METHOD

Haplotype threading: accurate polyploid phasing from long reads

Sven D. Schrinner[1†], Rebecca Serra Mari[2,3,4†], Jana Ebler[2†], Mikko Rautiainen[5,3,4], Lancelot Seillier[7], Julia J. Reimer[7], Björn Usadel[6,7,8], Tobias Marschall[2*†] and Gunnar W. Klau[1,8*†] (iD)

*Correspondence:
tobias.marschall@hhu.de;
gunnar.klau@hhu.de
†Sven D. Schrinner, Rebecca Serra
Mari and Jana Ebler are joint first
authors.
†Tobias Marschall and Gunnar W.
Klau are joint last authors.
²Institute for Medical Biometry and
Bioinformatics, Medical Faculty,
Heinrich Heine University
Düsseldorf, Moorenstraße 5, 40225
Düsseldorf, Germany
¹Algorithmic Bioinformatics,
Heinrich Heine University
Düsseldorf, Universitätsstr. 1, 40225
Düsseldorf, Germany
Full list of author information is
available at the end of the article

Abstract

Resolving genomes at haplotype level is crucial for understanding the evolutionary history of polyploid species and for designing advanced breeding strategies. Polyploid phasing still presents considerable challenges, especially in regions of collapsing haplotypes.

We present WHATSHAP POLYPHASE, a novel two-stage approach that addresses these challenges by (i) clustering reads and (ii) threading the haplotypes through the clusters. Our method outperforms the state-of-the-art in terms of phasing quality. Using a real tetraploid potato dataset, we demonstrate how to assemble local genomic regions of interest at the haplotype level. Our algorithm is implemented as part of the widely used open source tool WhatsHap.

Keywords: Polyploidy, Phasing, Haplotypes, Cluster editing, High-throughput nucleotide sequencing, Plant science, Sequence analysis

Background

Polyploid genomes have more than two homologous sets of chromosomes. Polyploidy is common to many plant species, including important food crops like potato (*Solanum tuberosum*), bread wheat (*Triticum aestivum*), and durum wheat (*Triticum durum*). Resolving polyploid genomes at the haplotype level, i.e., assembling the sequences of alleles residing on the same chromosome, is crucial for understanding the evolutionary history of polyploid species: Evolutionary events, such as whole-genome duplications, can be traced back and reveal the ancestry of polyploid organisms [1]. Beyond that, knowledge of haplotypes is key for advanced breeding strategies or genome engineering, especially for improving yield quality in important crop species [1–3].

In this work, we focus on phasing from long read information. Plant genomes typically exhibit many highly repetitive regions and frequently underwent structural variation events, rendering alignments from short reads alone problematic. Although long reads suffer from a higher number of sequencing errors, they align better to the reference
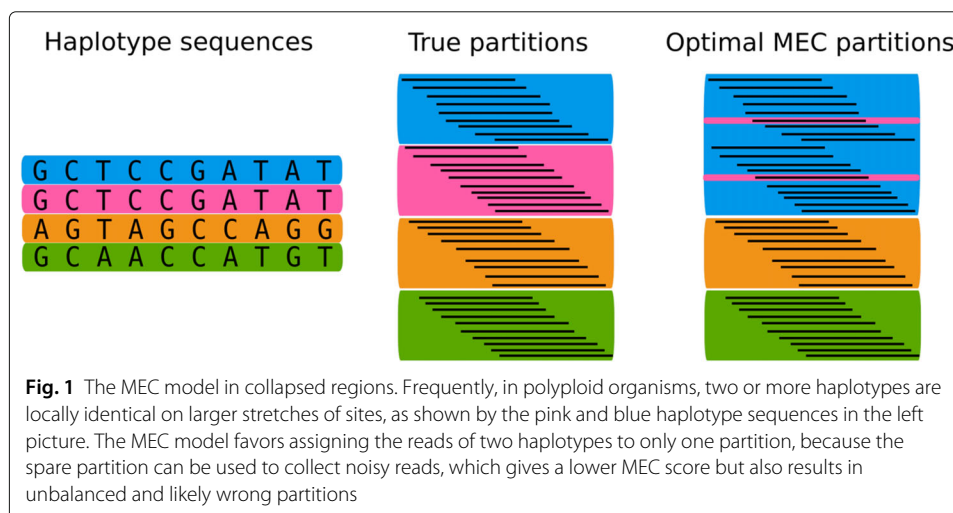
genome and span more variant positions. Consequently, there are larger overlaps between read pairs, which is the key information for molecular phasing methods. This is especially important for polyploid phasing, where the assignment must distinguish not only between two but between k haplotypes.

While phasing diploid genomes using long reads has become a routine step, polyploid phasing still presents considerable challenges [4]. Higher ploidy increases the complexity of the underlying computational problem: In the diploid case, assembling one haplotype over all heterozygous variants directly determines the complementary second haplotype. For genomes of higher ploidy, this is not the case. In addition, polyploid genomes usually exhibit larger regions of two or more identical haplotypes. The Minimum Error Correction (MEC) model [5], which is the most common and successful formalization for diploid haplotype assembly from sequence reads, is not suited to distinguish between locally identical haplotypes. It aims to minimize the number of corrections that are applied to the reads in order to partition them into distinct sets such that reads from the same partition belong to the same haplotype. The MEC score is the minimum number of necessary corrections. In the MEC model, it does not pay off to assign identical haplotypes. Hence, in regions of locally similar haplotypes, this model is likely to result in incorrect haplotype assignments, see Fig. 1. Consequently, MEC-based approaches for polyploid phasing struggle in such regions and, beyond that, face the challenge that exact optimization techniques for MEC [6] quickly become infeasible in practice due to the NP hardness of the problem.

Related work

Throughout the last years, a few polyploid phasing methods have already been proposed. In 2013, Aguiar et al. were the first to introduce a theoretical framework for polyploid haplotype assembly with the HapCompass [7, 8] model, which is based on spanning trees and uses the Minimum Weighted Edge Removal (MWER) criterion. In 2014, Berger et al. introduced HapTree [9], a maximum likelihood approach to discover the most likely haplotypes given aligned read data. To address the problem of computational complexity, the most likely haplotypes are assembled for a small set of SNP positions first and then iteratively extended, keeping only the most likely sub-solutions in each step. HapTree was

Fig. 1 The MEC model in collapsed regions. Frequently, in polyploid organisms, two or more haplotypes are locally identical on larger stretches of sites, as shown by the pink and blue haplotype sequences in the left picture. The MEC model favors assigning the reads of two haplotypes to only one partition, because the spare partition can be used to collect noisy reads, which gives a lower MEC score but also results in unbalanced and likely wrong partitions

shown to outperform HapCompass in terms of accuracy and runtime [9, 10]. Together with SDhaP [11], a semi-definite programming approach based on an approximate MEC criterion, HapCompass and HapTree were evaluated and compared to each other in a simulation study conducted by Motazedi et al. [10] in 2017. The study, where simulated data of the tetraploid potato genome as model organism was used, revealed that, out of the compared methods, HapTree provided the best results in terms of precision. However, it also showed the highest time and memory requirements and often suffered from low recall. SDhaP showed low performance in regions of locally similar haplotypes, which is probably related to the underlying MEC model. For ploidies above six, HapCompass was the only implementation to remain stable, although it showed an overall poor performance. As a result, none of the methods came out to be applicable for practical use due to computational inefficiency that prohibits scaling to large genomic regions as well as frequent failures and low overall accuracy [10].

