The Clark Phase-able Sample Size Problem: Long-range Phasing and Loss of Heterozygosity in GWAS*

Bjarni V. Halldórsson [1,3,4]
Derek Aguiar [1,2]
Ryan Tarpine [1,2]
Sorin Istrail [1,2]

[1] *Center for Computational Molecular Biology, Brown University*
[2] *Department of Computer Science, Brown University*
[3] *Department of Biomedical Engineering, Reykjavik University*
[4] *deCODE genetics*

Abstract

A phase transition is taking place today. The amount of data generated by genome resequencing technologies is so large that in some cases it is now less expensive to repeat the experiment than to store the information generated by the experiment. In the next few years it is quite possible that millions of cases and controls will have been genotyped. The premise of the paper is that long shared genomic regions (or tracts) are unlikely unless the haplotypes are identical by descent (IBD), in contrast to short shared tracts which may be identical by state (IBS). Here we estimate for populations, using the US as a model, what sample size of genotyped individuals would be necessary to have sufficiently long shared haplotype regions (tracts) that are identical by descent (IBD), at a statistically significant level. These tracts can then be used as input for a Clark-like phasing method to obtain a complete phasing solution of the sample. We estimate in this paper that for a population like the US and about 1% of the people genotyped (approximately 2 million), tracts of about 200 SNPs long are shared between pairs of individuals IBD with high probability which assures the Clark method phasing success. We show on simulated data that the algorithm will get an almost perfect solution if the number of individuals being SNP arrayed is large enough and the correctness of the algorithm grows rapidly with the number of individuals being genotyped. We also study a related problem that connects copy number variation with phasing algorithm success. A loss of heterozygosity (LOH) event is when, by the laws of Mendelian inheritance, an individual should be heterozygote but, due to a deletion polymorphism, is not. Such polymorphisms are difficult to detect using existing algorithms, but play an important role in the genetics of disease and will confuse haplotype phasing algorithms if not accounted for. We will present an algorithm for detecting LOH regions across the genomes of thousands of individuals. The design of the long-range phasing algorithm and the Loss of Heterozygosity inference algorithms was inspired by analyzing of the Multiple Sclerosis GWAS dataset of the International Multiple Sclerosis Consortium and we present in this paper some results of this analysis.

1 Introduction

Genome-wide association studies (GWAS) proceed by identifying a number of individuals carrying a disease or a trait and comparing these individuals to those that do not or are not known to carry the disease/trait. Both sets of individuals are then genotyped for a large number of Single Nucleotide Polymorphism (SNP) genetic variants which are then tested for association to the disease/trait. GWAS have been able to successfully identify a very large number of polymorphism associated to disease ([17, 2] etc.) and the amount of SNP data from these studies is growing rapidly. Studies

*corresponding authors are Bjarni V. Halldórsson bjarnivh@ru.is and Sorin Istrail sorin@cs.brown.edu
Authors are members of the International Multiple Sclerosis Genetics Consortium GWAS Analysis team.

using tens of thousands of individuals are becoming commonplace and are increasingly the norm in the association of genetic variants to disease [3, 17, 11]. These studies generally proceed by pooling together large amounts of genome-wide data from multiple studies, for a combined total of tens of thousands of individuals in a single meta-analysis study. It can be expected that if the number of individuals being genotyped continues to grow, hundreds of thousands, if not millions, of individuals will soon be studied for association to a single disease or trait.

SNPs are the most abundant form of variation between two individuals. However, other forms of variation exist such as copy number variation - large scale chromosomal deletions, insertions, and duplications (CNV). These variations, which have shown to be increasingly important and an influential factor in many diseases [15], are not probed using SNP arrays. A further limitation of SNP arrays is that they are designed to probe only previously discovered, common variants. Also, the SNPs must be well behaved: there are numerous gaps in the HapMap collection of 3.5 million SNPs that have no SNPs listed. In some cases no SNP assays were attempted because the region was a low or high copy-repeat region; in other cases, SNP assays tested by the HapMap project resulted in genotypes in the CEU samples but the SNP was listed as QC-negative because of inheritance errors or lack of Hardy-Weinberg equilibrium. These regions have been termed *SNP deserts*. In total these SNP deserts cover approximately 10% of the human genome.

In summary, there are two main reasons why current genome-wide associations studies have not reached their full potential: first, the sample size of genotypes is low and the number of individuals being genotyped still needs to grow; and secondly, characterization of variation patterns is too simple, that is, structural variation and rare SNPs are ignored.

The future direction of genetic association studies are mainly twofold: the testing of more individuals using genome-wide association arrays and the resequencing of a small number of individuals with the goal of detecting more types of genetic variations, both rare SNPs and structural variation [14]. Testing multiple individuals for the same variants using standard genome-wide association arrays is becoming increasingly common and can be done at a cost of approximately $100 per individual. This low cost implies that genome-wide association SNP array testing of individuals will soon become commonplace as the cost of acquiring the sample becomes less than the cost of storing the sample. In the next couple of years it is plausible that several million individuals in the US population will have had their genome SNP arrayed.

