair of observed SNPs corresponding to adjacent edges of the perfect phylogeny.

Any two observed SNPs correspond to edges in the tree, and the deletion of those edges induces three subtrees, which correspond to three different joint values of the two SNPs. For instance, in Figure 1A, the deletion of SNPs 1 and 5 induces three subtrees, where the first contains the haplotypes $S_1 = \{11000, 10000, 10100\}$, the second contains the haplotypes $S_2 = \{00000, 00010\}$, and the third contains the haplotype $S_3 = \{00011\}$. By our definition of an interaction (a partition of the set of haplotypes), an interaction between the two SNPs corresponds to a partition of the haplotypes into a set S_k versus the rest of the haplotypes. Thus, there are three possible interactions defined by a pair of SNPs (i, j). However, two of the three interactions correspond to testing one of the SNPs i or j . For instance, in the case of the pair (1,5) described above, testing the interaction ($S_3, S_1 \cup S_2$) is equivalent to testing SNP 5. In general, under the perfect-phylogeny assumption, every pair of SNPs has at most one nontrivial interaction that does not correspond to testing one of the SNPs. This can be shown by case analysis of all possible interactions of SNPs in such a scenario. When SNPs i and j correspond to adjacent edges, the nontrivial interaction corresponds to splitting a node and inserting an edge between the two copies. Indeed, our algorithm imputes precisely the unobserved SNPs that correspond to nontrivial interactions between adjacent edges.

Consider for example the case presented in Figure 1. Assume that SNPs 1, ..., 5 are genotyped and that SNP 6 is the causal SNP. In this case, testing the interaction between SNPs 4 and 5 in the original tree is equivalent to testing SNP 6. Similarly, testing the interactions

between SNPs 2 and 3 is equivalent to testing potential causal SNPs that has mutated after SNP 1 has mutated but before SNPs 2 and 3 have mutated. The CAMP algorithm restricts the set of tested interactions to interactions that correspond to such cases.

We note that there are other edge insertions that are not induced by the interaction of two SNPs. For instance, in the case of a starlike perfect phylogeny in which every leaf is adjacent to the root, any subset of the SNPs may correspond to a mutation that occurred after the root but before this set of SNPs. In CAMP, we do not consider such higher-order combinatorial interactions, although in theory they may potentially increase power.

The number of tests performed by our algorithm can be quadratic in the number of SNPs (e.g., if the perfect-phylogeny tree is a star). However, in practice the great majority of pairs of edges will not be adjacent. In particular, if the tree is degree bounded (i.e., the maximum number of edges that touch one vertex is below some constant number), the number of imputed unobserved SNPs will be linear in the number of observed SNPs.

Handling Recombination Events

The algorithm that we have described thus far is based on the perfect-phylogeny model, which assumes no recombination events. We now describe a modification of our algorithm that pulls back from this simplifying assumption and attempts to partially account for recombination events. One may view this modification as an approximation of the perfect-phylogeny model.

In the modified algorithm, in place of the perfect-phylogeny tree we instead construct a *perfect-phylogeny graph*. Each node in this graph represents a SNP. The edges in the graph are directed and are defined below.

There are two possible relationships between SNPs in the perfect phylogeny: (1) The haplotype 00 can have two descendant haplotypes, 01 and 10, which corresponds to a *brotherhood* relation between the two SNPs; or (2) the haplotype 00 can have a descendant 01, which has a descendant 11, which corresponds to a *parenthood* relation between the two SNPs.