H-PoP [12] was shown to outperform these previous approaches both in accuracy and runtime and is since then considered as the state-of-the-art method. It consists of a model called Polyploid Balanced Optimal Partition (PBOP) that creates k partitions of sequence reads with the aim to minimize two measures: Reads from one partition are supposed to be equal on as many variant loci as possible, whereas reads from different partitions should contain as many differences as possible. For $k = 2$, this equals the diploid MEC model and can thus be seen as a polyploid generalization of MEC. When genotype information is present, these constraints are added to the model; the appropriate extension is then referred to as H-PoPG.

More recent advances have not proven to be useful for whole-genome single-individual haplotyping, like PolyHarsh [13], a Gibbs sampling method that is also based on the MEC model and has only been shown to work on very small artificial examples, TriPoly [14] that infers haplotypes from trio data and thus requires family data, and SDA [15]. The latter provides two algorithms based on a discrete matrix completion approach and correlation clustering, respectively, and is used to resolve segmental duplications of higher ploidy during genome assembly. However, it is not designed to scale to the whole genome.

Other matrix-based models are SCGD-hap [16], a structurally constrained gradient descent approach, and AltHap [17], which builds on SCGD-hap and aims to solve an iterative sparse tensor decomposition problem. This model yielded results similar to those of H-PoP, but also relies on MEC.

Some tools have been proposed that do not work well with long read data. The work by [18] is based on minimum fragment removal. The long and relatively erroneous long reads would lead to a removal of too much data. Ranbow [1] uses allele co-occurrences on small sets of sampled positions in overlapping short reads. This approach is susceptible to high error rates found in long reads, as it seeds the phasing on local partitions of reads based on their allele combination on the small position samples. Thus, a large portion of the reads are clustered incorrectly and a lot of overlapping position samples are required to correct these mistakes.

Apart from the limitations of the underlying model, current methods do not give reliable information about the accuracy of the resulting haplotypes since these are either output in one consecutive sequence or in very long blocks. In particular, this means that there is no information about the positions of likely switch errors. Thus, large regions of the resulting haplotypes might be wrong, but it is not possible to identify these regions,

which makes the results very difficult to use in practice. In the H-PoPG algorithm, for example, the haplotypes are only divided into blocks if there is no read that covers the affected neighboring variant loci. Further uncertainties in the phasing are not considered in the model and thus not reported.

Contribution

We present an accurate model for polyploid phasing that produces reliable haplotype blocks and is computationally efficient and thus applicable in practice. We introduce WHATSHAP POLYPHASE, a method that departs from the MEC model in order to deal with the additional challenges arising in polyploid phasing. By taking coverage into account via a newly established threading step, WHATSHAP POLYPHASE is able to detect and properly phase regions where multiple haplotypes coincide. Additionally, our method is able to integrate information from input genotypes for accurate phasing results.

We introduce cuts within the haplotypes at positions with increased phasing uncertainty and thereby output phased blocks that ensure high accuracy within the fragments. We provide a sensible way to compute these block boundaries at varying, user-defined degrees of strictness. This way, we enable a configurable trade-off between longer blocks that potentially contain errors and shorter but highly accurate blocks.

We demonstrate on a simulated dataset of varying ploidy $k \in \{4, 5, 6\}$ that WHATSHAP POLYPHASE outperforms the state-of-the-art tool H-PoPG in terms of phasing quality, in particular in regions with two or more identical haplotypes, where our method phases with around seven times lower switch error rates than the competition. The efficient implementation of WHATSHAP POLYPHASE allows for scaling to gigabase-sized genomes, while being sufficiently fast: an artificial human tetraploid Chromosome 1 (249 Mb) is phased in less than 3.5 h on a single core of a standard desktop. We also show the usefulness of our method to phase real long read data from a tetraploid potato. We first correct the long reads, phase them with WHATSHAP POLYPHASE, and then locally assemble the partitioned reads. We show that this strategy is suitable to phase a major fraction of genes.
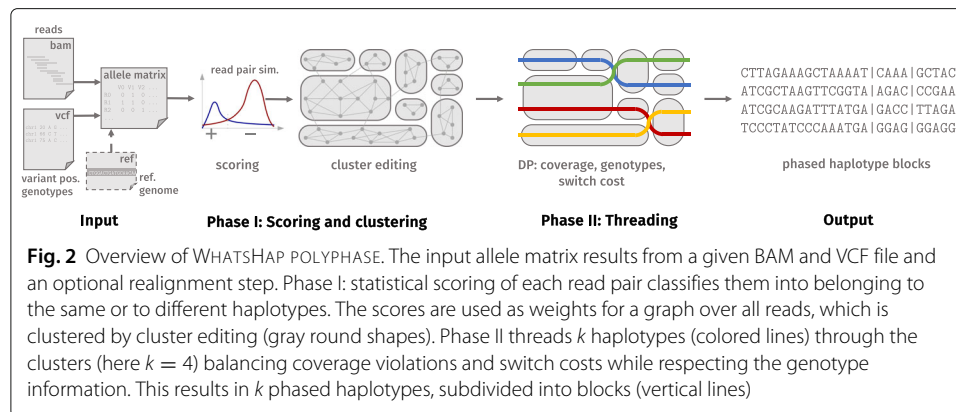
WHATSHAP POLYPHASE is implemented as part of the widely used open source tool WHATSHAP (https://whatshap.readthedocs.io) and therefore ready to be included in production settings. It offers convenient usage by supporting standard input and output formats (BAM and VCF). WHATSHAP POLYPHASE is available at https://github.com/whatshap/whatshap.

Results

Phasing model and algorithm

WHATSHAP POLYPHASE is a novel two-stage approach that produces accurate haplotypes for polyploid genomes using data from single-molecule sequencing technologies. See Fig. 2 for an overview of the method.

