phASER: Long range phasing and haplotypic expression from RNA sequencing

Stephane E. Castel[1,2]*, Pejman Mohammadi[1,2], Wendy K. Chung[3], Yufeng Shen[2,4], Tuuli Lappalainen[1,2]*

1 New York Genome Center, New York, NY
2 Department of Systems Biology, Columbia University, New York, NY
3 Departments of Pediatrics and Medicine, Columbia University, New York, NY
4 Department of Biomedical Informatics, Columbia University, New York, NY

* To whom correspondence should be addressed:
scastel@nygenome.org, tlappalainen@nygenome.org

Haplotype phasing of genetic variants is important for clinical interpretation of the genome, population genetic analysis, and functional genomic analysis of allelic activity. Here we present phASER, a fast and accurate approach for phasing variants that are overlapped by sequencing reads, including those from RNA-sequencing (RNA-seq), which often span multiple exons due to splicing. This provides 1) dramatically more accurate phasing of rare and *de novo* variants compared to population-based phasing; 2) phasing of variants in the same gene up to hundreds of kilobases away which cannot be obtained from DNA-sequencing reads; 3) high confidence measures of haplotypic expression, greatly improving power for allelic expression studies.

The phasing of rare and *de novo* variants is crucial for identifying putative causal variants in clinical genetics applications, for example by distinguishing compound heterozygotes from two variants on the same allele. Existing methods to phase variants are limited to phasing by transmission [1], only available in familial studies, population based phasing [2], which is ineffective for rare and *de novo* variants, phasing by sequencing long genomic fragments [3], which requires specialized and costly technology, phasing using expression data by inferring haplotype through allelic imbalance [4], which only applies to loci with well-detected allelic expression [5], and physical techniques, for example those that employ allelic probes and microscopy, which are low throughput but high confidence [6]. More recently "read backed phasing" using readily available short read DNA-sequencing (DNA-seq) has emerged [7], however it is limited by the relatively short distances which can be spanned by the reads. Our approach, called phASER (phasing and allele specific expression from RNA-seq), extends the idea of read backed phasing to RNA-seq reads, which due to splicing, enables phasing of variants over long genomic distances. Data from both DNA-seq and RNA-seq libraries can be integrated by phASER to produce high confidence phasing of proximal variants, primarily within the same gene, and when available population phasing can also be leveraged to produce full chromosome-length haplotypes (Figure 1a). Source code and complete documentation for phASER and its associated tools are available through github (https://github.com/secastel/phaser).

To phase variants phASER uses reads that overlap more than one variant to generate haplotype blocks, each composed of a network of read connections between variants that are observed in the data (Figure S1a). For each pair of connected variants a binomial test is performed to determine if the number of reads supporting configurations other than the chosen phase can be explained by the level of noise observed in the sequencing experiment, allowing for filtering of low confidence phasing (Figure S1b). Within a block each possible haplotype configuration is constructed (2^n variants) and the number of connections that support each configuration is counted. The configuration with the most support is selected as the likely phase. As a gold standard we compared phASER used with high coverage RNA-seq data generated from a lymphoblastoid cell line (LCL) [8] to Illumina's NA12878 Platinum Genome, sequenced at 200x and phased by transmission using parental data (http://www.illumina.com/platinumgenomes/), and found that phASER identified the correct phasing for 98.6% of variant pairs tested, and unlike population based phasing performed well for low frequency variants (Figure 1b). At large distances the performance decreased as a result of incorrect read mapping, however this could be easily addressed by filtering reads based on alignment score (Figure S2a).

When population phased data are available, haplotype blocks are phased relative to each other, producing a single genome wide phase through a method we call phase anchoring. Phase anchoring uses the population phase of each variant in a block weighted by their allele frequencies to assign a genome wide phase, since common variants are more likely to have correct population phasing (Figure S1c). Using this approach with combined RNA, whole exome sequencing (WES) and whole genome sequencing (WGS) data phASER could correct the genome wide phase of 1271 or 7.5% of rare variants (MAF $\leq$ 0.001) that were incorrectly phased with population data, while only assigning an incorrect phase for 54 or 0.32% of variants (Figure 1c).

In order to evaluate the improved long-range phasing of RNA-seq reads, we compared phASER results between WES, WGS, and two read lengths of paired-end (PE) RNA-seq (75bp and 250bp) from 4 Genotype Tissue Expression (GTEx) individuals where matched libraries were available [9]. As expected, long read RNA-seq

yielded the greatest proportion of distantly phased variants, with an average of 4300 equal to 5.8% of variants phased greater than 5kb, while at this distance WES and WGS phased 0 and 7 variants respectively (Figure 1d). Using RNA-seq reads we found that phasing remained accurate over a range of read lengths (Figure S2b), but that longer read lengths greatly increased both the distance and number of variants that could be phased (Figure S2c).

Having established that phASER is accurate and robust across sequencing assay types, we next benchmarked it against the commonly used GATK ReadBackedPhasing tool [10], which works only for DNA-seq reads, using both simulated and experimental WGS data as well as WES data. We found phASER to be more accurate in all cases (99.1-99.9% of variants phased correctly vs 71.1-78.5% for GATK, Figure S3a), more sensitive (phasing 1.23-1.36x more variants, Figure S3a), and significantly faster (1.6x-10.3x, Figure S3b) than the GATK tool, while producing haplotype blocks of the same size when DNA-seq reads were used (Figure S3c). In addition, phASER is highly scalable, since runtime is linearly proportional to both the number of reads and variants (Figure S3d-e), and unlike the GATK tool, it can be run using multiple threads, providing a dramatic increase in speed (Figure S3f).

We next sought to evaluate the utility of phASER for genetic studies. First we used GTEx data to demonstrate the number of variants that could be phased as a function of the number of tissues for which RNA-sequencing data is available. We began with whole blood, and progressively added libraries from up to 14 other distinct tissues. With joint phasing using 14 tissues, almost 50% of all heterozygous coding variants could be phased with at least one other variant (Figure S4a). When used individually, the total proportion of coding variants that could be phased for a given tissue was 15-27% (Figure S4b), and was dependent on transcriptome diversity [11] (Figure S4c), but not total read depth (Figure S4d).

With an understanding of the number of variants that can be phased, we next applied phASER to identify cases of compound heterozygosity using exome and RNA-seq data from 1000 Genomes individuals [12]. First we assessed the accuracy of compound heterozygote calls using population phasing compared to exome + RNA read backed phasing. As expected, protein-truncating and splice variants that are usually rare were enriched in cases where population phasing was incorrect, with cases involving stop gain variants being 2.9x more likely to be phased incorrectly than others (Figure S5a-b). Next, to demonstrate the advantage of using RNA-seq data over exome alone for phasing, we identified instances of compound heterozygosity involving at least one rare (MAF < 1%) variant with predicted loss of function (LoF) or damaging protein effects [13]. Including RNA-seq data from only one tissue increased the number of compound heterozygotes that could be identified in the most severe class (LoF x Damaging) by 1.3x over WES alone, and ranged between 1.19 – 1.15x for other combinations (Figure 2a).

Finally, we used paired exome and fibroblast RNA-seq from 20 patients with congenital diaphragmatic hernia [14] to illustrate phASER's application to a typical clinical workflow to prioritize putatively causal variants. Assuming that causative variants would be rare, recessive, damaging, and expressed in the tissue of disease relevance, phase information generated with phASER from exome and RNA-seq reads could prioritize a median of 25 alleles involved in cases of compound heterozygosity per individual (*trans*), while assigning a lower priority to a median of 44 alleles, where the alleles where on the same haplotype (*cis*) (Figure 2b). The inclusion of RNA-seq boosted the number of cases that could be identified by 2.6x for *trans* and 1.4x for *cis* interactions.