Whole genome resequencing is currently in its infancy. A few people have had their genome resequenced and the cost of sequencing a single individual is still estimated in the hundreds of thousands of dollars. However, whole genome sequencing is preferable for association studies as it allows for the detection of all genomic variation and not only SNP variation.

Due to the fact whole genome SNP arrays are becoming increasingly abundant and whole genome resequencing is still quite expensive, the question has been raised whether it would suffice to whole genome sequence a small number of individuals and then impute [5] other genotypes using SNP arrays and the shared inheritance of these two sets of individuals. It has been shown that this could be done most efficiently using the shared haplotypes between the individuals that are SNP arrayed and those that have been resequenced [8].

Problem. Current technologies, suitable for large-scale polymorphism screening only yield the genotype information at each SNP site. The actual haplotypes in the typed region can only be obtained at a considerably high experimental cost or via haplotype phasing. Due to the importance of haplotype information for inferring population history and for disease association, the development of algorithms for detecting haplotypes from genotype data has been an active research area for several years [1, 13, 16, 12, 8, 4]. However, algorithms for determining haplotype phase are still in their infancy after about 15 years of development (e.g. [1, 16, 7]). Of particular worry is the fact that the learning rate of the algorithm, i.e. the rate that the algorithms are able to infer more

correct haplotypes, grows quite slowly with the number of individuals being SNP arrayed.

Solution. In this paper we present an algorithm for the phasing of a large number of individuals. We show that the algorithm will get an almost perfect solution if the number of individuals being SNP arrayed is large enough and the correctness of the algorithm grows rapidly with the number of individuals being genotyped. We will consider the problem of haplotype phasing from long shared genomic regions (that we call tracts). Long shared tracts are unlikely unless the haplotypes are identical by descent (IBD), in contrast to short shared tracts which may be identical by state (IBS). We show how we can use these long shared tracts for haplotype phasing.

Problem. We further consider the problem of detecting copy number variations from whole genome SNP arrays. It is important that a phasing algorithm does not get confused when it encounters deviations from normal patterns of inheritance such as may be observed when there is a deletion. A loss of heterozygosity (LOH) event is when, by the laws of Mendelian inheritance, an individual should be heterozygote but due to a deletion polymorphism, is not. Such polymorphisms are difficult to detect using existing algorithms, but play an important role in the genetics of disease [15] and will confuse haplotype phasing algorithms if not accounted for.

Solution. We provide an exact exponential algorithm and a greedy heuristic for detecting LOH regions.

We run empirical tests and benchmark the algorithms on a GWAS dataset from the International Multiple Sclerosis Genetics Consortium [2] as well as simulated data derived using the Hudson simulator [6]. To determine LOH events we assume the data is given in trios, i.e. the genotypes of a child and both its parents are known.

2 Long range phasing and haplotype tracts

The haplotype phasing problem asks to computationally determine the set of haplotypes given a set of individual's genotypes. We define a *haplotype tract* (or *tract* for short) denoted $[i, j]$ as a sequence of SNPs that is shared between at least two individuals starting at the same position i in all individuals and ending at the same position j in all individuals. We show that if we have a long enough tract then the probability that the sharing is IBD is close to 1. Multiple sharing of long tracts further increases the probability that the sharing corresponds to the true phasing.

2.1 Probability of observing a long tract

We show that as the length of the tract increases the probability that the tract is shared IBD increases. Let t be some shared tract between two individual's haplotypes and l be the length of that shared tract. We can then approximate the probability this shared tract is identical by state (IBS) $p_{IBS}(l)$. Let $f_{M,i}$ be the major allele frequency of the SNP in position i in the shared tract t. Assuming the Infinite Sites model and each locus is independent,

$$p_{IBS}(l) = \prod_{i=1}^{l} \left((f_{M,i})(f_{M,i}) + (1 - f_{M,i})\,(1 - f_{M,i}) \right)$$

We can approximate $p_{IBS}(l)$ by noticing $f_{M,i} * f_{M,i}$ dominates $(1 - f_{M,i})(1 - f_{M,i})$ as $f_{M,i} \to 1$, $p_{IBS}(l) \approx \prod_{i=1}^{l} (f_{M,i})^2$. Let f_{avg} be $\frac{1}{l} f_{M,i}\ \forall i \in t$. Then $p_{IBS}(l) \approx (f_{avg})^{2l}$. Given $f_{M,i}$ is some high frequency, say 95%, then a sharing of 100 consecutive alleles is very unlikely, $p_{IBS}(100) \approx 0.95^{200} = 10^{-5}$. For very large datasets we will need to select the length of the tract being considered to be large enough so that the probability that the sharing is identical by state is small.

The probability two individuals separated by $2(k+1)$ meiosis (kth-degree cousins) share a locus IBD is 2^{-2k} [8]. As k increases, the probability kth-degree cousins share a particular locus IBD decreases exponentially. However, if two individuals share a locus IBD then they are expected to share about $\frac{200}{2k+2}$ cM [8]. Relating $P(IBD)$ to length of tract l,

$$P(IBD|sharing\ of\ length\ l) = \frac{2^{-2n}}{2^{-2n} + \left((f_{M,i})^{2l} + (1 - f_{M,i})^{2l}\right)}$$

which is shown in Figure 1.