The first phase of the algorithm uses cluster editing [19] to find clusters of reads which are likely to originate from identical haplotypes. In short, this is done by computing a statistical similarity score for each pair of reads and constructing a graph using the reads as nodes and the scores as edge weights [20]. The size of the graph makes it infeasible to solve cluster editing to optimality in reasonable time, so we rely on an iterative heuristic to produce accurate clusters. We deliberately make no assumptions on the ploidy at the

Fig. 2 Overview of WHATSHAP POLYPHASE. The input allele matrix results from a given BAM and VCF file and an optional realignment step. Phase I: statistical scoring of each read pair classifies them into belonging to the same or to different haplotypes. The scores are used as weights for a graph over all reads, which is clustered by cluster editing (gray round shapes). Phase II threads k haplotypes (colored lines) through the clusters (here $k = 4$) balancing coverage violations and switch costs while respecting the genotype information. This results in k phased haplotypes, subdivided into blocks (vertical lines)

clustering stage. In particular, reads of multiple haplotypes that are locally identical end up in the same cluster.

The second phase consists of the actual haplotype assembly by threading k haplotypes through the set of clusters obtained in the first phase. We take the position-wise read coverage of each cluster into account to determine the number of haplotypes threaded through each cluster. In contrast to MEC-based models, this handles genomic regions where some haplotypes are locally identical by allowing that multiple haplotypes run through the same cluster. During the threading step, we further expect haplotypes to stay in the same cluster for as long as possible and ensure that the consensus genotype fits the input genotype, if provided. We cut the phasing into blocks at variant pairs showing insufficient phasing confidence to increase its accuracy at the cost of decreased phasing block lengths. See the "Methods" section for details of WHATSHAP POLYPHASE.

WHATSHAP POLYPHASE produces accurate results

To demonstrate that WHATSHAP POLYPHASE works well in practice, we ran it on an artificial tetraploid dataset at different coverages and compared our results to those of H-POPG, the state-of-the-art method for polyploid phasing, as well as to HapCompass, HapTree, and Ranbow. Since HapCompass, HapTree, and Ranbow were not able to process whole chromosomes, we applied them to an exemplary 1-Mb-long region from Chromosome 1. The results can be found in Additional file 1: Table S1, showing that HapCompass, HapTree, and Ranbow performed considerably worse than WHATSHAP POLYPHASE and H-POPG. We therefore focus our further comparative analysis on H-POPG. We used common evaluation statistics that capture different properties of haplotype sequences to compare the solutions computed by both tools to ground truth haplotypes available for our datasets.

Evaluation statistics

For ploidy k, a set of ground truth haplotype sequences $h = \{h_1, ..., h_k\}$ and predicted haplotypes $h^* = \{h_1^*, ..., h_k^*\}$, we compute the number of Hamming errors HE as

$$\text{HE} = \min_{\sigma \in S_k} \frac{1}{k} \sum_{i=1}^{k} d_H\left(h_i, h_{\sigma(i)}^*\right)$$

where S_k represents the permutation group on $\{1, ..., k\}$ and $d_H()$ the Hamming distance between two sequences. The *Hamming rate* (HR) is then defined as the sum of Hamming

errors divided by the total number of all phased variants. If subtracted from 1, the Hamming rate is equivalent to the *reconstruction rate* and the *correct phasing rate* presented in [14] and [12], respectively.

A well-established evaluation metric for diploid phasing is the *switch error rate* (SER), for which we use a polyploid version. Instead of counting the number of incorrect alleles on each haplotype, the SER counts the minimum number of switches, i.e., how often the assignment between predicted and true haplotypes must be changed in order to reconstruct the true haplotypes from the predicted ones. The polyploid extension of the switch error was already introduced as the *vector error rate* in [9].

More formally, for every position j, let Π_j be the set of one-to-one mappings between h and h^*, such that for each $\pi \in \Pi_j$ it holds that $h_i[j] = h^*_{\pi(i)}[j]$ for all haplotypes h_i. The switch error rate is then defined as:

$$\text{SER} = \min_{(\pi_1,\dots,\pi_m) \in \Pi_1 \times \dots \times \Pi_m} \frac{1}{k(m-1)} \sum_{i=1}^{m-1} d_S(\pi_i, \pi_{i+1})$$

where m is the number of variants and $d_S(\pi_i, \pi_{i+1})$ the number of different mappings between π_i and π_{i+1}.

If the genotype of h^* is not equal to the genotype of h for every position, the set $\Pi_1 \times \dots \times \Pi_m$ is empty and the vector error cannot be computed. Therefore, we compare only those positions, on which the predicted genotype is correct, and report the fraction of *missing variants* (MV), that is, either unphased or incorrectly genotyped variants, separately.

Phasing tools may not phase the entire input region as one set of haplotypes. If the phasing between two consecutive variants is too uncertain (e.g., if not enough reads cover both variants), the phasing might be split into *blocks*. In our evaluation, we applied the HR and SER on all reported phasing blocks separately and aggregated them. More precisely, we summed up the number of respective errors and divided them by the total number of variants (HR) or by the total number of variants excluding the first variant in every block (SER). Since this favors shorter blocks, we also included the N50 block length into our evaluation, which is the smallest block length needed to cover 50% of the considered genomic region when using only blocks of that size and larger.

Testing on artificial polyploid humans

We generated a tetraploid, pentaploid, and hexaploid versions of human Chromosome 1 by combining sequencing data of three individuals (NA19240, HG00514, and HG0733), for which high-quality trio-based haplotype information is available [21]. We refer to these haplotypes as ground truth haplotypes. We merged PacBio sequencing data for the first two samples to produce tetraploid data at different coverages ($40\times$ and $80\times$). Using the read simulator PBSIM [22], we additionally generated equivalent simulated tetraploid, pentaploid, and hexaploid datasets with the same coverages and known read origin.

We ran WHATSHAP POLYPHASE and H-POPG and compared the resulting phasings to the ground truth haplotypes. H-POPG defines phased blocks based on the connected components of the underlying reads by introducing cuts between pairs of variants not connected by any sequencing reads. Per default, WHATSHAP POLYPHASE uses a more sensitive approach (see the "Methods" section), typically leading to shorter but more accurate haplotype blocks. Additionally, our algorithm supports different levels of block

cut sensitivities, which allow to balance block length against block accuracy. In order to provide a better comparison of both tools, we ran WHATSHAP POLYPHASE with different configurations, which can be seen in Fig. 3.

Even when forcing our tool to yield block lengths as computed in H-POPG, we observe around 30% lower switch error rates among the tested datasets (Fig. 3, see Additional file 1: Figure S1 for Hamming error rates). As expected, higher coverage has a positive effect on the error rates. More sensitive block cuts, and in particular the default setting for WHATSHAP POLYPHASE, lead to a significant decrease in switch error rates.