Outside of clinical genetics, variant phasing is important for allelic expression (AE) analysis, which aims to quantify the relative expression of one allele versus another [5], and has emerged as a powerful method to study diverse biological processes including cis-regulatory variation [15], parent of origin imprinting [16], and protein-truncating variants [17]. AE is typically measured at single heterozygous variants, however the unit of interest is often a gene or transcript, which may contain many variants. Integrating read counts across phased variants can greatly improve the power to detect AE, but simply adding up allele counts can lead to double counting of reads

(if variants are covered by the same read), and both false positives and negatives as a result of incorrect phasing. To address this limitation, when used with RNA-seq data, phASER quantifies the expression of phased haplotypes by reporting the number of unique reads that map to each. To demonstrate this utility we generated haplotypic counts at genes with known expression quantitative trait loci (eQTL) for Geuvadis samples using either single variants with population based phasing alone, or phased haplotype blocks generated by phASER (Figure S5). By improving phase and preventing double counting phASER reduced false positives at 56.2% of genes tested, while uncovering false negatives at 7.3% (Figure 2c).

In summary, phASER provides scalable and high confidence variant phasing, incorporating RNA-seq and DNA-seq data with population phasing, allowing phasing over longer distances than previous read-based methods. We have demonstrated that this method has direct applications in medical genetics, where improved resolution of compound heterozygotes and identification of parental origin of de novo variants can lead to changes in their interpretation. Furthermore, phASER improves the accuracy of haplotypic expression when integrating allelic counts across variants by reducing false positives. Our approach will complement the existing repertoire of phasing methods [18] and makes use of a readily available experimental data type that has become trivial to produce, allowing for phasing of rare and distant variants at high accuracy. As RNA-seq experiments become commonplace in clinical and population scale studies, phASER will become a valuable tool for rare variant phasing.

Figures

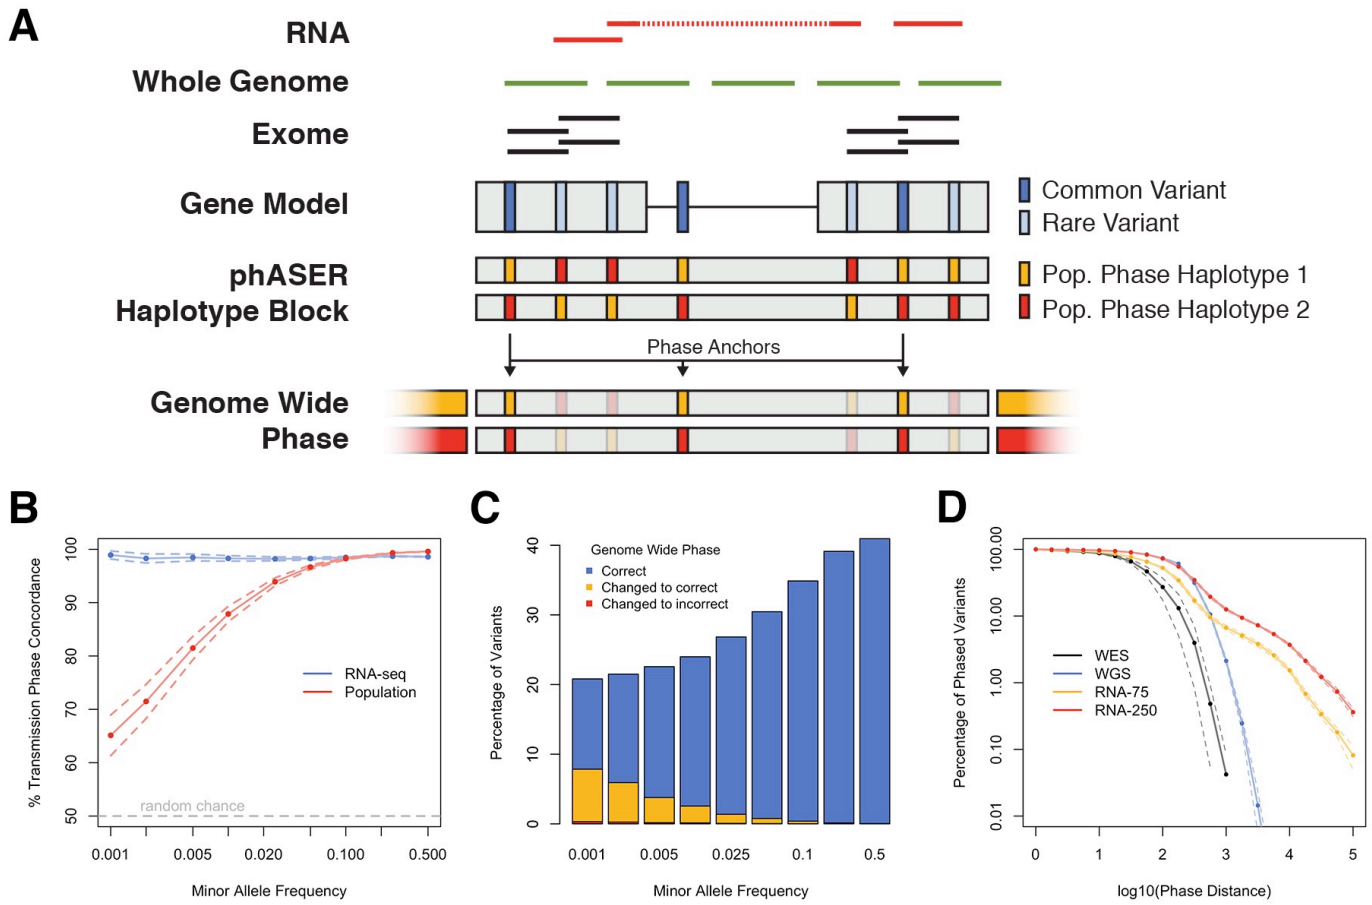

Figure 1. Read backed haplotype phasing that incorporates RNA-seq using phASER. A) phASER produces accurate variant phasing through the use of combined DNA and RNA read backed phasing integrated with population phasing. Due to splicing, RNA-seq reads often span exons and UTRs, allowing read backed phasing over long ranges, while high coverage exome and whole genome sequencing can phase close proximity variants. A local haplotype is produced by testing all possible phase configurations, and selecting the configuration with the most support. Local haplotype blocks can be phased relative to one another when population data is available by anchoring the phase to common variants, where the population phase is likely correct. B) Concordance between either population or RNA-seq based phasing and phasing by transmission using the Illumina NA12878 Platinum Genome as a function of variant minor allele frequency. Concordance is defined per variant as the percentage of variant – variant phase events that are correct as compared to the known transmission phase. C) Percentage of all heterozygous variants as a function of minor allele frequency that can be assigned a genome wide phase by phASER using phase anchoring and combined RNA, WES, and WGS for NA12878. Variants are broken into those where population and read backed phasing assign the correct phase (correct), those where read backed phasing corrected the population phase (changed to correct) or those where read backed phasing made the phasing incorrect (changed to incorrect), all as compared to transmission phasing. D) Percentage of phased variants that can be phased at greater than or equal to increasing genomic distances using WES, WGS, paired-end 75 and 250 RNA-seq data in four GTEx samples. Solid lines represent the means, and dotted lines the standard error.