We say that two SNPs i and j are in a brotherhood relation if $C_{0,0}(i,j)C_{1,1}(i,j) < C_{0,1}(i,j)C_{1,0}(i,j)$; otherwise, these SNPs are in a parenthood relation. If SNPs i and j are in a parenthood relation, then i is defined to be an *ancestor* of j if $C_{1,0}(i,j) > C_{0,1}(i,j)$. It is easy to see that, in the case of a perfect phylogeny in which the root is the haplotype for which all alleles have value 0, this definition agrees precisely with the notion of ancestry in the phylogeny tree. Similarly, if i and j are in a brotherhood relation, then neither of them is an ancestor of the other. We now define the edges in the perfect-phylogeny graph as follows: There is a directed edge from vertex v_i to v_j if v_i is an ancestor of v_j and there is no other vertex v_x such that v_i is an ancestor of v_x and v_x is an ancestor of v_j . Such a graph can be built via a topological sorting of the vertices.

In this construction we assume that the root of the tree is the haplotype for which all alleles have value 0; we can justify this by rooting the tree in one of the existing haplotypes and renaming the alleles of each SNP so that the root will satisfy this assumption.

Similarly to the way we perform the original algorithm described in the section titled *The Algorithm*, we test the interaction of two SNPs if they have a common parent in the perfect-phylogeny graph or if one of them is the parent of the other. In constructing the perfect-phylogeny graph, we do not consider relations of pairs of SNPs whose physical separation in the genome is higher than some threshold c . We call this threshold the *linkage upper bound*.

Observe that if there are no recombination events, the modified algorithm described in this section is equivalent to the algorithm described in the section titled [The Algorithm](#). In a perfect phylogeny, at least one of the four $C_{0,0}(i,j)$, $C_{1,1}(i,j)$, $C_{0,1}(i,j)$, or $C_{1,0}(i,j)$ equals zero (i.e., the four-gamete test holds), whereas here we do not require this property in order to decide the relationship of the SNPs.

Results

Data Sets

In order to test our approach, we needed a large data set that contains a sequence of several megabases for thousands of individuals. Currently, such a data set does not exist, and therefore we generated simulated population data as follows. We used the SNPs obtained from the HapMap data set as a starting point. To amplify these data, we assumed a fixed population size of 10,000, a mutation rate of 10^{-8} and a recombination rate of 10^{-8} . In each generation, individuals are mated randomly to produce the next generation. The number of children generated by two individuals is a random variable with a predefined distribution. We used 30,000 generations to generate the final population sample. This process was done for 15 Mbp along one chromosome. Note that we did not use an approximation-based approach (such as the coalescent model with recombination events or the Li and Stephens model¹⁵) to simulate the population, but rather we used an explicit forward simulation of the population, and this simulation initiates from a real data set.

We used a multiplicative model to generate samples of cases and controls. We simulated experiments with 1000 cases and 1000 controls. A *panel* is defined to be one experiment. For each panel, a SNP was randomly chosen to be the causal SNP and was then removed from the panel. We set the disease prevalence to 0.01 and the relative risk to 1.5. We set the linkage upper bound (c) to 50 kb, which has been shown to be a good estimate in humans (e.g.,⁴). We used the perfect-phylogeny-graph algorithm described in the section titled [Handling Recombination Events](#).

The running time of CAMP for 500 cases and 500 controls for 38,864 SNPs on chromosome 1 (corresponding to the Affymetrix SNP chip) on a Sun workstation (with a Quad 2.4 GHz AMD Opteron 850 Processor) is 4 min to calculate the scores for each SNP and an additional 18 min for a standard-permutation test.

Evaluation of Previous Methods

Many of the existing genealogy-based methods are computationally inefficient, and thus a large-scale evaluation of these methods is prohibitive. Our experiments involved thousands of panels, each containing thousands of haplotypes with thousands of SNPs, and thus we concentrated on the evaluation of methods that are efficient enough to handle data sets of this size. In particular, coalescent-based methods such as LATAG⁵ are not computationally feasible for large-scale data sets. The Margarita⁶ algorithm

is also too computationally intensive: It took more than two weeks to analyze a data set of 500 cases and controls with 10,000 SNPs (we used the recommended parameters by the developers: 30 ARGs and 100 permutations).