Table 1 shows all used evaluation metrics on H-POPG and WHATSHAP POLYPHASE for their default settings. We can see that WHATSHAP POLYPHASE phases more accurately, with at least three times lower switch error rates than H-POPG on the varying datasets when using the default settings. For the Hamming rate, the differences are even larger. Among other reasons, this is caused by the block cut policy of H-POPG, leading to switch errors on sparsely connected variants, which have a big impact on the global correctness of the phasings. The configurable block cut strategy of WHATSHAP POLYPHASE allows to maintain accurate blocks with low Hamming rates. For a comparison with more similar block lengths, we also tested the datasets using a different setting of our tool, which can be achieved by running the tool with "-B 1" as an additional parameter. The N50 block lengths and Hamming rates become almost identical, while the switch error rates remain

Fig. 3 N50 block lengths and the respective block-wise switch error rates for different block cut strategies of WHATSHAP POLYPHASE (default strategy marked by a circle) on the real tetraploid read dataset (top) and the simulated tetraploid dataset (bottom) with 40× and 80× coverage
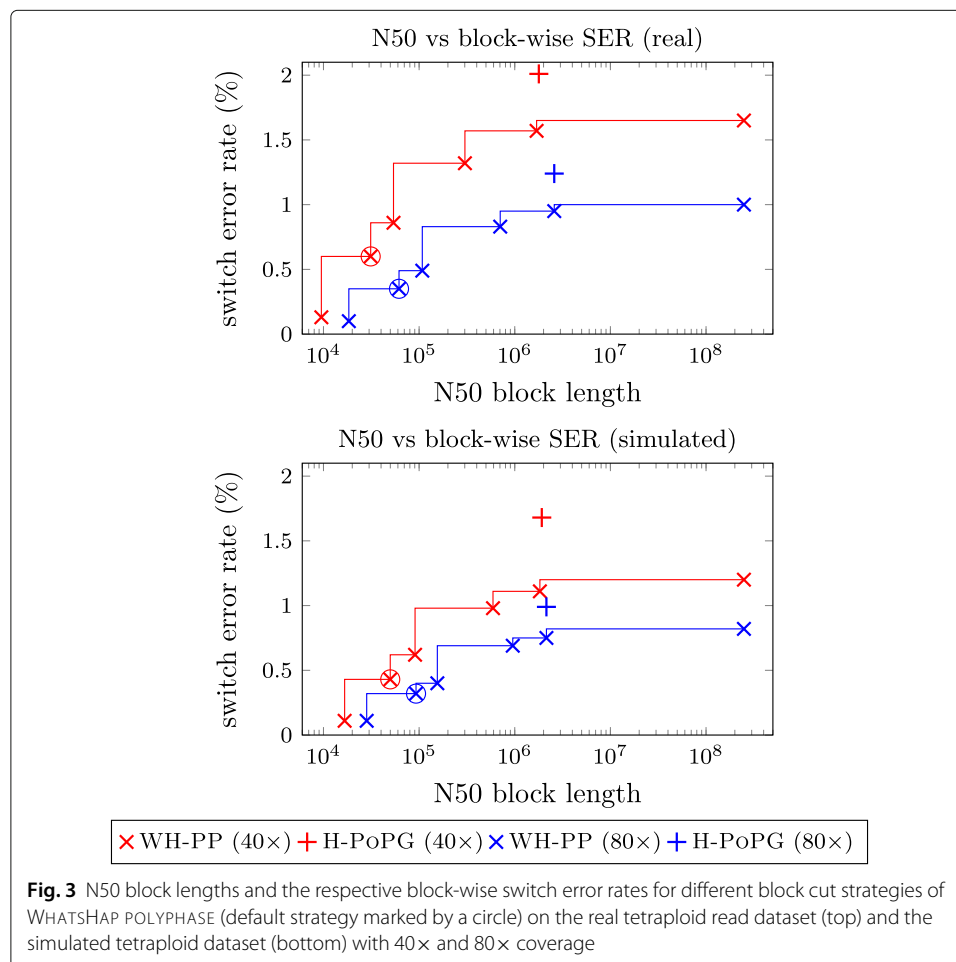
Table 1 Comparison of WHATSHAP POLYPHASE and H-PoPG on tetraploid real (a) and simulated (b) datasets, pentaploid simulated dataset (c), and hexaploid simulated dataset (d). Performances are based on the switch error rate (SER), block-wise Hamming rate (HR), and N50 for the block size. For better comparability with H-PoPG, a second setting (WH-PP*) with less block cuts was used. The total length of the chromosome is 249 Mb

Coverage	Method	SER (%)	HR (%)	N50 (bp)	Runtime (s)	Memory (GB)
		(a) Real tetraploid read data				
40×	WH-PP	0.58	1.48	29,529	3333	1.41
	WH-PP*	1.39	28.72	1,692,352	3433	1.42
	H-PoPG	2.01	27.53	1,785,293	2230	9.97
80×	WH-PP	0.31	1.43	54,434	12694	2.52
	WH-PP*	0.74	28.27	2,587,104	13042	2.89
	H-PoPG	1.24	27.66	2,587,104	4368	9.99
		(b) Simulated tetraploid read data				
40×	WH-PP	0.42	1.74	48,815	1960	1.10
	WH-PP*	1.00	26.57	1,830,943	2004	1.17
	H-PoPG	1.67	26.37	1,917,094	1414	9.96
80×	WH-PP	0.29	2.51	86,227	5738	1.78
	WH-PP*	0.68	25.23	2,142,893	5865	2.04
	H-PoPG	0.98	25.65	2,142,893	2843	9.97
		(c) Simulated pentaploid read data				
40×	WH-PP	0.86	1.57	22,625	2331	1.05
	WH-PP*	2.01	25.34	1,361,459	2377	1.07
	H-PoPG	3.50	24.78	1,453,040	2357	9.97
80×	WH-PP	0.47	1.18	33,438	5031	1.69
	WH-PP*	1.33	23.64	1,701,753	5118	1.87
	H-PoPG	2.24	24.76	1,748,404	4849	9.96
		(d) Simulated hexaploid read data				
40×	WH-PP	1.12	1.82	16,785	25841	1.30
	WH-PP*	2.35	27.03	3,877,456	25860	1.79
	H-PoPG	3.85	26.75	4,490,129	5450	9.96
80×	WH-PP	0.48	0.97	26,711	10331	1.98
	WH-PP*	1.34	25.63	4,540,968	10827	2.63
	H-PoPG	2.37	25.93	4,721,421	11563	10.89

30–40% lower than those of H-PoPG . Both configurations need noticeably more time to compute the phasings than H-PoPG . Higher coverages increase the running time more than linearly for WHATSHAP POLYPHASE as well as ploidy does for both tools. H-PoPG consumed more memory while running our experiments.