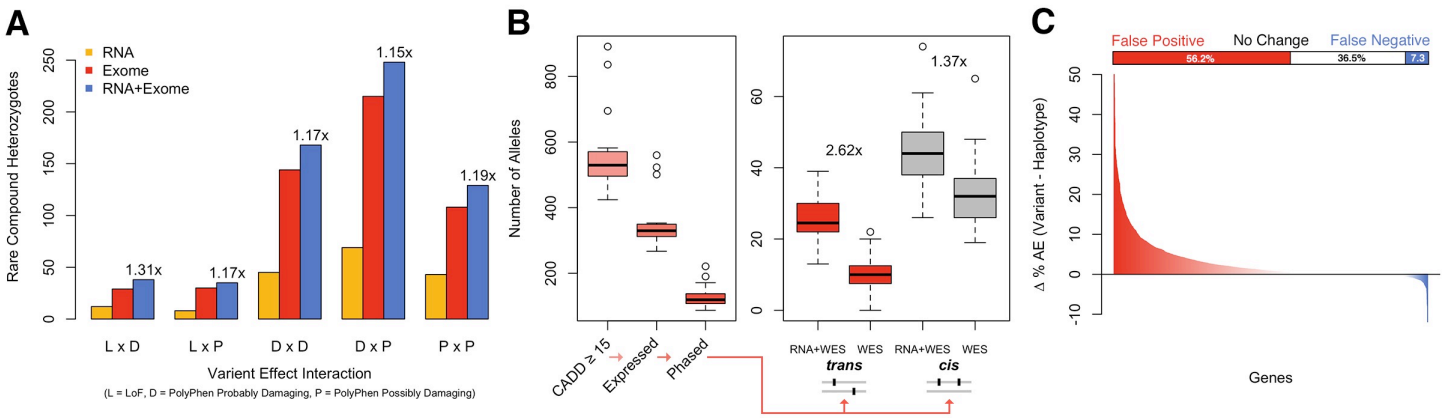

Figure 2. Application of RNA-seq based haplotype phasing to studies of functional variants and allelic expression analysis. A) Instances of compound heterozygosity involving rare (MAF < 0.01) loss of function (L), probably damaging (D) or possibly damaging (P) coding variants called using phase data generated by phASER with either RNA-seq reads, exome-seq reads, or both for 1000 Genomes European individuals with Geuvadis LCL RNA-seq data. The fold increases in the number of compound heterozygotes when RNA-seq data is included are indicated. B) Example application of phASER to prioritize rare (alternative AF < 0.01 in 1000 Genomes) recessive alleles in a clinical study that includes both exome-seq and RNA-seq in a tissue of clinical relevance [14]. Boxplots show the number of heterozygous alleles per individual after these successive filtering steps were applied: CADD phred score ≥ 15, expressed in fibroblast RNA-seq data, phased with read backed phasing, involved in either *trans* or *cis* interactions with another deleterious variant (CADD ≥ 15) using RNA and exome data (RNA+WES) or exome alone (WES). The fold increases from including RNA-seq data are indicated. C) The difference in percentage of individuals with significant (binomial test, q < 0.05) allelic expression for each gene with a known heterozygous *cis* eQTL calculated by either summing all single variant read counts across haplotypes using population phasing, or by summing phASER haplotype blocks phased relative to each other with phase anchoring (see Figure S6). Genes where an increase in the percentage of individuals with significant allelic expression is observed are colored red, representing false negatives when summing single variant counts, while those with a decrease, representing false positives, are colored blue. The bar plot above indicates the percentage of the 1118 genes where allelic expression was measured that fall into each category.

Supplemental Figures

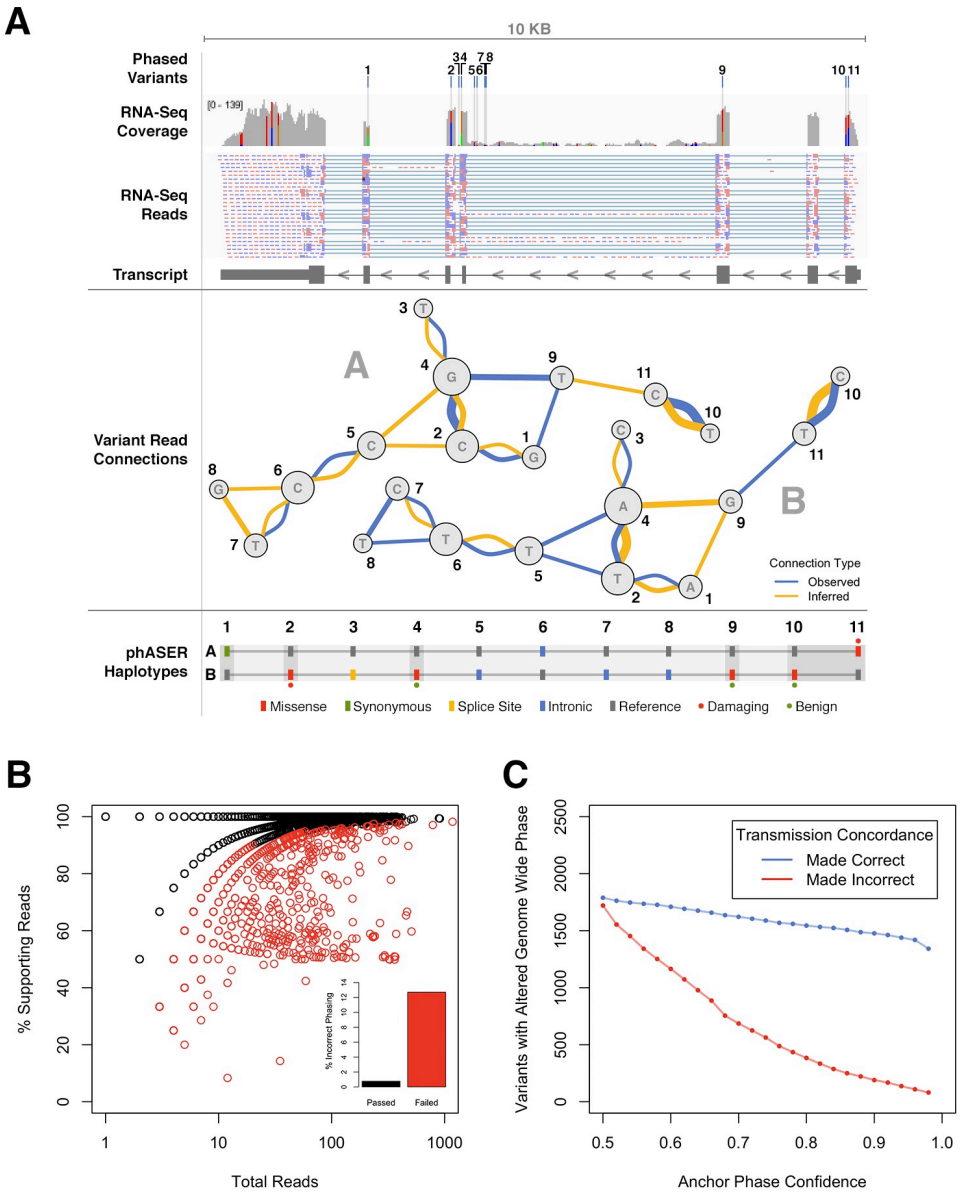

Figure S1. Haplotype phasing using RNA-seq reads and measures of phase confidence with phASER. A) RNA-seq reads often span exons containing heterozygous variants, as illustrated for the gene *RPUSD4* with coverage from a Geuvadis [15] LCL overlaid, and variants labeled 1 through 10. Reads that overlap more than one variant are used to produce haplotype blocks composed of a network of connected variants. The haplotype configuration with the most support is selected as the phase, as illustrated here using both observed connections (blue, read spans variants) and inferred connections (yellow, read spans variants on other haplotype). B) Percentage of reads supporting the chosen phase as a function of total reads at each variant – variant connection for phASER run using NA12878 RNA [8] + WES data [12]. Each point represents a variant – variant phasing, points in black passed the phase confidence test, while those in red failed (evidence for a conflicting phase configuration, $p < 0.01$). The inset bar plot shows the percentage of variant pairs phased incorrectly (versus transmission phasing) for variant connections that passed the test (black) and those that failed (red). C) The anchor confidence statistic robustly removes incorrect genome wide phasing as shown by the number of variants with genome wide phase either made correct or made incorrect as a function of anchor phase confidence with phasing by transmission as a ground truth for NA12878 run with 1000 Genomes exome [12], Illumina Platinum Genomes WGS (5x) and LCL RNA-seq data [8].