We did, however, test the power of Margarita on a small number of SNPs and individuals. To our surprise, the power of Margarita was not as good as the standard χ^2 test under several of the scenarios that we studied; because Margarita assumes the coalescent model, we tested it on null data produced with a coalescent model, upon which phenotypes were randomly assigned (i.e., in which there is no causal SNP). We generated 100 different panels of size 25 kb by using the *ms* software (Hudson), which simulates data under the coalescent model. In these experiments we generated 50 individuals, who were randomly assigned to be cases or controls. We found that the false-positive rate (using a p value cutoff of 0.05) was 1%. We used the following strong association model: We arbitrarily set one of the SNPs as a strong causal SNP by declaring an individual to be a case if the corresponding SNP was either heterozygous or homozygous 1 and otherwise declaring the individual to be a control. Using this model, we observed that the power of Margarita is 17%, whereas the power that is obtained by a standard-permutation test was much higher: 69%.

We repeated the same experiments described above by simulating additional panels of SNPs by using the *ms* software with recombination events. We found that the false-positive rate in this case was 89%. The power for the strong-association model was 69%, versus 75% obtained by the standard-permutation test.

These results show that there are serious problems with existing methods either with respect to running time or power. Indeed, our results show that these methods are dominated by the standard χ^2 approach in the scenarios that we studied. We thus used the standard χ^2 approach as a baseline in our experimental study of CAMP. We also compared CAMP to CLADHC¹⁰ because CLADHC is computationally feasible in our scenarios.

Power

To study the power of the methods, we applied them to case-control panels by using the multiplicative model for generating cases and controls as described in the section titled [Data Sets](#). A panel is generated from the whole sequence, and then the SNPs are sampled according to the specified sparsity (i.e., density of the sampled SNPs). Therefore, in many cases (but not always) the causal SNP does not exist in the SNP-sampled panel. Under a specific significance level (which controls the type I error), the *relative power* is defined as a ratio in which the numerator is the number of SNP-sampled panels defined to be significant and the denominator is the number of original panels with the entire sequence, including the causal SNP, defined to be significant.

We simulated panels with 1,000 cases and 1,000 controls. We fixed the sparsity of the sampled SNPs to each