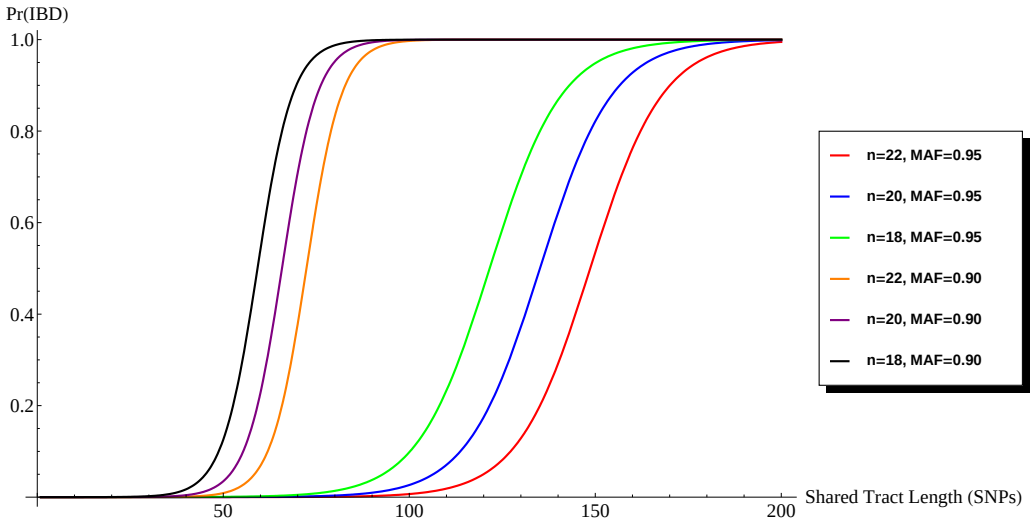

Figure 1: *Probability of IBD as a function of shared tract length (measured in SNPs) and plotted for several n and major allele frequencies (MAF). n is the number of meiosis between the two individuals. The smaller the MAF or n the faster $P(IBD)$ converges to 1.*

2.2 The Clark phase-able sample size problem

Given the large tract sharing, we can construct the *Clark consistency graph* having individuals as vertices and an edge between two individuals if they share a tract [13]. Figure 2 shows the Clark consistency graph for different *minimum significant tract lengths* (or window sizes) in the MS dataset. At what minimum significant tract lengths will the graph become dense enough so that phasing can be done properly? What percentage of the population needs to be genotyped so that the Clark consistency graph becomes essentially a single connected component? We call this "The Clark sample estimate: the size for which the Clark consistency graph is connected, C."

We computed the average number of edges in the haplotype consistency graph as a function of window size to get a sense when the Clark consistency graph of the MS data becomes connected. Based on Figure 3 and $P(IBD)$ we can propose an algorithmic problem formulation from the Clark consistency graph. Preferably we would like to solve either one of the below problems.

Problem 1 *Remove the minimum number of the edges from the Clark consistency graph so that the resulting graph gives a consistent phasing of the haplotypes.*

Problem 2 *Maximize the joint probability of all the haplotypes given the observed haplotype sharing.*

We believe that both of these problem formulations are NP-hard and instead propose to solve these problems using a heuristic. Our benchmarking on simulated data shows that this heuristic works quite well.

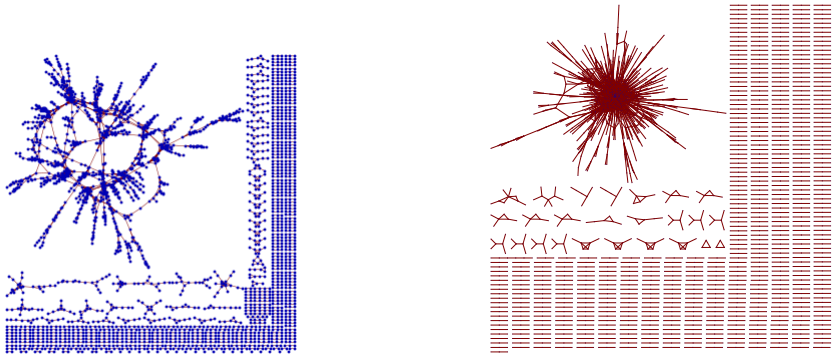

Figure 2: *Left: The Clark consistency graph for region [1400,1600). A large fraction of individuals share consistent haplotypes of length 200 suggesting many are IBD. Right: The Clark consistency graph for a smaller window size of 180 base pairs. We observe a more dense connected component in part due to the smaller windows size but also because of the specific genomic region.*

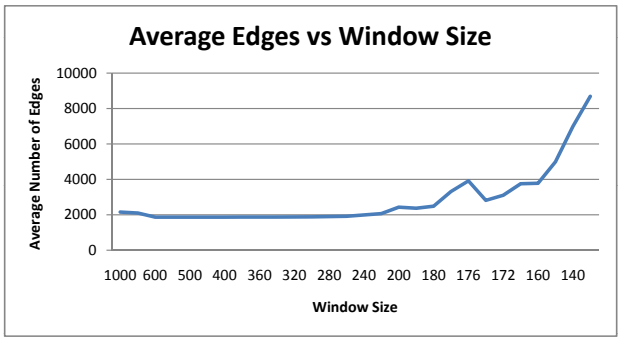

Figure 3: *The average number of edges per window size stays relatively constant until a window size of about 180. The graph becomes more connected at this point likely because the window size is small enough to not be largely affected by recombination (but still large enough for the shared tracts to not likely be IBS).*