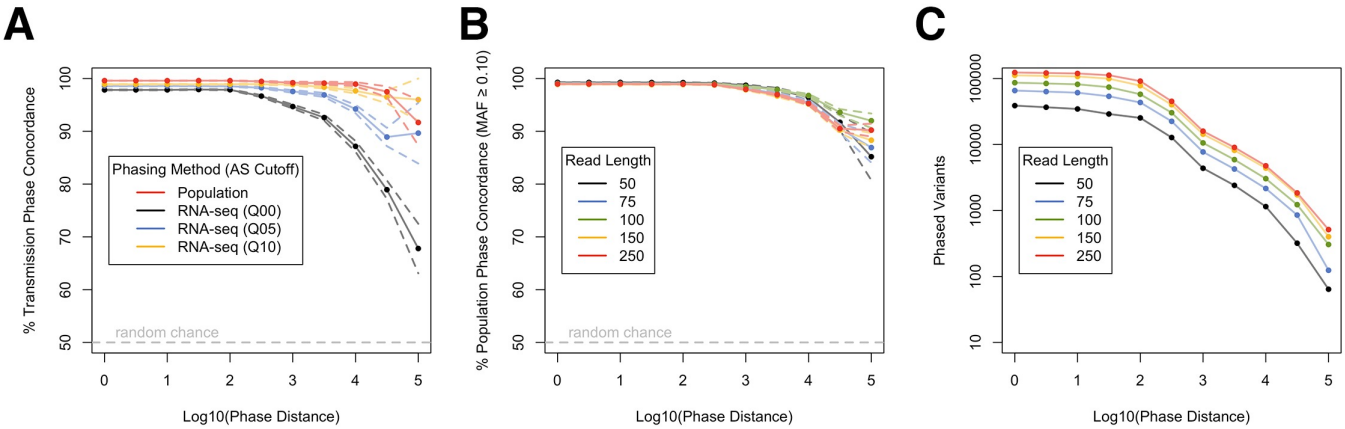

Figure S2. Effects of alignment quality and read length on RNA-seq read backed phasing. A) Comparison of RNA-seq phasing to phasing by transmission using either no alignment score cutoff (Q00), or cutoffs equal to the bottom 5% (Q05), and 10% (Q10), with population based phasing shown for comparison. Generated using NA12878 LCL RNA-seq data [8], 1000 Genomes Phase 3 population phasing [12], and Illumina Platinum Genome transmission phasing. Comparison of RNA-seq phasing to population phasing for variants with MAF > 10% at increasing read lengths (b), and number of variants phased at equal to or greater than increasing genomic distances for increasing read lengths (c), using a GTEx [9] long read RNA-seq library (GTEX-WFON-0001-SM-5S2SE) clipped to various lengths. Solid lines represent the means, and dotted lines the standard error.