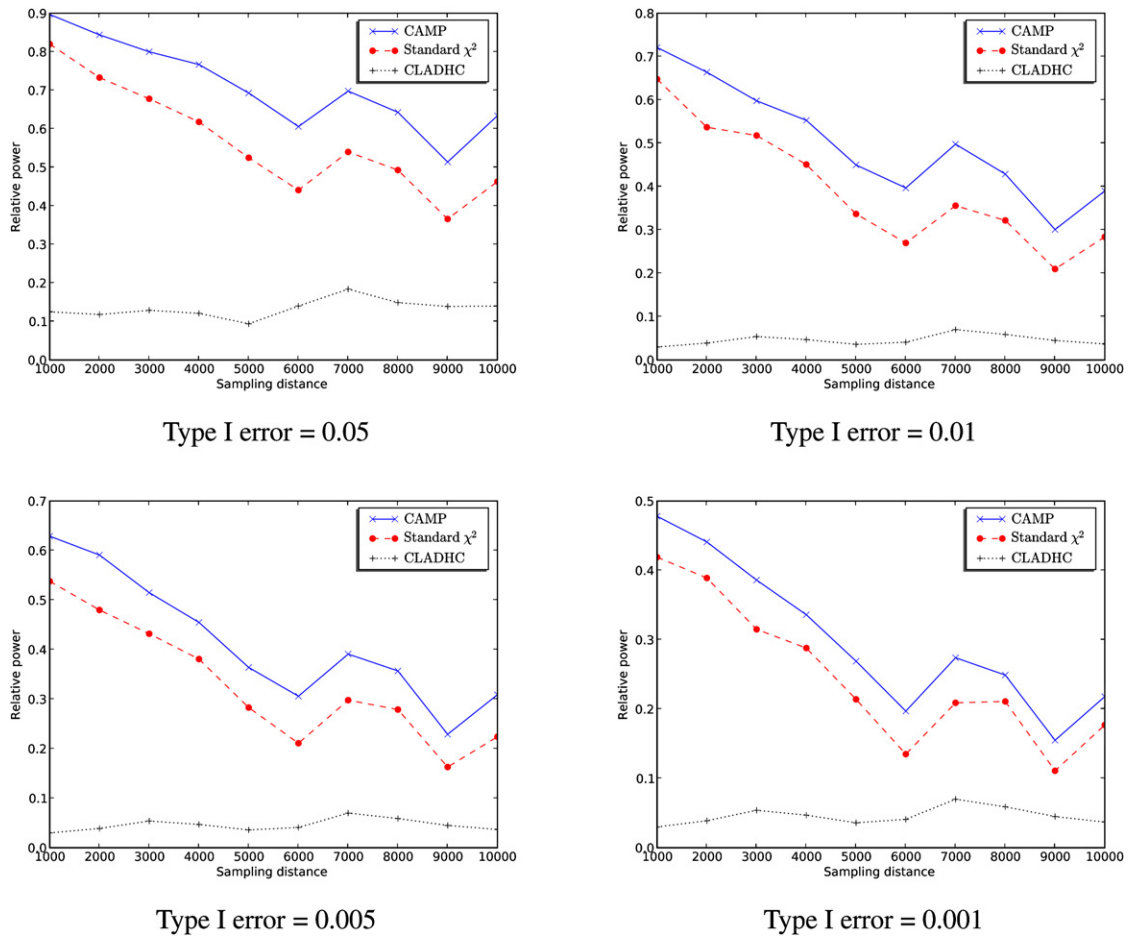


Figure 2. The Relative Power for Different Sampling Distances of SNPs and for Four Different Significance Levels

of the following values: (1,000, 2,000, ..., 10,000). For each value of sampling sparsity, 5,000 different panels were generated. Results for the different SNP sampling densities are presented in Figure 2, and results for different numbers of cases and controls are presented in Figure 3. The difference in the relative power between CAMP and standard χ^2 testing reaches more than 10%. An even more prominent difference is observed between CAMP and CLADHC and ranges up to 52.7%.

Localization of the Causal SNP

We also tested the accuracy in localizing the causal SNP. To obtain an estimate of location from the output of CAMP, note first that CAMP may output more than one SNP. These SNPs presumably represent mutations on the genealogical tree near the causal SNP; thus, we use the average location of the SNPs found by CAMP to estimate the position of the causal SNP. A comparison to the χ^2 test is presented in Figure 4 for different numbers of individuals. The advantage of CAMP over the standard χ^2 is quite notable; e.g., for 3,000 controls and cases, the percentage of panels for which the distance between the found location and the true causal SNP is below 100 Kb was 86% for CAMP and 79% for the standard χ^2 .

Measuring the Advantage of the Coalescent Approach

Our algorithm tests a subset of all possible interactions. This subset, as described before, is determined according to the approximated genealogical relations between the SNPs. Does testing all possible interactions within the linkage upper bound give similar results? To answer this