The haplotype blocks output by our tool remain quite short when applying the default settings for accurate blocks. One reason for this is the fact that the phaser always discontinues all haplotypes when it decides that two variants cannot be connected with enough confidence. This is due to a limitation of the VCF format, where it is common to only supply one phase set identifier per variant. WHATSHAP POLYPHASE is also able to output these phase set identifiers per haplotype in a custom format (HS). This results in haplotype-level blocks that are between 1.8 and 2.4 times as long as the original blocks, because not all haplotypes need to be interrupted in the polyploid case, if there is just a possible switch between two of them. We show the effect of this alternative phase block definition in Additional file 1: Table S3.

Identifying collapsing regions

We define regions in the genome where two or more haplotypes share the same sequence for at least 50 variant positions as *collapsing regions*. For MEC-based approaches, these parts are particularly difficult to phase since different configurations of haplotypes with locally identical sequences are not distinguishable based on their MEC scores and the MEC model exploits this to explain sequencing errors with "noise" haplotypes.

We evaluated the ability to correctly assemble haplotypes in these regions. We compared WHATSHAP POLYPHASE and H-POPG on Chromosome 1 of the simulated and real tetraploid datasets with 40× and 80× coverage, respectively. Collapsing regions take up a large part (17.28%) of the simulated Chromosome 1.

Table 2 shows the results. It can be seen that the differences between switch error rates achieved by H-POPG and by WHATSHAP POLYPHASE are remarkably higher in the case of collapsing regions than for the rest of the genome. In comparison to WHATSHAP POLYPHASE, the switch error rate of H-POPG is around 7 times higher in collapsing regions, while on average throughout the whole chromosome, this factor is only 3.37. For higher coverage, these values are further increased to 7.5 and 3.5, respectively. The closest results are achieved in non-collapsing regions, i.e., regions where either all haplotype sequences are unique or coincide on fragments shorter than 50 variants. In these regions, H-POPG results in 3.13 times more switch errors.

For the simulated data (see Table 2), the differences are even more striking, especially on 80× coverage. In regions with coinciding haplotypes, WHATSHAP POLYPHASE outperforms H-POPG by a factor of up to 11.75. Compared to the average quotient of 3.09, WHATSHAP POLYPHASE thereby yields an almost 4 times higher reduction in switch error rates in collapsing regions. On lower coverage, similar results are obtained.

As for the previous experiments, we repeated this analysis with block lengths computed as in H-POPG. The results of this second run are presented in Additional file 1: Table S2.

Table 2 Comparison between the resulting switch error rates of WHATSHAP POLYPHASE (WH-PP) and H-POPG on collapsing regions over at least 50 variants as compared to non-collapsing regions and the average throughout the genome. Results (switch error rates in %) are presented for Chromosome 1 of the real and simulated tetraploid dataset on both 40× and 80× coverage. The third row marks the quotient between the switch error rate of H-POPG and that of WHATSHAP POLYPHASE to highlight by which magnitude the results differ

Coverage	Method	Collapsing regions	Non-collapsing regions	Total
		(a) Real read data (tetraploid)		
40×	WH-PP	0.29	0.69	0.60
	H-POPG	2.02	2.16	2.02
	SER($\frac{\text{H-POPG}}{\text{WH-PP}}$)	6.97	3.13	3.37
80×	WH-PP	0.14	0.46	0.35
	H-POPG	1.05	1.30	1.24
	SER($\frac{\text{H-POPG}}{\text{WH-PP}}$)	7.50	2.83	3.54
		(b) Simulated read data (tetraploid)		
40×	WH-PP	0.18	0.45	0.43
	H-POPG	2.01	1.63	1.68
	SER($\frac{\text{H-POPG}}{\text{WH-PP}}$)	11.17	3.62	3.91
80×	WH-PP	0.08	0.37	0.32
	H-POPG	0.94	0.98	0.99
	SER($\frac{\text{H-POPG}}{\text{WH-PP}}$)	11.75	2.65	3.09

Potato data

We applied our algorithm to real sequencing data for tetraploid potato (*Solanum tubero-sum*), for which we generated paired-end short Illumina and long Oxford Nanopore reads. In the first step, we aligned the reads produced by the different technologies to the potato reference genome published by the Potato Genome Sequencing Consortium (PGSC) [24]. We observed unbalanced coverage distributions for the alignments, especially for the short Illumina reads, hinting towards a high number of structural variations and rearrangements being present in the data (Fig. 4a). Thus, the Illumina reads are ill-suited for reliable variant calling as their short length makes it more difficult to unambiguously
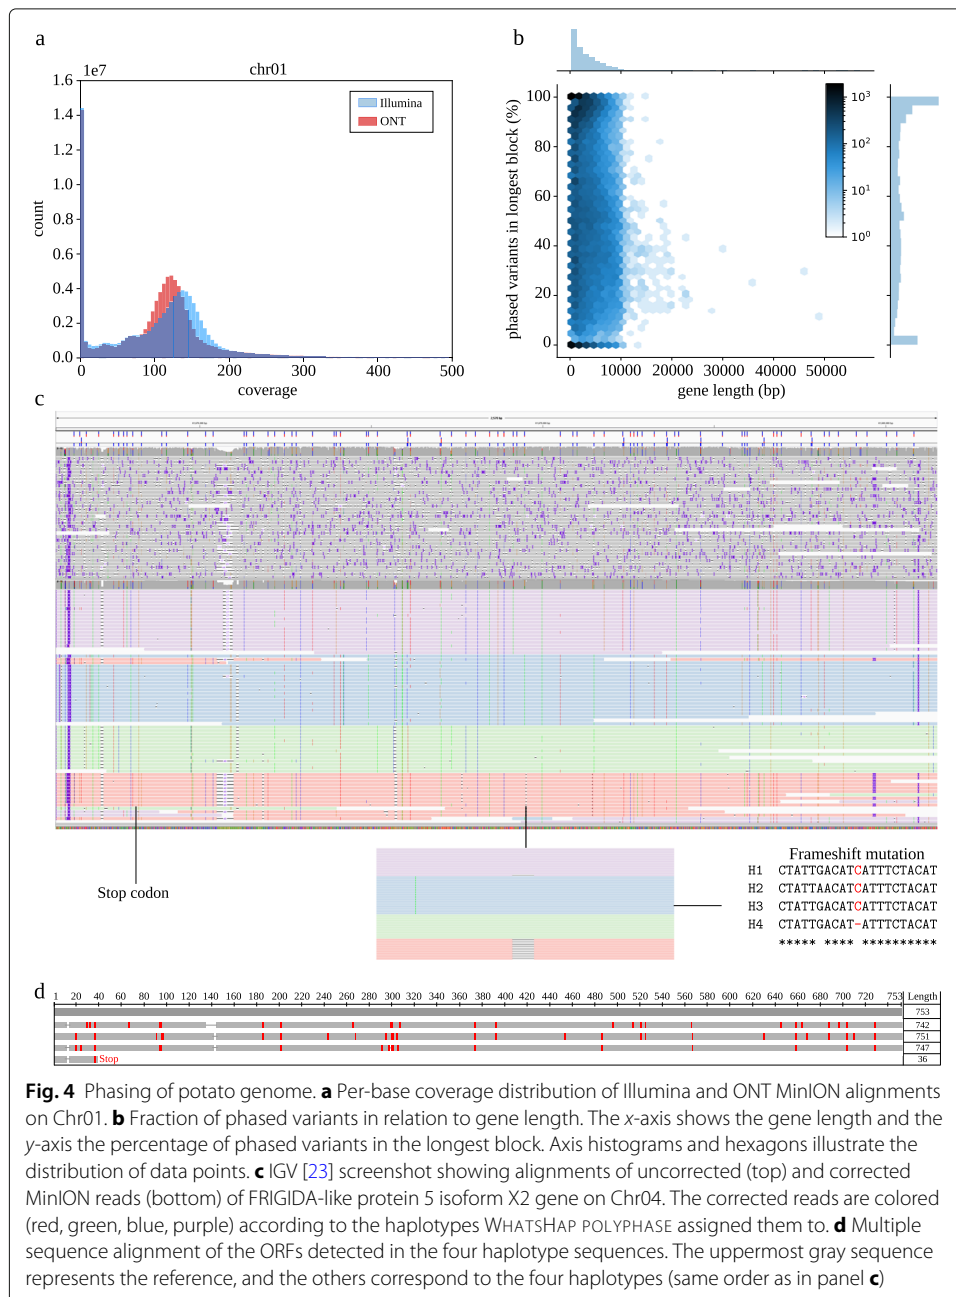