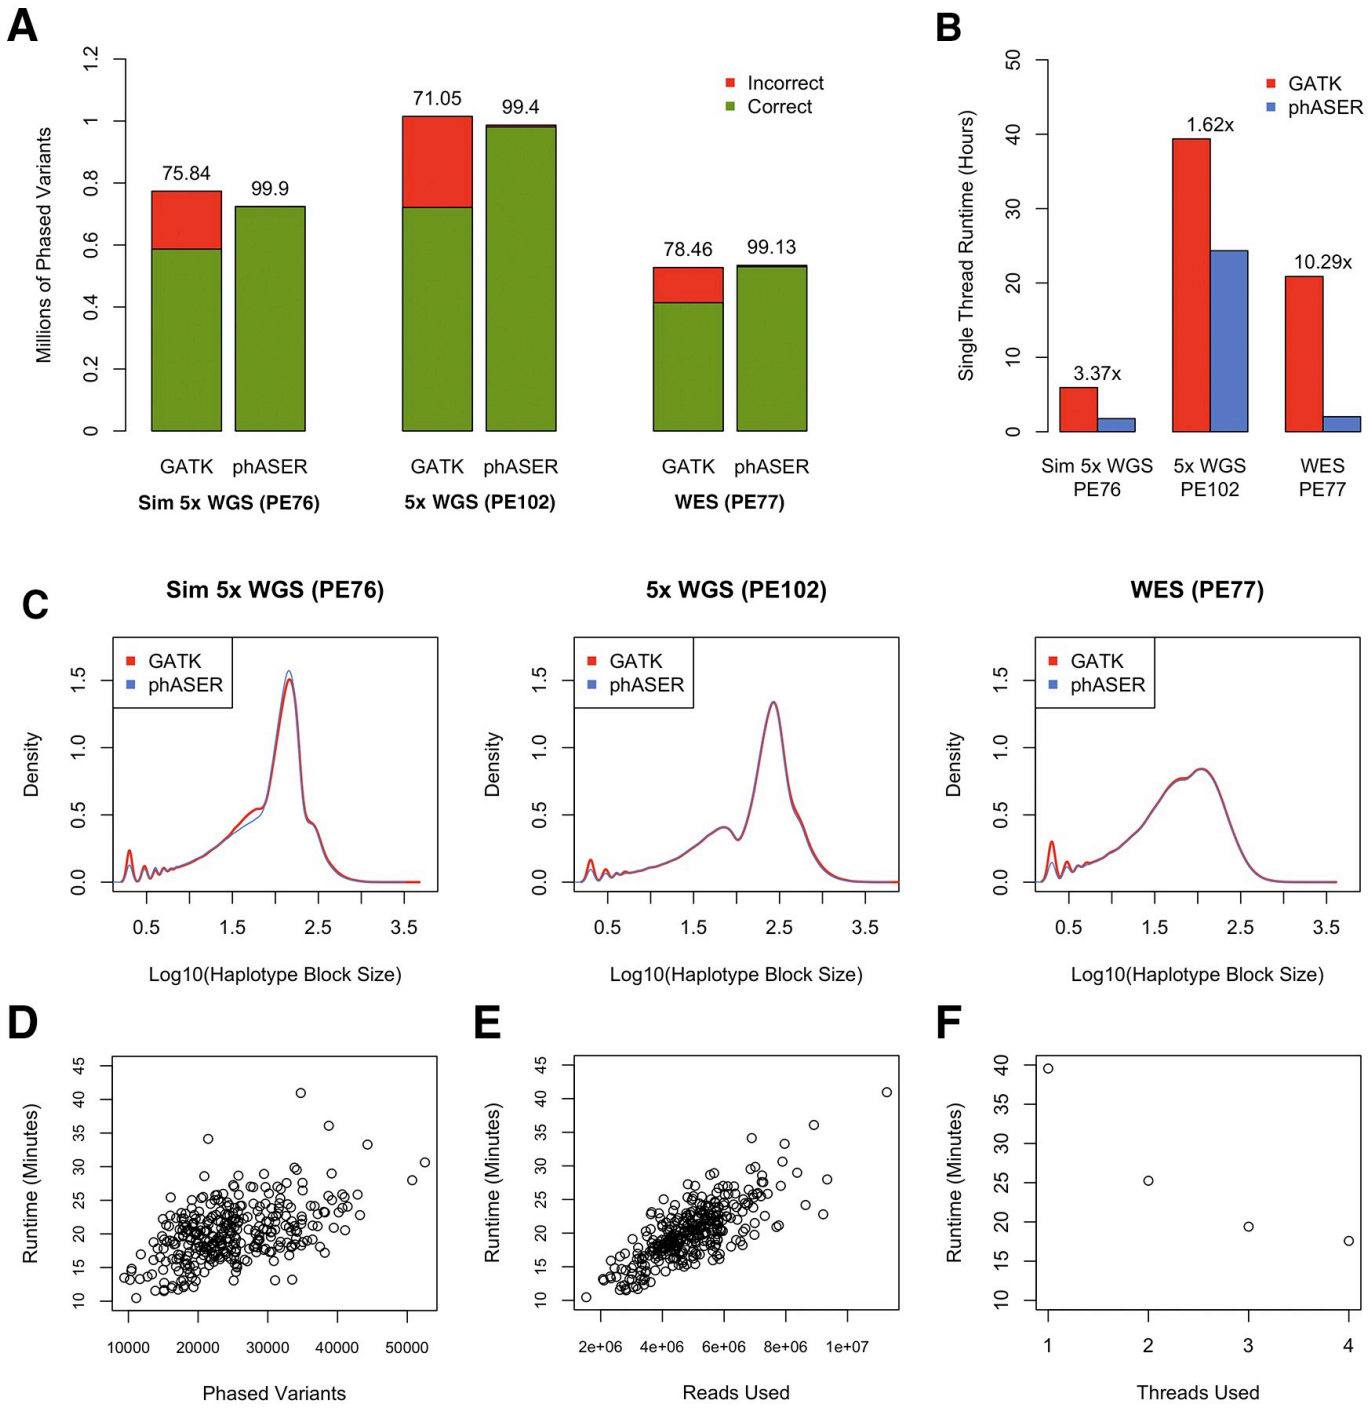

Figure S3. Variant phasing is efficient and accurate using phASER. Comparison of phASER to the GATK Read Backed Phasing tool, using either simulated 5x WGS (NA06986, PE76), experimental 5x WGS (NA12878 Illumina Platinum Genome, PE102), or WES (NA12878 1000 Genomes phase 3 [12], PE77). phASER correctly phases more variants (A), and runs much faster (B) than the GATK tool while phasing variants over the same distance (C). The runtime of phASER scales efficiently both with the number of variants phased (D) and the number of read used (E). In addition, phASER can be run using multiple threads for a dramatic increase in speed (F). A variant was considered correctly phased if 100% of observed variant – variant allele configurations were concordant with the known true phase (either the haplotypes used to simulate reads, or the transmission phased NA12878 haplotypes).