2.3 Phasing the individuals that are part of the largest component

We now proceed with an iterative algorithm working on the connected components in the Clark haplotype consistency graph. First we construct the graph according to some length of haplotype consistency (Figure 3 and $P(IBD)$ help define this length). We iterate through each site of each individual to find the tracts. After we find a site with some long shared region, we look at its neighbors in the connected component and apply a voting scheme to decide what the value is for each heterozygous allele. After each individual has been processed we iterate with having resolved sites in the original matrix.

Observation 1 *If the Clark consistency graph is fully connected all edges are due to IBD sharing and all individuals can be perfectly phased up to the point were all individuals are heterozygote at a particular site.*

Therefore, phasing individuals in a connected component of the graph should be easy, but in practice there will be some inconsistencies for a number of reasons. If the a node in the Clark consistency graph has a high degree then the phasing of that node will be ambiguous if the neighbors

are not consistent. At some times this may be due to genotyping error and at times this may be due to identical by state sharing to either one or both of an individuals haplotypes. The identical by state sharing may because the haplotype has undergone recombination, possibly a part of the haplotype is shared identical by descent and a part is identical by state.

Our alphabet for genotype data is $\Sigma = \{0, 1, 2, 3\}$. 0s and 1s represent the homozygote for the two alleles of a SNP. A 2 represents a heterozygous site and a 3 represents missing data. Given a set of n-long genotype strings $G = \{g_1, g_2, \ldots, g_{|G|}\}$ where $g_i \in \Sigma^n$, we represent this in a matrix M with $m = 2\,|G|$ rows and n columns:

$$M = \begin{bmatrix} M_{1,1} & M_{1,2} & \cdots & M_{1,n} \\ M_{2,1} & M_{2,2} & \cdots & M_{2,n} \\ \vdots & \vdots & \ddots & \vdots \\ M_{m,1} & M_{m,2} & \cdots & M_{m,n} \end{bmatrix}$$

Each genotype g_i is represented by the two rows $2i-1$ and $2i$. Initially, $M_{2i-1,j} = M_{2i,j} = g_i[j]$.

We define allele consistency to be:

$$c(a,\, b) = \begin{cases} 1 & \text{if } a = b \text{ or } a \in \{2,\, 3\} \text{ or } b \in \{2,\, 3\} \\ 0 & otherwise \end{cases}$$

Rows r and s of M are consistent along a tract $[i,\, j]$ (i.e. have a shared tract) is written

$$C_{[i,\,j]}(r,\, s) = \prod_{k \in [i,\, j]} c\,(M_{r,k},\, M_{s,k})$$

The length of a tract is written $|[i,\, j]| = j - i + 1$.

A shared tract $[i,\, j]$ between rows r and s is *maximal shared tract* if it cannot be extended to the left or right; i.e., $i = 1$ or $c(M_{r,i-1},\, M_{s,i-1}) = 0$ and $j = n$ or $c(M_{r,j+1},\, M_{s,j+1}) = 0$. The maximal shared tract between rows r and s at position i is written $S_i^{r,s}$. It is unique. Note that if $S_i^{r,s} = [j,\, k]$ then $\forall_{l \in [j,\, k]} S_l^{r,s} = S_i^{r,s}$.

2.4 Tract finding and phasing algorithm

Given that there are some loci for which individuals share IBD and that these sharings are expected to be large, we developed an algorithm to detect and use these sharings to resolve the phase at heterozygous sites. Each site is resolved by determining if there are any other individuals that likely share a haplotype by descent. SNPs that do not have their phase determined during any given iteration will be processed in succeeding iterations. If there are enough long IBD loci, this algorithm should unambiguously determine the phase of each individual.

If we know that the data contains trios, a child and both of its parents, we start by phasing the trios using Mendelian laws of inheritance. This replaces many of the heterozygote sites (whenever at least one member of a family is homozygous) and even a few of the sites having missing data (i.e., when the parents are both homozygous and the child's genotype is missing).

To phase using long shared tracts, we start by fixing a minimum significant tract length L. We run several iterations, each of which generate a modified matrix M' from M, which is then used as the basis for the next iteration.

First, we set $M' := M$.

For each row r we examine position i. If $M_{r,i} \in \{0,\, 1\}$ then we move to the next i. Otherwise $M_{r,i} \in \{2,\, 3\}$, and we count "votes" for whether the actual allele is a 0 or 1.

$$V_0^r = |\{s \,|\, s \neq r \text{ and } |S_i^{r,s}| \geq L \text{ and } M_{s,i} = 0\}|$$

V_1^r is defined analogously (the difference being the condition $M_{s,i} = 1$). If $V_0^r > V_1^r$ then we set $M'_{r,i} := 0$. Similarly for $V_1^r > V_0^r$. If $V_0^r = V_1^r$ then we do nothing.