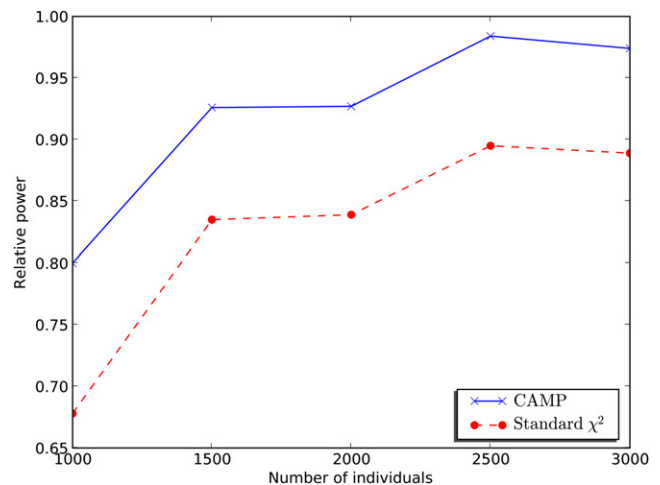


Figure 3. Relative Power for Different Numbers of Individuals

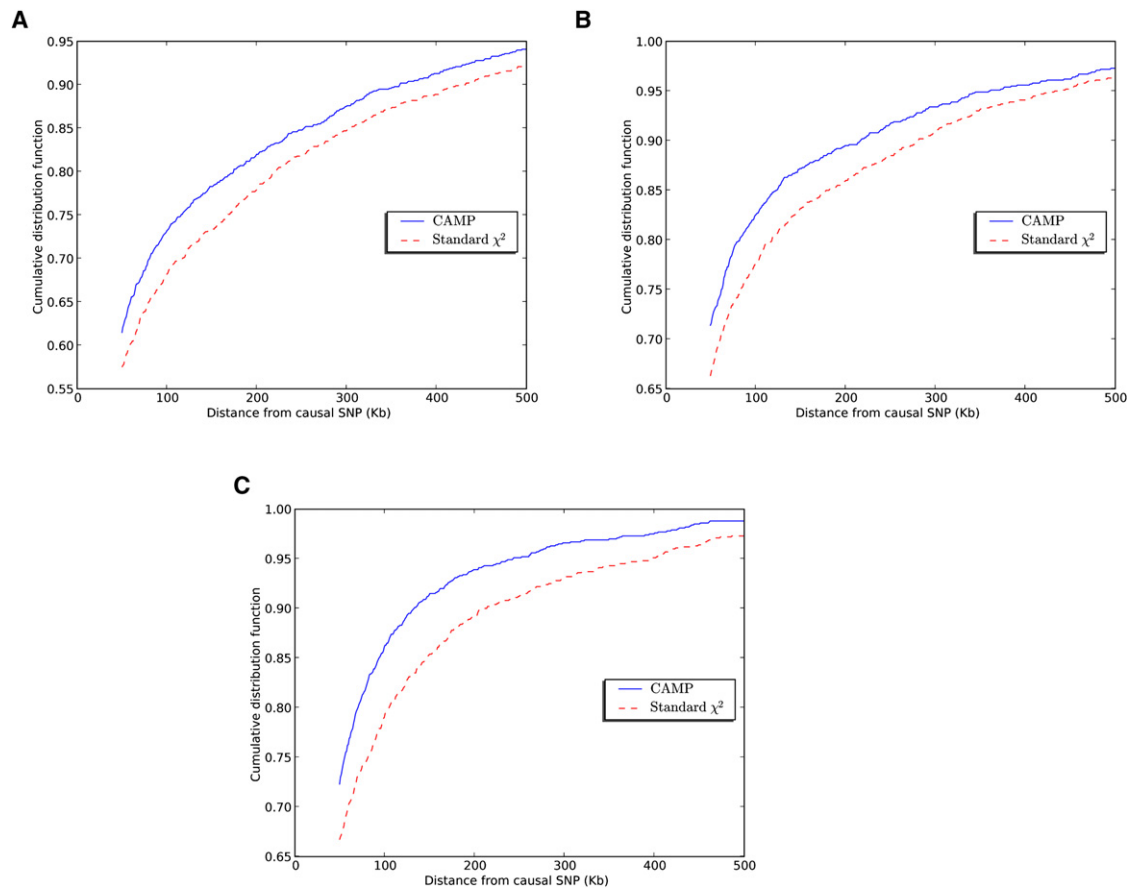


Figure 4. A comparison of the Cumulative Distribution Functions for the Distance between the Discovered SNP and the True Causal SNP

The three figures represent the results obtained from different numbers of controls and cases: A - 1,000, B - 2,000, C - 3,000.

question, we compared CAMP to a procedure that tests all pairwise interactions within the linkage upper bound. For two SNPs, this corresponds to testing the association between the haplotypes generated by the SNPs and the phenotype, which is calculated by a standard χ^2 test for these two vectors.