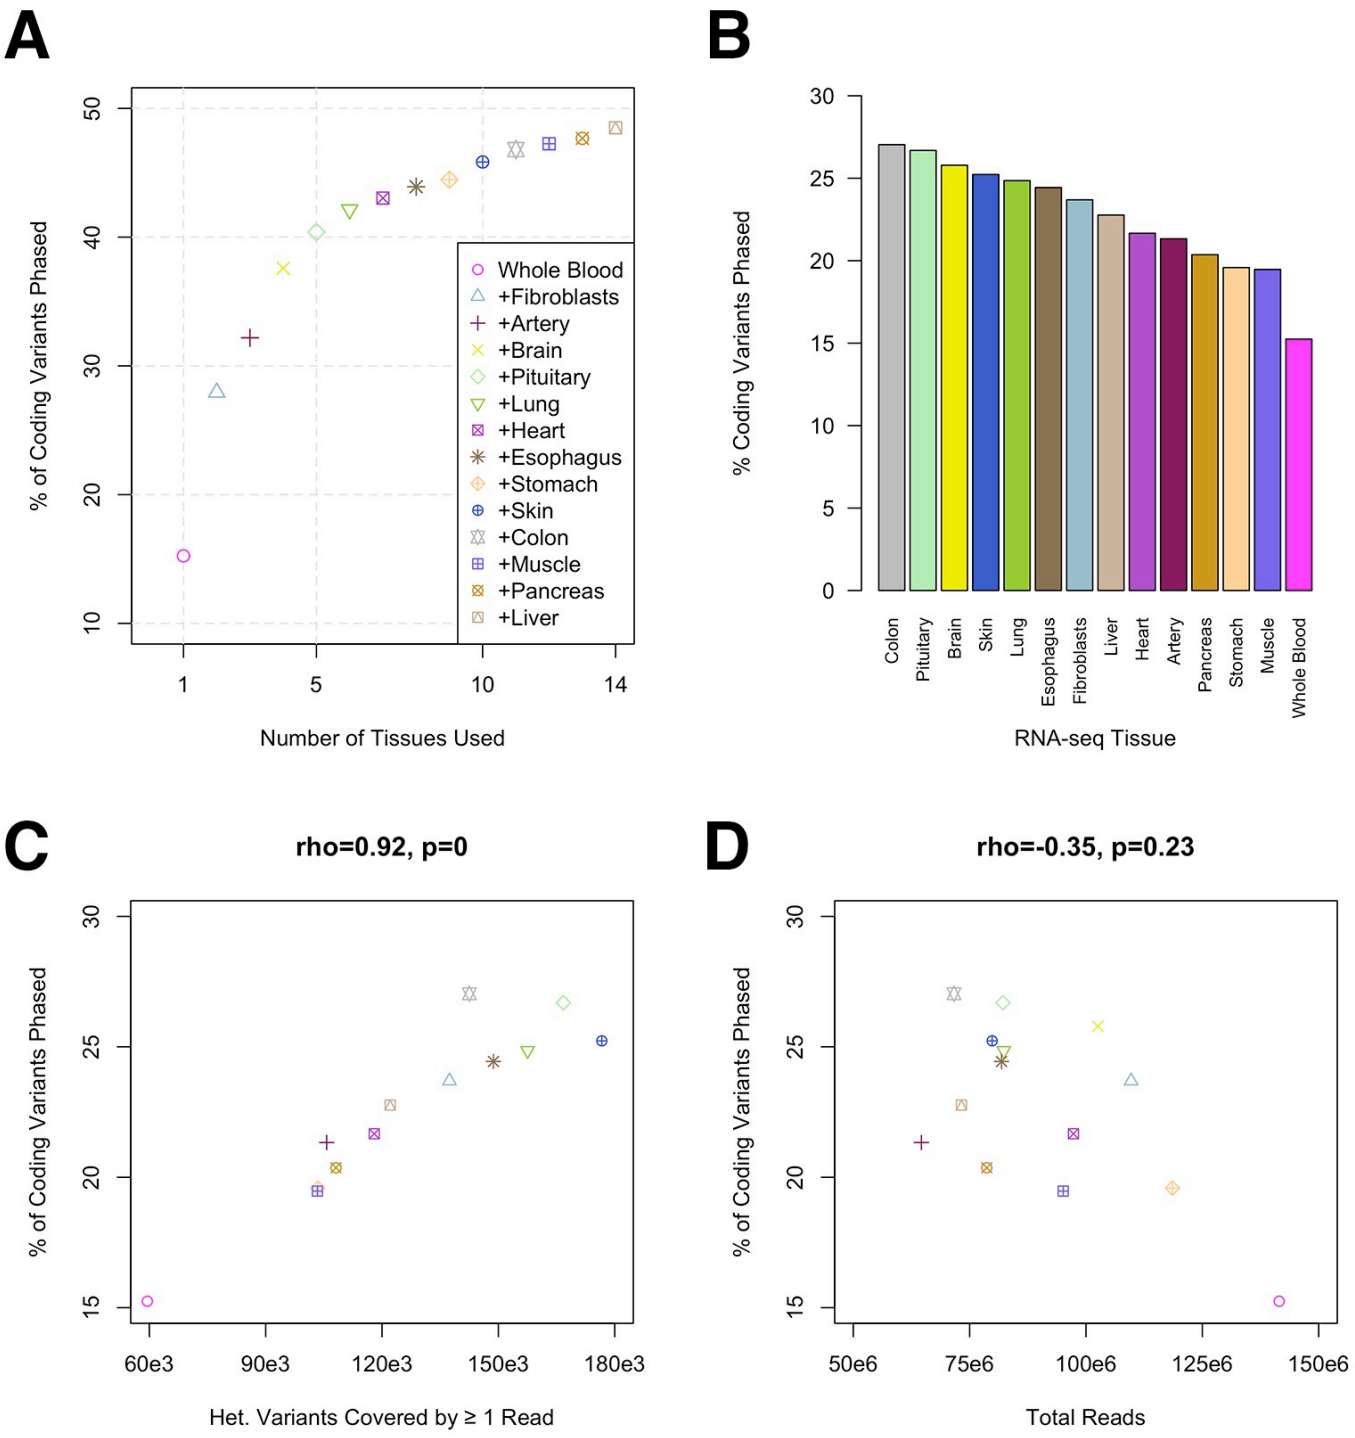

Figure S4. Joint RNA-seq based phasing over multiple tissues greatly improves the number of variants that can be phased in an individual. A) Percentage of coding variants that can be phased beginning with RNA-seq from whole blood, and progressively adding data from up to 14 other distinct tissues from GTEx individual ZAB4 [9]. B) Percentage of coding variants that can be phased using each tissue from (A) individually. C-D) Percentage of coding variants as a function of the number of heterozygous variants covered by at least one read (C) or total library reads (D) for each tissue from (A) individually.

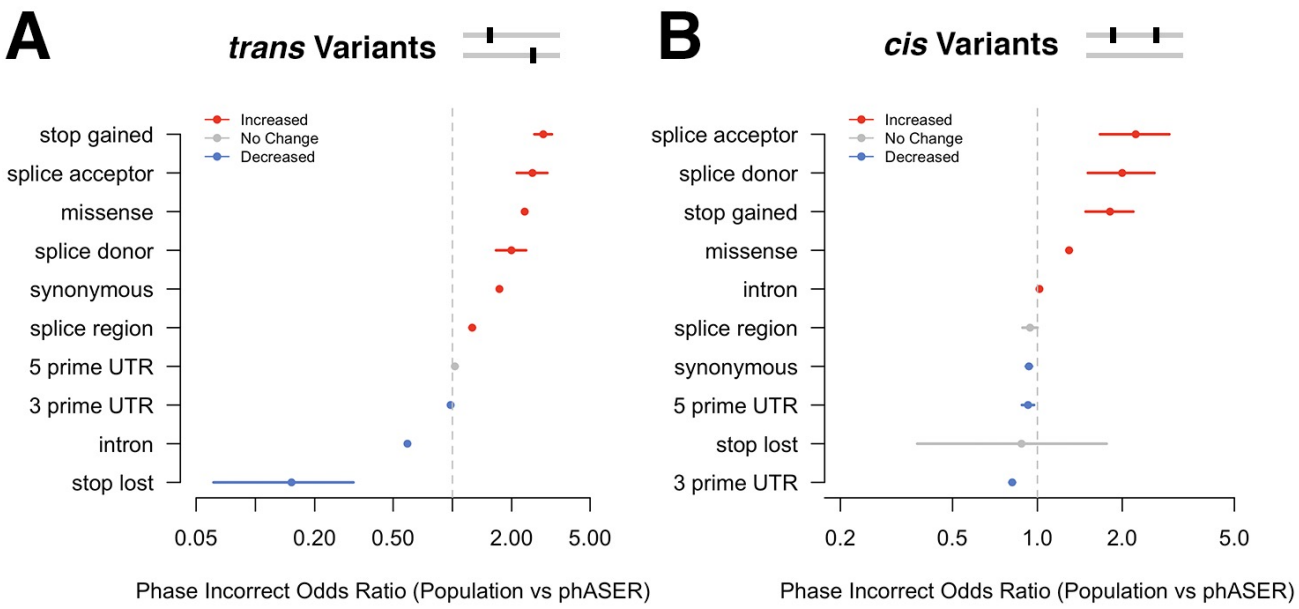

Figure S5. Read backed phasing improves the ability to correctly identify instances of compound heterozygosity at rare variants. Cases of compound heterozygosity were called using either 1000 Genomes Phase 3 population phasing or exome [12] + Geuvadis LCL RNA-seq [15] read backed phasing with phASER. The odds ratio for cases involving each variant type being incorrect in population data versus other types in either *trans* (A) or *cis* (B) interactions was calculated using Fisher's exact test. Variant types with a significantly increased probability of being incorrect when population phasing is used are shown in red, while those with a decrease are shown in blue ($p < 0.05$). Error bars indicate the 95% confidence interval of the odds ratio.