A more complex case is when $M_{r,i} = 2$. We make sure the complementary haplotypes are given different alleles by setting the values of both haplotypes simultaneously. This does not cause a dependency on which haplotype is visited first because we have extra information we can take advantage of. We count votes for the complementary haplotype and treat them oppositely. That is, votes for the complementary haplotype having a 1 can be treated as votes for the current haplotype having a 0 (and vice versa). So letting r' be the row index for the complementary haplotype, we actually compare $V_0^r + V_1^{r'}$ and $V_1^r + V_0^{r'}$. This is helpful when SNPs near position i (which therefore will fall within shared tracts involving i) have already been phased (by trio pre-phasing or previous iterations). It also helps in making the best decision when both haplotypes receive a majority of votes for the same allele, e.g., both have a majority of votes for 0. In this case, taking into account votes for the two haplotypes simultaneously will result in whichever has *more* votes getting assigned the actual value 0. If they each receive the exact same number of votes, then no allele will be assigned. This also avoids the above-mentioned dependency on the order in which the haplotypes are visited – the outcome is the same since votes for both are taken into account.

In this manner, M' is calculated at each position. If $M' = M$ (i.e. no changes were made) then the algorithm terminates. Otherwise, $M := M'$ (M is replaced by M') and another iteration is run.

2.5 Phasing the individuals that are not a part of the largest component

Individuals that are part of small connected components will have a number of ambiguous sites once they have been phased using the edges in their connected component. For these individuals, we compute a minimum number of recombinations and mutations from their haplotypes to others that have better phasing (belong to larger components). We then assign these haplotypes phase based on minimizing the number of mutations plus recombinations in a similar manner as the approach of Minichiello Durbin [10].

Alternatively this could be done in a sampling framework, where we sample the haplotype with a probability that is a function of the number of mutations and recombinations.

2.6 Experimental results on simulated data

We tested the correctness of our algorithms using a simulated dataset consisting of 4000 individuals assuming no recombination. We used the Hudson Simulator [6] to generate 4000 haplotypes consisting of 3063 SNPs from chromosomes of length 10^5. We estimated a population size of 10^6 with a neutral mutation rate of 10^{-9}. We randomly sample from the distribution of simulated haplotypes with replacement such that each haplotype was sampled on average 1, 2, 3, 4, and 5 times. The phasing solution was almost perfect for the simulated dataset – decreasing from 0.5% error to 0.1% error with more samples.

3 Loss of heterozygosity regions

We call the loss of the normal allele a Loss of Heterozygosity (LOH) which may be a genetic determinant in the development of disease [9, 15]. The detection of CNVs, such as deletions, is an important aspect of GWAS to find LOH events, and yet, it is commonly overlooked due to technological and computational limitations. In some situations, individuals that are heterozygous at a particular locus can possess one normal allele and one deleterious allele.

LOH can be inferred using data from SNP arrays. The SNP calling algorithm for SNP arrays cannot distinguish between an individual who is homozygous for some allele a and an individual who has a deletion haplotype and the allele a (Figure 4, Left). LOH events can then be inferred by finding such genotypic events throughout the dataset. We will present an algorithm for computing putative LOH regions across GWAS datasets.

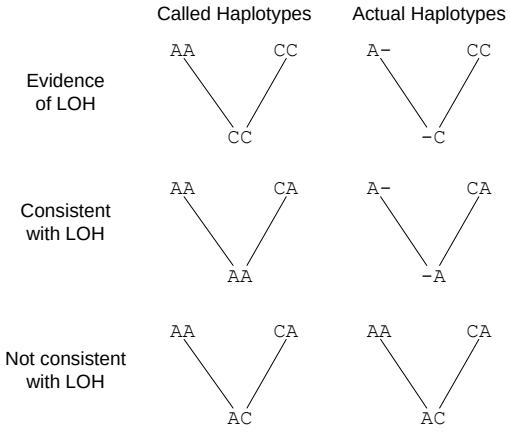

Father Genotype

Mother Genotype	0	1	2	3	Child Genotype
0	C^P	E^F	C^M	C^P	0
	N	E^M	E^M	E^M	1
	N	N	N	N	2
	C^P	C^P	C^M	C^P	3
1	E^M	N	E^M	E^M	0
	E^F	C^P	C^M	C^P	1
	N	N	N	N	2
	C^P	C^P	C^M	C^P	3
2	C^F	E^F	N	C^F	0
	E^F	C^F	N	C^F	1
	N	N	N	N	2
	C^F	C^F	N	C^F	3
3	C^P	E^F	C^M	C^P	0
	E^F	C^P	C^M	C^P	1
	N	N	N	N	2
	C^P	C^P	C^M	C^P	3

Figure 4: *Left: Three examples of inheritance patterns in GWAS data in the context of LOH. The Evidence of LOH (ELOH) shows strong correlation between LOH and a SNP site because the only possible explanation involves introducing a deletion haplotype. A consistent with LOH (CLOH) is consistent with a deletion haplotype but can also be explained with normal inheritance patterns. A not consistent with LOH (NCLOH) occurs when a deletion haplotype cannot be introduced to explain the trio inheritance pattern. Right: The correlation between inheritance pattern and ELOH, CLOH, and NCLOH. We define E to be ELOH, C to be CLOH, and N to be NCLOH. The superscript defines for which parent the putative deletion haplotype is associated. We define the superscript F to be consistent with a deletion haplotype inherited from the father, M for mother, and P for both parents.*