Fig. 4 Phasing of potato genome. **a** Per-base coverage distribution of Illumina and ONT MinION alignments on Chr01. **b** Fraction of phased variants in relation to gene length. The *x*-axis shows the gene length and the *y*-axis the percentage of phased variants in the longest block. Axis histograms and hexagons illustrate the distribution of data points. **c** IGV [23] screenshot showing alignments of uncorrected (top) and corrected MinION reads (bottom) of FRIGIDA-like protein 5 isoform X2 gene on Chr04. The corrected reads are colored (red, green, blue, purple) according to the haplotypes WHATSHAP POLYPHASE assigned them to. **d** Multiple sequence alignment of the ORFs detected in the four haplotype sequences. The uppermost gray sequence represents the reference, and the others correspond to the four haplotypes (same order as in panel **c**)

align them to the reference. We therefore relied on the much longer nanopore reads to identify SNPs that we can later use for phasing. However, Oxford Nanopore reads typically come with high sequencing error rates, complicating the calling process. In order to obtain reliable variant positions and genotypes from these error-prone reads, we ran an error correction pipeline [25] to reduce the number of sequencing errors (see the "Methods" section). Figure 4c shows an exemplary IGV [23] screenshot of uncorrected reads (top) and corrected (bottom) for the FRIGIDA-like protein 5 isoform X2 gene. Next, we ran minimap2 [26] to align the corrected nanopore reads to the potato reference genome and called variants using FreeBayes [27]. To verify the genotypes produced in this way, we added an additional step to WHATSHAP POLYPHASE that re-genotypes the positions based on the nanopore reads prior to phasing and only keeps those variants, for which the newly predicted genotype matches the one reported by FreeBayes (see the "Methods" section).

We focused on the potato genes [24] as they are biologically interesting for phasing. Of the total 36274 genes containing heterozygous variants after calling and retyping, 91% could be (at least partially) phased by WHATSHAP POLYPHASE. On average, about 2.13 phased blocks were produced per gene. Furthermore, for each gene, we determined the number of phased variants inside the longest phased block. We observe that a large fraction of genes, including many long genes, can be fully phased, see Fig. 4b. We also evaluated the percentage of phased genes in relation to their level of heterozygosity, but could not observe a strong dependency, see Additional file 1: Figure S2.

We used the FRIGIDA-like protein 5 isoform X2 (accession: XP_015169713) gene as an example to demonstrate how WHATSHAP POLYPHASE enables haplotype-resolved assembly. We extracted the phasing of the longest phasing block reported for this gene and separated the reads by haplotype. In order to do so, we extended the commands `whatshap haplotag` and `whatshap split`, previously implemented in the diploid version of `whatshap`, to higher ploidies. Briefly, the idea is to tag each sequencing read according to the computed haplotype sequence it is most similar to and separate the reads based on these tags (see "Methods" section). The reads shown in Fig. 4c are colored according to the resulting haplotype assignments. In the next step, we separately ran wtdbg2 [28] on each haplotype-specific read set to produce local assemblies of the four haplotypes. Additional file 1: Figure S3 shows a visualization of a multiple sequence alignment of these haplotypes. We ran the NCBI ORFfinder [29] on each of the assemblies and detected a long ORF in the first three haplotypes representing the FRIGIDA like coding sequence. For the fourth haplotype, we could not detect a corresponding ORF, as the putative FRIGIDA gene in the fourth phase showed an early STOP codon highlighted in Fig. 4c. Interestingly, the fourth phase showed an additional frameshift mutation shown in the inset of Fig. 4c where only the phasing information provides the information that this is linked to the premature STOP codon highlighting the necessity of (local) phasing to understand gene architecture. Using COBALT [30], we generated multiple sequence alignments of the amino acid sequences resulting from these three ORFs and the corresponding reference sequence (Fig. 4d). The three sequences show an overall good alignment with the reference with small differences, which may serve as an input for functional follow-up studies.

Runtimes

We show the runtimes of WHATSHAP POLYPHASE and H-POPG for phasing the artificial human Chromosome 1 in Table 1. Both programs were run on a single core on a dual socket machine (2×Intel Xeon E5-2670 v2) with 256 GB of memory. At coverage 40×, WHATSHAP POLYPHASE took about 49 min to phase the real data, while H-POPG took about 30 min. WHATSHAP POLYPHASE phased the simulated dataset in about 30 min and H-POPG in 20 min. At coverage 80×, WHATSHAP POLYPHASE took 3.2 h on the real data and H-POPG 1 h. On the simulated data, WHATSHAP POLYPHASE took 1.5 h and H-POPG 41 min for phasing.