The results are presented in Figure 5. As can be seen, CAMP yields significantly greater power than the procedure that tests all pairwise interactions. Reaching more than 50%, the difference between these procedures is much larger than the difference between CAMP and the standard χ^2 algorithm.

Using the Phased Haplotypes

Because CAMP uses phased data to construct the coalescent graph, we tested the effect of phasing errors on our algorithm by considering phasing error rates of 3%, 30%, and 50%. The value of 3% corresponds roughly to the error rate reported in the literature for phasing algorithms¹⁶, and 50% corresponds to randomly phasing each one of the heterozygous sites.

Results are presented in Figure 6. As can be observed in the graph, even when the phasing error rate is 50% (which

is very unlikely), CAMP has a relative power that is 9% larger than the standard χ^2 test (for a significance level of 5%). With phasing error rates of 3% and even 30%, no significant reduction in the power is observed.

Discussion

We have presented a method that leverages the coalescent model to conduct association mapping in whole-genome association studies. We exploit the unobserved genealogy of the chromosomes in order to evaluate more accurately the significance and location of causal SNPs. The genealogy defines a set of haplotypes, and our method consists of a strategy for selecting tests on the basis of these haplotypes and the genealogy. As we have demonstrated, selecting these tests carefully gives a large advantage in the power and in the localization of the causal SNP. We have also shown that several existing methods that aimed to address this problem either suffer from low power or suffer from a very high false-positive rate when they are compared to a standard approach in which each of the SNPs is tested separately with a χ^2 test. We have also shown

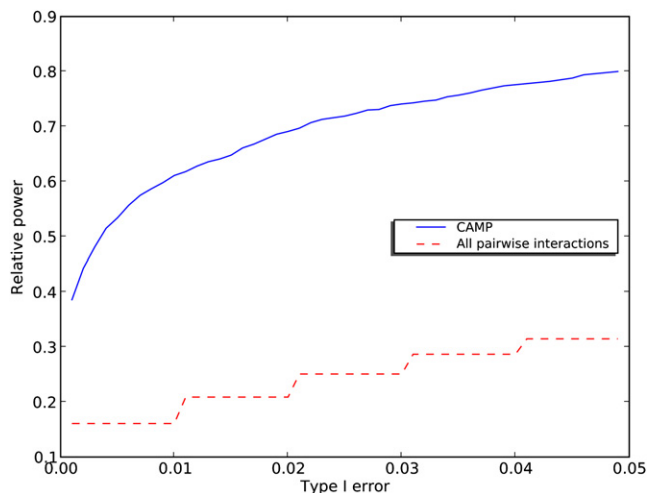


Figure 5. A Comparison of CAMP to a Naive Pairwise-Interaction Algorithm, in Which All Pairwise Interactions of SNPs with Distances Smaller Than the Threshold Used in CAMP Are Tested

that considering all SNP interactions reduces the power considerably.

Interestingly, we observed that introducing very high rates of phasing errors (30%) does not reduce the power of our method. This can be explained by the fact that when enough individuals are given, the genealogical relationship between SNPs can be determined accurately, even if the heterozygous sites are ignored. The signal is strong enough in the homozygous sites so that phasing accuracy has a minor effect on the results.

Accurate localization of the causal SNP is as important as significance estimation. We have shown that CAMP estimates the location of the real causal SNP more accurately than other methods. Note, moreover, that this was achieved by a relatively naive approach of taking the average of the interacting SNPs. Most likely, the location can be determined even more accurately by a more sophisticated algorithm that uses properties of the coalescent.