3.1 Definitions

A *trio* consists of three individual's genotypes and is defined by the inheritance pattern of parents to child. As before, let M denote the matrix of genotypes, let $M_{m,i}$ and $M_{f,i}$ denote the parent genotypes, and $M_{c,i}$ denote the child genotype for some trio of length l where $i = 1...l$. At any site i the trio may have 4^3 possible genotype combinations for which the trio can either be *consistent with LOH* (CLOH), *not consistent with LOH* (NCLOH), or *show evidence of LOH* (ELOH) (see Figure 4, Left). A trio at site i shows ELOH if the inheritance pattern can only be explained with the use of a deletion haplotype (or a genotyping error). A trio at site i is NCLOH if the inheritance pattern cannot be explained with the use of a deletion haplotype, and CLOH if it *may* be explained with the use of a deletion haplotype.

3.2 The LOH inference problem

We are given a set of n SNPs and a set of m trios genotyped at those SNPs. For each SNP/trio pair the SNP can have one of three labels:

- X - The marker is inconsistent with having a loss of heterozygosity (Figure 4, Left: Not Consistent with LOH).

- 0 - The marker is consistent with having a loss of heterozygosity (Figure 4, Left: Consistent with LOH).
- 1 - The SNP shows evidence of loss of heterozygosity, (Figure 4, Left: Evidence of LOH).

For any trio, a contiguous sequence of at least one 1 and an unbounded number of 0 sites is called an *putative deletion*. We call two putative deletions overlapping if any portion of one trio's putative deletion overlaps with another trio's deletion. For any given putative deletion, all 1's must either define a deletion in the same parent or a genotyping error. Also, all consistencies within the putative deletion must share the same parent (Figure 4, Right). Finding all deletions in the input matrix can be formulated as finding the minimum exact cover of all 1's for all disjoint overlapping putative deletion sets throughout each chromosome.

3.3 LOH inference algorithms

We present an exponential algorithm and a greedy heuristic for computing putative deletions. Both algorithms begin by parsing M and removing SNPs in which the Mendelian error rate is above 5% to remove artifacts from genotyping. We then calculate the LOH site vector for each trio in the dataset which corresponds to using that table defined in Figure 4 (Right) to translate each SNP site. This new matrix is denoted $N^{\left(\frac{|M|}{3} \times l\right)}$. In order to identify the putative LOHs, we define two operations on N: error correction call and deletion haplotype call. An error correction call will take an ELOH and call it a genotyping error effectively removing it from any particular deletion haplotype. A deletion haplotype call will identify an putative deletion as containing an inherited deletion haplotype.

The problem of inferring inherited deletion haplotypes can now be summarized by the objective function

$$min_N(\text{number of genotype error corrections}) + k_1 * (\text{number of deletion haplotypes})$$

where k_1 is some weighing factor. The k_1 factor in the objective function can be a simple constant factor or a more complex distribution. For complex distributions, we weigh the score to prefer deletions that are shared by many people. In these cases we could define k_1 to be k_2(number of conserved individuals) + k_3(length of overlapping region + k_4(ELOH/CLOH)). k_1 can be tuned depending on the situation. For example, association tests will tune the parameter to favor many conserved individuals, that is, large overlapping putative deletions. This problem is NP-hard for general N (proof omitted). In the case of the Multiple Sclerosis dataset, the matrix N contains small overlapping putative deletions and over 95% of N is non-putative deletions, that is, N is very sparse.

Algorithm 1. We start by giving an exact exponential algorithm for this objective function. Let x_i denote a set of overlapping putative deletions. For sparse N we can reduce the minimization function from min_N to $min_{x_1..x_s}$ where $x_1..x_s \in N$ and $\{x_1..x_s\} \subseteq N$. Since any particular putative deletion is defined by the ELOH sites, we can enumerate all feasible inclusion/exclusion combinations of any x_i. Computing this for all putative deletions demands work proportional to $\sum_{i=1}^{s} B(e_i)$ where e_i is the number of ELOH sites in x_i and B is the Bell number. In practice, we found that e_i is always bounded by a small constant.

Algorithm 2. We've developed a greedy algorithm to run for cases where e_i is an unbounded constant or shows dependence on n (Figure 5). For any $x_i \in N$ the algorithm selects the component with the maximum *trio sharing*, that is, the possibly overlapping putative deletions that include the most ELOH sites. For each trio involved in this component, the algorithm will attempt to

extend the putative deletion call to ELOH sites on either side. If extending the deletion haplotype improves the score the algorithm includes the sites. When no ELOH inclusion improves the score, the algorithm calls the site(s) a deletion haplotype or genotyping error depending on the objective function and then recomputes adjacent LOH sites.