Discussion

We introduce WHATSHAP POLYPHASE, a novel two-stage algorithm enabling accurate haplotype phasing of polyploid genomes. Our model consists of two phases performing a clustering of the reads based on their similarity and assembling the final haplotypes through the resulting clusters. Unlike approaches based on solving the MEC problem, WHATSHAP POLYPHASE takes coverage of the read clusters into account to resolve regions with multiple coinciding haplotypes. Additionally, the phasing can be cut at low confident positions to maximize phasing accuracy. In our view, such stringent block boundaries are warranted to avoid phase connections that essentially represent guessing, for instance when there is only one read connecting two variants in a tetraploid setting while three haplotypes remain unobserved (Additional file 1: Figure S4).

Applying our algorithm to Chromosome 1 of polyploid datasets created of human samples NA19240, HG00514, and HG0733 shows that in comparison with H-POPG, the current state-of-the-art phasing method, WHATSHAP POLYPHASE returns 3.4-fold to 4.9-fold lower switch error rates on real and simulated data at varying ploidies and coverages (Table 1). As default setting, we prescribe a stringent block boundary criterion in WhatsHap and hence the returned blocks are shorter than for H-POPG. Relaxing this criterion, even though not recommended in our view, leads to block sizes similar to H-POPG at still 1.4-fold to 1.7-fold lower switch error rates.

The Hamming rate is more sensitive than the switch error rate, because a single switch on two haplotypes in the middle of a block can potentially cause 50% of the variants being phased wrongly on the two affected haplotypes. While WHATSHAP POLYPHASE shows low Hamming error rates ranging from 0.97 to 2.51% (Table 1), H-POPG displays much higher rates between 24.76 and 28.72%, highlighting again that overly aggressive block-connecting strategies lead to errors. When WHATSHAP POLYPHASE is run with relaxed block boundary criteria, the Hamming error rates become comparable to H-POPG.

Furthermore, we show that the coverage-aware approach of haplotype threading is able to resolve regions where multiple haplotypes coincide, which occur frequently in polyploid genomes. A comparison to H-POPG shows that WHATSHAP POLYPHASE performs particularly well in these regions. The switch error rates are 7 times higher in H-POPG for the real data and more than 11 times higher for the simulated dataset. When using larger blocks according to the block definition of H-POPG, the switch error rate of H-POPG is still more than 3 times higher in these collapsing regions as opposed to 1.22 times, on average. Within Chromosome 1 of our simulated dataset, with a total length of 249 MB, we found around 17% of the genome to be part of long collapsing regions over at least

50 variants. These results clearly highlight the limitation of MEC-based approaches with regard to these regions and the need for phasing methods that address this problem.

Finally, we present a typical use case of polyploid phasing using real sequencing data of potato. Due to the high genomic diversity and lack of high-quality reference sequences, large-scale polyploid phasing remains challenging. We restricted our analysis to the gene regions and use the FRIGIDA-like protein 5 isoform X2 gene as an example to demonstrate that our polyploid phasing tools enable haplotype-resolved assembly of polyploid organisms.

Conclusions

Polyploid phasing is a difficult technological and computational problem. Current state-of-the-art tools rely on the Minimum Error Correction model, which is successful for diploid phasing, but has limitations in the conceptually and computationally far more complex polyploid case. We provide an implementation that departs from the MEC paradigm and instead uses a novel clustering and threading method, taking coverage and genotype information into account. Doing so, it represents the first algorithm designed to specifically handle locally identical haplotypes and, in consequence, performs significantly better in such regions than the state-of-the-art. To our knowledge, it is also the first approach that offers a configurable trade-off between the lengths of phased haplotype blocks and phasing accuracy to fit the user's individual needs. Our implementation scales to whole genomes while being sufficiently fast.

Current challenges lie in resolving more switch error locations, as they either lead to block cuts or to switch errors, which have a high impact on the Hamming rate. Also, the running time of our approach scales exponentially with increasing ploidy. To mitigate this, we have implemented heuristics that allow handling even higher ploidies than the hexaploids evaluated here. We plan to test and validate these techniques in future work. Furthermore, our haplotype threading framework is well suited to be extended to use local ploidy estimates. That would allow phasing genomic regions where the observed ploidy deviates, for instance due to aneuploidies, large deletions on a subset of haplotypes or collapses of segmental duplications in the used reference genome, and we plan to add this functionality to future releases.

Of note, alignment-based phasing methods heavily depend on the quality of the alignments and the subsequent variant calls. In case of strong deviations from the reference genome, as, for example, in large regions of our proof-of-concept potato phasing study, any alignment-based method that relies on the reference genome will struggle. On good quality reference genomes such as the artificial tetraploid benchmark genome proposed in this paper, we show that our method WHATSHAP POLYPHASE delivers haplotype reconstructions with significantly lower error rates compared to the state-of-the-art tool H-POPG. Our algorithm is implemented as part of the widely used open source tool WHATSHAP and is hence ready to be included in production settings.

Methods

Here, we present the phasing algorithm in detail. We denote with k the ploidy of the phased genome, with n the number of heterozygous variants in the genomic region of interest, and with m the number of reads. We assume that all variants are biallelic, denoting the major allele with 0 and the minor allele with 1. Each read r is represented by a

sequence $r_0, \ldots, r_{n-1}$ of length n over the alphabet $\Sigma = \{0, 1, -\}$ such that r_i is the allele for the ith variant and "$-$" indicates an uncovered variant. We use $\text{olp}rs = |\{i \mid r_i, s_i \in \{0, 1\}\}|$ to denote the size of the overlap (number of shared variants) between two reads r, s and $\text{dis}rs = |\{i \mid r_i, s_i \in \{0, 1\}, r_i \neq s_i\}|$ for the number of disagreements between r and s. The ratio between these values is a value between 0 and 1 and called the *Hamming rate* between two reads. The true (and to us unknown) haplotype of a read r is denoted as $H(r) \in \{0, \ldots, k-1\}$. The objective is to find k sequences $H'_0, \ldots, H'_{k-1}$ of length n over Σ, which are close or identical to the original haplotypes $H_0, \ldots, H_{k-1}$.

Clustering

The first step of our algorithm is to cluster reads that are likely to originate from the same haplotype. The clustering is based on pairwise similarity of overlapping reads. The similarity scores of the read pairs are then used in the clustering process. Two reads with an overlap of less than 2 variants are not considered as overlapping and always get a neutral score of 0.