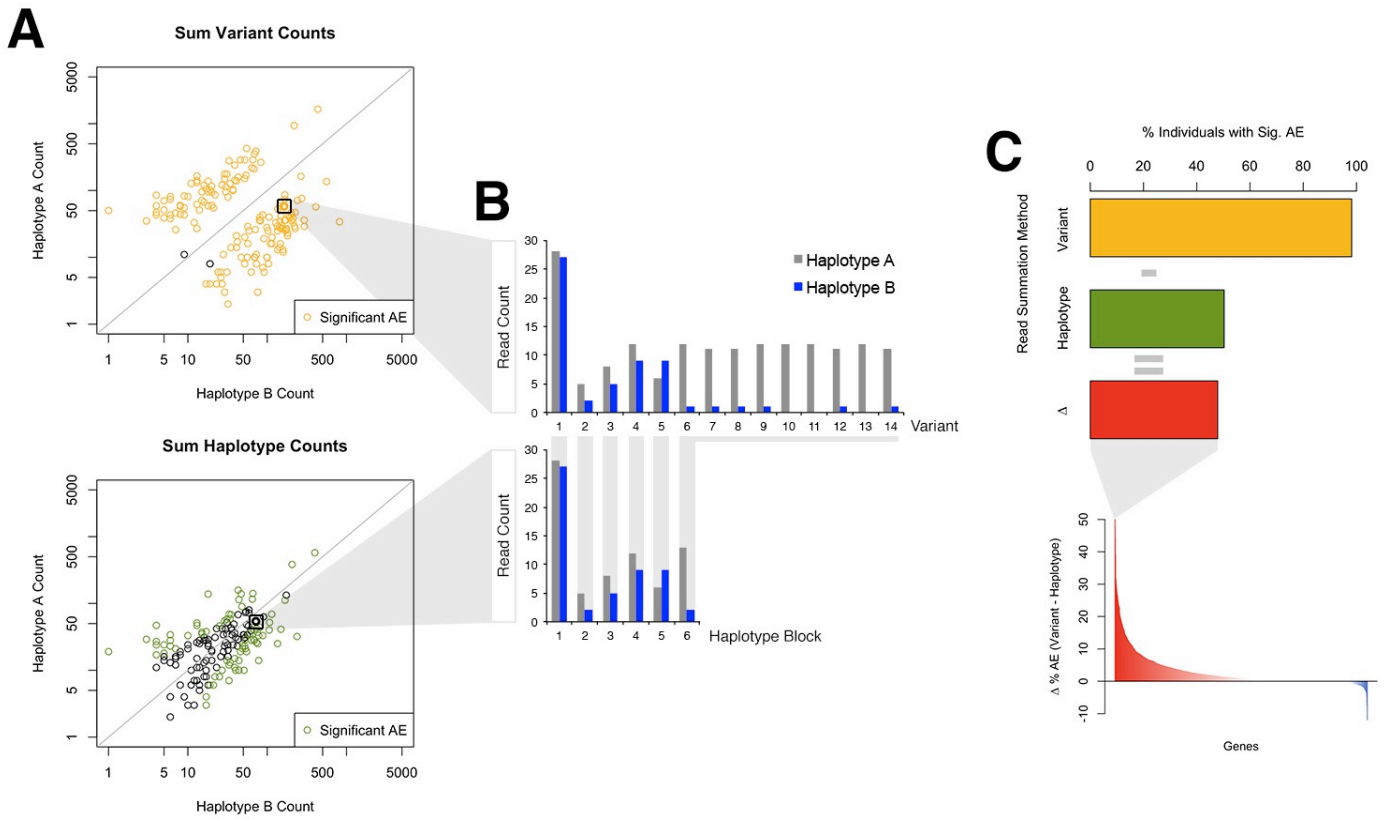

Figure S6. Integrating allelic counts over variants using accurate phasing reduces false positives in allelic expression studies. A) Haplotypic counts for Geuvadis individuals at an example gene (ENSG00000162654 or *GBP4*) calculated by either summing counts from individual variants using 1000 Genomes Phase 3 population phasing [12] (top plot, yellow) or with phASER haplotypic counts (bottom plot, green). Each point represents one individual. B) Example illustrating that summing variant counts for the individual highlighted in (A) leads to double counting of variants 7 through 14 in this haplotype (top plot) and is prevented when haplotypic counts are used (bottom plot). C) For this gene the difference (red) in the percentage of individuals showing significant (5% FDR, binomial test) allelic imbalance calculated either by either summing variants (yellow) or using haplotypic counts (green). This value was calculated for all genes with allelic expression data from at least 30 individuals heterozygous for the top Geuvadis *cis*-eQTL (bottom plot and figure 2c).

Methods

Implementation of read backed phasing in phASER

phASER is written in Python and requires the following libraries: IntervalTree, pyVCF, SciPy, NumPy. In addition it requires Samtools and Bedtools to be installed. Aligned reads (BAM format) are mapped to heterozygous variants through the use of an interval tree, and for each heterozygous variant a hashed set of all overlapping reads is produced, which allows for quick comparison of overlapping reads between variants. Connections between variants are established whenever a read (or read mate) overlaps more than one variant. Each variant – variant connection is tested for the presence of a conflicting phase configuration (see below), and those connections that fail, based on a user defined significance threshold (by default nominal p = 0.01) are removed *. Haplotype blocks are generated by starting with a single unphased variant and recursively adding any other variants with read connections. Block construction is completed when no further variants can be added. Once all haplotype blocks have been generated, phasing between alleles in each block is determined by summing up the total number of connections between alleles that are observed supporting each possible phase configuration (2^n). The number of configurations tested is limited by a user defined setting, which is by default 15 variants = 2^{15} configurations. When the number of variants in a block exceeds this number, the block is split into smaller sub-blocks not exceeding that limit. Once the most supported configuration is established within each sub-block, sub-blocks are phased relative to each other, by again selecting the configuration with the highest support. Variant parsing, read variant mapping, haplotype block construction, and haplotype block phasing can all by parallelized for an increase in speed.

* *It should be noted that this test addresses instances where due to sequencing error individual bases on a read have been misread, it does not address issues with read mapping, which is prone to error at sites with genetic variation. For a discussion of methods to address allelic read mapping issues please see Castel et al., 2015.*

Statistical test for conflicting phasing between two variants

For each SNP pair covered by at least one read a test is performed that determines if the number of reads supporting alternative phasing (any phase other than the configuration chosen by phASER) could be observed by chance from sequencing noise alone. In this test, significance indicates more conflicting reads than would be expected from noise alone, and thus suggests that there may be an error in the phase selected by phASER. A conservative approach is to filter out any variant connections with a nominal p value of < 0.01 (see Figure S1b). Lower p value thresholds can be used to retain more blocks, with the caveat that some may have incorrect phasing. Filtered connections will not be used during the haplotype block construction process.

The test is based on a uniform error model in which a true allele nucleotide can be substituted randomly to any other nucleotide. All pairwise substitutions in this model are assumed to be equally probable. We denote this pairwise substitution probability with ε. Let us assume a pair of SNPs $a_1|a_2$ and $b_1|b_2$ with a haplotype structure $a_1b_1|\ a_2b_2$. Let a_1:b_1 denote a read spanning alleles a_1 and b_1. Reads supporting the true haplotype in this case are a_1:b_1, and a_2:b_2, and any other configurations correspond to conflicting evidence. Let us only consider reads generated from the first haplotype. The probability of observing a read supporting the correct haplotype structures, p_s, is $(1-3\varepsilon)^2+\varepsilon^2$, where the first term corresponds to the probability of observing a_1:b_1 (the case in which either of the sites being affected by noise in the read), and the second term is the probability of observing a_2:b_2 (the case in which both sites were altered by noise in the read and happened to have the other second allele). Binomial distribution is used to evaluate the probability of observing equal or less supporting reads for two given SNP sites:

$$p(x \leq n_s|n,\varepsilon) = \text{binomial}_{cdf}(n_s, n, p_s),$$

where n_s and n are the number of reads supporting the chosen haplotype structure and total number of reads respectively. The pair-wise substitution rate ε is calculated from over all SNP sites as:

$$\varepsilon = 0.5\frac{\sum \text{mismatches}}{\sum \text{total coverage}},$$

where mismatches correspond to all cases that a nucleotide other than the reference or the alternative allele where observed at the site. Only variants where > 50% of the reads come from the reference and alternative alleles are used to calculate the substitution rate to avoid inflation of noise estimates as a result of genotyping error.

Genome wide phasing using phase anchoring

When population phased data is available phASER will attempt to determine the genome wide phase for each haplotype block using phase anchoring. Phase anchoring operates on the assumption that common variants are more likely to be phased correctly using population data, so for each variant within a haplotype the genome wide phase determined by population phasing is weighted by the variant allele frequency. The genome wide phase across all variants with the most support after weighting is selected for each haplotype block.

For each haplotype block:
α (phase support for configuration 1) = Σ(MAF of variants supporting configuration 1)
β (phase support for configuration 2) = Σ(MAF of variants supporting configuration 2)
if $\alpha > \beta$ genome wide phase = configuration 1
if $\alpha < \beta$ genome wide phase = configuration 2
Anchor Phase Confidence = $\max(\alpha, \beta) / (\alpha+\beta)$

phASER settings

The following settings were used unless otherwise noted. For all libraries: alignment score quantile cutoff of 0.05, BASEQ of 10, and conflicting configuration threshold of 0.01, indels ignored. For RNA-seq libraries: no maximum insert size, MAPQ of 255 (indicates unique in STAR). For exome-seq libraries: 500 bp maximum insert size, MAPQ of 40. For whole genome sequencing libraries: 1000 bp maximum insert size, MAPQ of 40. All variants in HLA genes were filtered.

Software availability

Source code and complete documentation for phASER and its associated tools are available through github (https://github.com/secastel/phaser). In addition to the phASER core software we provide two scripts: one which given an input VCF that has been phased using phASER will identify all interactions between alleles (phASER Annotate), and retrieve information such as allele frequency and predicted variant effect if supplied with the appropriate files, and second, a script which will use haplotypic counts produced by phASER in combination with population phasing to produce gene level haplotypic read counts for use in allelic expression studies (phASER Gene AE).