There are several important issues that need further attention. We have shown that CAMP is more powerful than the standard χ^2 test, but we have not shown its optimality. The question of whether there exists a more powerful algorithm or strategy for choosing a subset of interactions of SNPs should be explored. In particular, it is intriguing to study the following optimization problem: For a fixed disease model (say, the multiplicative model with given penetrance and relative risk) and significance level, find the strategy that determines which SNP interactions are tested, such that the power is maximized.

It is also important to study the generalizations of methods such as CAMP to the case where multiple populations may be participating in the study. One of the challenges in drawing causal inferences from whole-genome case-control association studies is the confounding effect of population structure.^{17–24} This issue has received much attention in recent years in the literature (e.g.,^{24,25}). Currently, CAMP assumes no stratification effect, i.e., the




Figure 6. A Comparison of the Relative Power of CAMP in the Presence of Different Rates of Phasing Errors

controls and cases are presumed to be from one population. There is a clear need to explore methods for taking population stratification into account in CAMP.

Because CAMP performs transformation on the SNP data, it can be naturally extended to handle other types of phenotypes such as continuous traits (QTLs), although this was not tested experimentally in this work.

The cost of genotyping is continually decreasing, and technology is evolving toward genome resequencing (e.g., Illumina/Solexa 1G and Roche/454). However, it is still very expensive to conduct resequencing of the whole genome as a tool for association studies. Consequently, it is clear that genotype data of common genetic variants such as SNPs will be the leading approach in association studies in the coming years. Algorithms such as CAMP that yield high statistical power by exploiting aspects of genealogy will play an important role in the analysis of these data.

Acknowledgments

G.K., E.H. and R.M.K. were supported by National Science Foundation grant IIS-0513599. E.H. was also supported by National Science Foundation grant IIS-0713254. M.I.J. was supported by a grant from Microsoft Research and by an appointment as a Miller Research Professor in the Miller Institute for Basic Research in Science.

Received: April 21, 2008

Revised: August 31, 2008

Accepted: October 21, 2008

Published online: November 20, 2008

Web Resources

The URLs for data presented herein are as follows:

Hapmap project, <http://www.hapmap.org>

Illumina/Solexa, <http://www.illumina.com>

Roche/454, <http://www.454.com>

ms – A program for generating samples under neutral models, <http://home.uchicago.edu/rhudson1/source/mksamples.html>

CAMP – coalescent-based association mapping (the algorithm described in this paper), <http://www2.icsi.berkeley.edu/kimmel/software/camp>