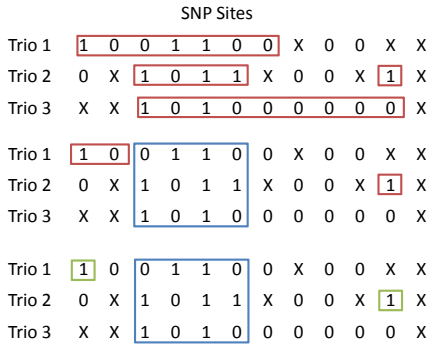

Figure 5: *This figure shows the steps of the greedy algorithm for finding putative deletions (consistencies with particular parents are omitted for simplicity). First select the SNP sites with the largest trio sharing: SNP sites 3-6. Can we extend any trios putative deletion in any direction to improve our score? No, so make SNP sites 3-6 a deletion, update intervals, and continue. Select SNP sites with largest trio sharing: SNP sites 1 or 11. Does making these SNP sites deletions minimize our objective function? No, so call them genotyping errors.*

We tested the algorithm using the trio dataset from a recent Multiple Sclerosis (MS) GWAS study [2]. The MS dataset consists of 2793 family trios - parents and affected child - and 334923 SNPs. Let a conserved LOH component of size k be a set of overlapping putative deletions in k individuals. Tests show thousands of putative LOH regions in the genotypes of the trios where the maximum size of a conserved LOH component is 16 and the average size 8.86.

4 Conclusion and future work

Identifying putative deletions is an especially important preprocessing step in order to clean the data. Instead of assuming all Mendelian errors are truly due to incorrect genotyping, the LOH algorithm calls many of these errors part of a deletion haplotype. Mendelian errors attributed to deletion haplotypes can also be factored into the phasing algorithm when applied to GWAS data.

We have shown that long range phasing using Clark consistency graphs is practical for very large datasets and the accuracy of the algorithm improves rapidly with the size of the dataset. We have also given an algorithm that removes most Mendelian inconsistencies and distinguishes between genotyping errors and deletion events.

All algorithms are available via sending a request to the corresponding authors.

5 Acknowledgments

Thanks to the International Multiple Sclerosis Genetics Consortium for sharing the Multiple Sclerosis GWAS dataset.

References

[1] AG Clark, *Inference of haplotypes from PCR-amplified samples of diploid populations*, Mol Biol Evol **7** (1990), no. 2, 111–122.

[2] The International Multiple Sclerosis Genetics Consortium, *Risk alleles for multiple sclerosis identified by a genomewide study*, N Engl J Med **357** (2007), no. 9, 851–862.

[3] Daniel F. Gudbjartsson, G. Bragi Walters, Gudmar Thorleifsson, Hreinn Stefansson, Bjarni V. Halldorsson, Pasha Zusmanovich, Patrick Sulem, Steinunn Thorlacius, Arnaldur Gylfason, Stacy Steinberg, Anna Helgadottir, Andres Ingason, Valgerdur Steinthorsdottir, Elinborg J. Olafsdottir, Gudridur H. Olafsdottir, Thorvaldur Jonsson, Knut Borch-Johnsen, Torben Hansen, Gitte Andersen, Torben Jorgensen, Oluf Pedersen, Katja K. Aben, J. Alfred Witjes, Dorine W. Swinkels, Martin den Heijer, Barbara Franke, Andre L. M. Verbeek, Diane M. Becker, Lisa R. Yanek, Lewis C. Becker, Laufey Tryggvadottir, Thorunn Rafnar, Jeffrey Gulcher, Lambertus A. Kiemeney, Augustine Kong, Unnur Thorsteinsdottir, and Kari Stefansson, *Many sequence variants affecting diversity of adult human height*, Nat Genet **40** (2008), no. 5, 609–615.

[4] Bjarni V. Halldórsson, Vineet Bafna, Nathan Edwards, Shibu Yooseph, and Sorin Istrail, *A survey of computational methods for determining haplotypes*, (2004).

[5] Bryan N. Howie, Peter Donnelly, and Jonathan Marchini, *A flexible and accurate genotype imputation method for the next generation of genome-wide association studies*, PLoS Genet **5** (2009), no. 6, e1000529.

[6] Richard R. Hudson, *Generating samples under a wright-fisher neutral model of genetic variation*, Bioinformatics **18** (2002), no. 2, 337–338.

[7] Sorin Istrail, *The haplotype phasing problem*, Symposium in Honor of Mike Waterman's 60th Birthday, 2002.

[8] Augustine Kong, Gisli Masson, Michael L. Frigge, Arnaldur Gylfason, Pasha Zusmanovich, Gudmar Thorleifsson, Pall I. Olason, Andres Ingason, Stacy Steinberg, Thorunn Rafnar, Patrick Sulem, Magali Mouy, Frosti Jonsson, Unnur Thorsteinsdottir, Daniel F. Gudbjartsson, Hreinn Stefansson, and Kari Stefansson, *Detection of sharing by descent, long-range phasing and haplotype imputation*, Nat Genet **40** (2008), no. 9, 1068–1075.