Data processing, usage and availability

For analyses involving 1000 Genomes individuals, phase 3 genotypes and population phasing where used with hg19 aligned exome-seq data, both downloaded from the 1000 Genomes website (http://www.1000genomes.org). Raw (FASTQ) RNA-sequencing data from 1000 Genomes individual derived LCLs was downloaded from the European Nucleotide Archive (ERP001942), and aligned with STAR to hg19. For comparison of phase statistics between sequencing experiments the following GTEx individuals were used: S32W, T5JC, T6MN, WFON. Both short and long read RNA-seq was obtained for whole blood and LCLs, and aligned using STAR to hg19. WES reads were aligned using Bowtie 2 to hg19. GTEx data is available through dbGaP for authorized users (phs000424.v6.p1). GTEx individual ZAB4 was used for comparison of number phased variants versus number of RNA-seq tissues used. For comparison to transmission phasing the following data from the 1000 Genomes individual NA12878 was used: exome-seq downloaded from 1000 Genomes

website, whole genome sequencing data (NCBI SRA ERS179577), RNA-seq from a LCL (NCBI GEO GSM1372331), transmission phased genotypes (Illumina Platinum Genome, http://www.illumina.com/platinumgenomes/).

Benchmarking

Benchmarking was run on CentOS 6.5 with Java version 1.6 and Python 2.7 on an Intel Xeon CPU E7- 8830 @ 2.13GHz, with GATK v3.4 and phASER v0.2. The GATK tool was run with default settings, with the exception of: min_mapping_quality_score = 40, maxPhaseSites = 15, min_base_quality_score = 10, phASER settings were matched to those of GATK. Simulated PE 75 WGS data was produced with ART Chocolate Cherry Cake 03-19-2015 [19] from a NA12878 1000 Genomes Phase 3 population phased reference. WES and WGS libraries used were those listed above.

Acknowledgements

We would like to thank the Geuvadis Consortium, the GTEx Consortium, the members of the Lappalainen lab, and the bioinformatics team of the New York Genome Center. The Genotype-Tissue Expression (GTEx) Project was supported by the Common Fund of the Office of the Director of the National Institutes of Health (commonfund.nih.gov/GTEx). Additional funds were provided by the NCI, NHGRI, NHLBI, NIDA, NIMH, and NINDS. Donors were enrolled at Biospecimen Source Sites funded by NCI\SAIC-Frederick, Inc. (SAIC-F) subcontracts to the National Disease Research Interchange (10XS170), Roswell Park Cancer Institute (10XS171), and Science Care, Inc. (X10S172). The Laboratory, Data Analysis, and Coordinating Center (LDACC) was funded through a contract (HHSN268201000029C) to The Broad Institute, Inc. Biorepository operations were funded through an SAIC-F subcontract to Van Andel Institute (10ST1035). Additional data repository and project management were provided by SAIC-F (HHSN261200800001E). The Brain Bank was supported by a supplement to University of Miami grant DA006227. Statistical Methods development grants were made to the University of Geneva (MH090941), the University of Chicago (MH090951 & MH090937), the University of North Carolina - Chapel Hill (MH090936) and to Harvard University (MH090948).

We are thankful to all the CDH families for their generous contributions. We are grateful for the technical assistance provided by Patricia Lanzano, Jiancheng Guo, Liyong Deng, Badri Vardarajan and Jing He from Columbia University. We also thank Dr. George B Mychaliska and Jeannie Kreutzman from University of Michigan; Dr. Robert Cusick and Sheila Horak from University of Nebraska; Dr. Douglas Potoka, Laurie Luther and Min Shi from University of Pittsburgh; Dr. Gundrun Aspelund and Julia Wynn from Columbia University; Dr. Marc S. Arkovitz from Tel Hashomer Medical Center; Dr. Charles J. Stolar from California Pediatric Surgery Group; Dr. Kenneth S. Azarow from Oregon Health Science Center for enrolling participants. This work was supported by National Institute of Health grant [HD057036] and was supported in part by Columbia University's Clinical and Translational Science Award (CTSA); grant [UL1 RR024156] from National Center for Advancing Translational Sciences/National Institutes of Health (NCATS-NCRR/NIH) a grant from CHERUBS, a grant from the National Greek Orthodox Ladies Philoptochos Society, Inc. and generous donations from The Wheeler Foundation, Vanech Family Foundation, Larsen Family, Wilke Family and many other families.

Author Contributions

S.E.C. and T.L. wrote the manuscript, S.E.C. analyzed the data, S.E.C. and P.M. developed the statistical model, W.K.C and Y.S. provided clinical data.

References

1. Roach, J. C. *et al*. Chromosomal haplotypes by genetic phasing of human families. *Am. J. Hum. Genet*.

89, 382–397 (2011).
2. Delaneau, O., Marchini, J. & Zagury, J.-F. A linear complexity phasing method for thousands of genomes. *Nat. Methods* **9,** 179–181 (2012).
3. Kuleshov, V. *et al*. Whole-genome haplotyping using long reads and statistical methods. *Nat. Biotechnol*. **32,** 261–266 (2014).
4. Berger, E., Yorukoglu, D. & Berger, B. in *Research in Computational Molecular Biology* **9029,** 28–29 (Springer International Publishing, 2015).
5. Castel, S. E., Levy-Moonshine, A., Mohammadi, P., Banks, E. & Lappalainen, T. Tools and best practices for data processing in allelic expression analysis. *Genome Biol*. **16,** 195 (2015).
6. Regan, J. F. *et al*. A rapid molecular approach for chromosomal phasing. *PLoS ONE* **10,** e0118270 (2015).
7. Yang, W.-Y. *et al*. Leveraging reads that span multiple single nucleotide polymorphisms for haplotype inference from sequencing data. *Bioinformatics* **29,** 2245–2252 (2013).
8. Li, X. *et al*. Transcriptome sequencing of a large human family identifies the impact of rare noncoding variants. *Am. J. Hum. Genet*. **95,** 245–256 (2014).
9. GTEx Consortium. Human genomics. The Genotype-Tissue Expression (GTEx) pilot analysis: multitissue gene regulation in humans. *Science* **348,** 648–660 (2015).
10. McKenna, A. *et al*. The Genome Analysis Toolkit: a MapReduce framework for analyzing next-generation DNA sequencing data. *Genome Res*. **20,** 1297–1303 (2010).
11. Melé, M. *et al*. Human genomics. The human transcriptome across tissues and individuals. *Science* **348,** 660–665 (2015).
12. Abecasis, G. R. *et al*. A global reference for human genetic variation. *Nature* **526,** 68–74 (2015).
13. Adzhubei, I. A. *et al*. A method and server for predicting damaging missense mutations. *Nat. Methods* **7,** 248–249 (2010).
14. Yu, L. *et al*. Increased burden of de novo predicted deleterious variants in complex congenital diaphragmatic hernia. *Hum. Mol. Genet*. **24,** 4764–4773 (2015).
15. Lappalainen, T. *et al*. Transcriptome and genome sequencing uncovers functional variation in humans. *Nature* **501,** 506–511 (2013).
16. Baran, Y. *et al*. The landscape of genomic imprinting across diverse adult human tissues. *Genome Res*. **25,** 927–936 (2015).
17. Rivas, M. A. *et al*. Human genomics. Effect of predicted protein-truncating genetic variants on the human transcriptome. *Science* **348,** 666–669 (2015).
18. Browning, S. R. & Browning, B. L. Haplotype phasing: existing methods and new developments. *Nat. Rev. Genet*. **12,** 703–714 (2011).
19. Huang, W., Li, L., Myers, J. R. & Marth, G. T. ART: a next-generation sequencing read simulator. *Bioinformatics* **28,** 593–594 (2012).