References

1. Wellcome Trust Case Control Consortium (2007). Genome-wide association study of 14,000 cases of seven common diseases and 3,000 shared controls. *Nature* 447, 661–683.
2. Bonnen, P.E., Pe'er, I., Plenge, R.M., Salit, J., Lowe, J.K., Shapero, M.H., Lifton, R.P., Breslow, J.L., Daly, M.J., Reich, D.E., et al. (2006). Evaluating potential for whole-genome studies in Kosrae, an isolated population in Micronesia. *Nat. Genet.* 38, 214–217.
3. Slager, S.L., Huang, J., and Vieland, V.J. (2000). Effect of allelic heterogeneity on the power of the transmission disequilibrium test. *Genet. Epidemiol.* 18, 143–156.
4. Kimmel, G., and Shamir, R. (2006). A fast method for computing high significance disease association in large population-based studies. *Am. J. Hum. Genet.* 79, 481–492.
5. Zollner, S., and Pritchard, J.K. (2005). Coalescent-based association mapping and fine mapping of complex trait loci. *Genetics* 169, 1071–1092.
6. Minichiello, M.J., and Durbin, R. (2006). Mapping trait loci by use of inferred ancestral recombination graphs. *Am. J. Hum. Genet.* 79, 910–922.
7. Nordborg, M. (2001). *Handbook of Statistical Genetics* (Chichester, West Sussex, UK: John Wiley and Sons).
8. Griffiths, R.C., and Marjoram, P. (1996). Ancestral inference from samples of DNA sequences with recombination. *J. Comput. Biol.* 3, 479–502.
9. Fearnhead, P., and Donnelly, P. (2001). Estimating recombination rates from population genetic data. *Genetics* 159, 1299–1318.
10. Durrant, C., Zondervan, K.T., Cardon, L.R., Hunt, S., and Deloukas, P. (2004). Linkage disequilibrium mapping via clastic analysis of single-nucleotide polymorphism haplotypes. *Am. J. Hum. Genet.* 75, 35–43.
11. Bafna, V., Gusfield, D., Lancia, G., and Yooseph, S. (2003). Haplotyping as perfect phylogeny: A direct approach. *J. Comput. Biol.* 10, 323–340.
12. Halperin, E., and Eskin, E. (2004). Haplotype reconstruction from genotype data using imperfect phylogeny. *Bioinformatics* 20, 1842–1849.
13. Kimmel, G., and Shamir, R. (2005). The incomplete perfect phylogeny haplotype problem. *J. Bioinform. Comput. Biol.* 3, 359–384.
14. Li, N., and Stephens, M. (2003). Modelling linkage disequilibrium and identifying recombinations hotspots using SNP data. *Genetics* 165, 2213–2233.
15. Kimmel, G., and Shamir, R. (2005). GERBIL: Genotype resolution and block identification using likelihood. *Proc. Natl. Acad. Sci. USA* 102, 158–162.
16. Freedman, M.L., Reich, D., Penney, K.L., McDonald, G.J., Mignault, A.A., Patterson, N., Gabriel, S.B., Topol, E.J., Smoller, J.W., Pato, C.N., et al. (2004). Assessing the impact of population stratification on genetic association studies. *Nat. Genet.* 36, 388–393.
17. Clayton, D.G., Walker, N.M., Smyth, D.J., Pask, R., Cooper, J.D., Maier, L.M., Smink, L.J., Lam, A.C., Ovington, N.R., Stevens, H.E., et al. (2005). Population structure, differential bias and genomic control in a large-scale, case-control association study. *Nat. Genet.* 37, 1243–1246.
18. Lander, E.S., and Schork, N.J. (1994). Genetic dissection of complex traits. *Science* 265, 2037–2048.
19. Helgason, A., Yngvadottir, B., Hrafnkelsson, B., Gulcher, J., and Stefansson, K. (2005). An Icelandic example of the impact of population structure on association studies. *Nat. Genet.* 37, 90–95.
20. Campbell, C.D., Ogburn, E.L., Lunetta, K.L., Lyon, H.N., Freedman, M.L., Groop, L.C., Altshuler, D., Ardlie, K.G., and Hirschhorn, J.N. (2005). Demonstrating stratification in a European American population. *Nat. Genet.* 37, 868–872.
21. Hirschhorn, J.N., and Daly, M.J. (2005). Genome-wide association studies for common diseases and complex traits. *Nat. Rev. Genet.* 6, 95–108.
22. Lohmueller, K.E., Pearce, C.L., Pike, M., Lander, E.S., and Hirschhorn, J.N. (2003). Meta-analysis of genetic association studies supports a contribution of common variants to susceptibility to common disease. *Nat. Genet.* 33, 177–182.
23. Price, A.L., Patterson, N.J., Plenge, R.M., Weinblatt, M.E., Shadick, N.A., and Reich, D. (2006). Principal components analysis corrects for stratification in genome-wide association studies. *Nat. Genet.* 38, 904–909.
24. Kimmel, G., Jordan, M., Halperin, E., Shamir, R., and Karp, R. (2007). A randomization test for controlling population stratification in whole-genome association studies. *Am. J. Hum. Genet.* 81, 895–905.
25. Devlin, B., and Roeder, K. (1999). Genomic control for association studies. *Biometrics* 55, 997–1004.