[9] Steven A. McCarroll, Finny G. Kuruvilla, Joshua M. Korn, Simon Cawley, James Nemesh, Alec Wysoker, Michael H. Shapero, Paul I. W. de Bakker, Julian B. Maller, Andrew Kirby, Amanda L. Elliott, Melissa Parkin, Earl Hubbell, Teresa Webster, Rui Mei, James Veitch, Patrick J. Collins, Robert Handsaker, Steve Lincoln, Marcia Nizzari, John Blume, Keith W. Jones, Rich Rava, Mark J. Daly, Stacey B. Gabriel, and David Altshuler, *Integrated detection and population-genetic analysis of snps and copy number variation*, Nat Genet **40** (2008), no. 10, 1166–1174.

[10] Mark J. Minichiello and Richard Durbin, *Mapping trait loci by use of inferred ancestral recombination graphs*, **79** (2006), no. 5, 910–922.

[11] F. Rivadeneira, U. Styrkarsdottir, K. Estrada, B. Halldorsson, Y. Hsu, J. B. Richards, M. C. Zillikens, F. Kavvoura, N. Amin, Y. Aulchenko, L. Cupples, P. Deloukas, S. Demissie, E. Grundberg, A. Hofman, A. Kong, D. Karasik, J. van Meurs, B. Oostra, T. Pastinen, H. A. P. Pols, G. Sigurdsson, N. Soranzo, G. Thorleifsson, U. Thorsteindottir, F. Williams, S. Wilson, Y. Zhou, S. Ralston, C. van Duijn, T. Spector, D. Kiel, K. Stefansson, J. Ioannidis, A. G. Uitterlinden, and In behalf of the GEFOS Consortium, *Bone*, vol. 44, ch. Twenty loci associated with bone mineral density identified by large-scale meta-analysis of genome-wide association datasets, pp. S230–S231, Elsevier Science, Jun 2009.

[12] Paul Scheet and Matthew Stephens, *A fast and flexible statistical model for large-scale population genotype data: Applications to inferring missing genotypes and haplotypic phase*, **78** (2006), no. 4, 629–644.

[13] Roded Sharan, Bjarni V. Halldórsson, and Sorin Istrail, *Islands of tractability for parsimony haplotyping*, IEEE/ACM Transactions on Computational Biology and Bioinformatics **3** (2006), no. 3, 303–311.

[14] Nayanah Siva, *1000 genomes project.*, Nature biotechnology **26** (2008), no. 3, 256.

[15] Hreinn Stefansson, Dan Rujescu, Sven Cichon, Olli P. H. Pietilainen, Andres Ingason, Stacy Steinberg, Ragnheidur Fossdal, Engilbert Sigurdsson, Thordur Sigmundsson, Jacobine E. Buizer-Voskamp, Thomas Hansen, Klaus D. Jakobsen, Pierandrea Muglia, Clyde Francks, Paul M. Matthews, Arnaldur Gylfason, Bjarni V. Halldorsson, Daniel Gudbjartsson, Thorgeir E. Thorgeirsson, Asgeir Sigurdsson, Adalbjorg Jonasdottir, Aslaug Jonasdottir, Asgeir Bjornsson, Sigurborg Mattiasdottir, Thorarinn Blondal, Magnus Haraldsson, Brynja B. Magnusdottir, Ina Giegling, Hans-Jurgen Moller, Annette Hartmann, Kevin V. Shianna, Dongliang Ge, Anna C. Need, Caroline Crombie, Gillian Fraser, Nicholas Walker, Jouko Lonnqvist, Jaana Suvisaari, Annamarie Tuulio-Henriksson, Tiina Paunio, Timi Toulopoulou, Elvira Bramon, Marta Di Forti, Robin Murray, Mirella Ruggeri, Evangelos Vassos, Sarah Tosato, Muriel Walshe, Tao Li, Catalina Vasilescu, Thomas W. Muhleisen, August G. Wang, Henrik Ullum, Srdjan Djurovic, Ingrid Melle, Jes Olesen, Lambertus A. Kiemeney, Barbara Franke, Chiara Sabatti, Nelson B. Freimer, Jeffrey R. Gulcher, Unnur Thorsteinsdottir, Augustine Kong, Ole A. Andreassen, Roel A. Ophoff, Alexander Georgi, Marcella Rietschel, Thomas Werge, Hannes Petursson, David B. Goldstein, Markus M. Nothen, Leena Peltonen, David A. Collier, David St Clair, and Kari Stefansson, *Large recurrent microdeletions associated with schizophrenia*, Nature **455** (2008), no. 7210, 232–236.

[16] Matthew Stephens, Nicholas J. Smith, and Peter Donnelly, *A new statistical method for haplotype reconstruction from population data*, **68** (2001), no. 4, 978–989.

[17] Unnur Styrkarsdottir, Bjarni V. Halldorsson, Solveig Gretarsdottir, Daniel F. Gudbjartsson, G. Bragi Walters, Thorvaldur Ingvarsson, Thorbjorg Jonsdottir, Jona Saemundsdottir, Jacqueline R. Center, Tuan V. Nguyen, Yu Bagger, Jeffrey R. Gulcher, John A. Eisman, Claus Christiansen, Gunnar Sigurdsson, Augustine Kong, Unnur Thorsteinsdottir, and Kari Stefansson, *Multiple Genetic Loci for Bone Mineral Density and Fractures*, N Engl J Med **358** (2008), no. 22, 2355–2365.