We make two assumptions about the reads for the scoring scheme. First, all true haplotypes are expected to be equally frequent among the reads. Second, the Hamming rate between all pairs of haplotypes is expected to be the same (i.e., all haplotypes are equally different from each other). While the first assumption is reasonable, the second one is a simplification, as in practice the dissimilarity between even a fixed pair of haplotypes can vary heavily depending on evolutionary history and chromosomal region. The idea is to estimate the expected Hamming rate between reads from the same haplotype, which we call d_{same}, and the expected Hamming rate for reads from different haplotypes, called d_{diff}. The former depends only on the sequencing error rate, while the latter additionally includes the differences between the true haplotypes. With d_{all}, we further denote the expected Hamming rate over all overlapping read pairs.

For two reads r and s, the probability of observing the same allele at a shared variant locus equals d_{same} if $H(r) = H(s)$ or d_{diff} if $H(r) \neq H(s)$. Since the variants are independent from each other in our model, $\text{dis}rs$ should follow one of the two binomial distributions $B_{\text{olp}rs, d_{\text{same}}}$ or $B_{\text{olp}rs, d_{\text{diff}}}$ with $\text{olp}rs$ being the number of attempts and d_{same} or d_{diff} being the success probability. For each individual read pair, we can then decide which of the two possible distributions is the most likely one.

According to our first assumption, a $\frac{1}{k}$-fraction of all possible read pairs include reads from the same haplotype each, as for a read r there is a $\frac{1}{k}$ chance that another one is from the same haplotype. In order to estimate d_{same}, we compute the Hamming rate over all overlapping read pairs and use the average of the lower $\frac{1}{k}$-fraction as estimate. As alternative, one could also use the sequencing error rate to compute d_{same}, since the corresponding reads contain only sequencing errors. These error rates are, however, not always available, especially when preprocessing steps like error correction are included. Since d_{all} can be simply computed and $d_{\text{all}} \approx \frac{1}{k}d_{\text{same}} + \frac{k-1}{k}d_{\text{diff}}$, we get an estimate on d_{diff} as well. Finally, the similarity score of reads r and s is defined as

$$\log \left(\frac{f(\text{dis}rs, \text{olp}rs, d_{\text{same}})}{f(\text{dis}rs, \text{olp}rs, d_{\text{diff}})} \right)$$

with f being the binomial probability density function for $\text{dis}rs$ successes.
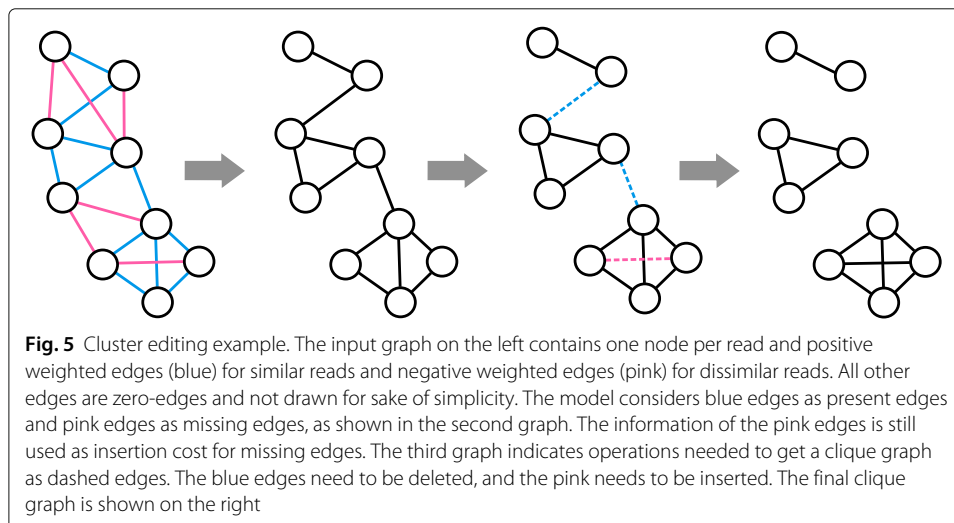
Please note that the score is negative if the read pair appears to be from different haplotypes and positive in the opposite case.

In our studies, we noticed that the disagreement rate between haplotypes varied between different regions. In order to increase the accuracy of our model, we partition the variants into windows $w_1, \ldots, w_l$ of average read length and compute d_{same} and d_{diff} independently for each window. If the overlap region of a read pair spans multiple windows, we use the weighted average of the d-values.

As clustering model, we chose cluster editing [19], which takes a complete graph with real edge weights as input and finds the most cost-efficient way to transform it into a graph only consisting of disjoint cliques. Therefore, positive weighted edges are interpreted as present edges and negative ones as missing edges. The absolute value of a weight is the cost to either insert a missing or delete a present edge. A small example of this model can be found in Fig. 5. For our algorithm, we model each read as a node of the input graph and use the similarity score for each read pair to obtain edge weights. Non-overlapping read pairs are defined to have an edge weight of 0, which we call a *zero-edge*. The resulting cliques can be interpreted as clusters of reads with high confidence of originating from the same haplotype.

The number of clusters depends on the data and is not an input parameter. In practice, we get much more than k clusters for two reasons: First, the distance between variants can vary and can become too large for some variant pairs such that enough reads connect them with sufficient confidence. Second, collapsed regions lead to clusters with reads from multiple haplotypes, forcing single-haplotype clusters to be discontinued. Restricting the cluster editing model to k clusters would force clusters to span poorly connected variants and split up reads from locally identical sequences. As this would likely introduce errors, we instead postpone the problem of reducing the clusters to k haplotypes to the second part of our algorithm.

Due to the NP hardness of the cluster editing problem, it is infeasible to solve it to optimality on large real-world instances as given by the comparison of all read pairs. Instead, we use a heuristic that greedily picks an edge in each iteration and decides whether it should be present in the resulting clique graph or not, potentially inserting or deleting

Fig. 5 Cluster editing example. The input graph on the left contains one node per read and positive weighted edges (blue) for similar reads and negative weighted edges (pink) for dissimilar reads. All other edges are zero-edges and not drawn for sake of simplicity. The model considers blue edges as present edges and pink edges as missing edges, as shown in the second graph. The information of the pink edges is still used as insertion cost for missing edges. The third graph indicates operations needed to get a clique graph as dashed edges. The blue edges need to be deleted, and the pink needs to be inserted. The final clique graph is shown on the right

edges. We denote the first case as making an edge *permanent* and the second one as making an edge *forbidden*. If an edge (u, v) is made permanent, for all other nodes w, it must hold that either both (u, w) and (v, w) must be in the final clique graph or none of them. Similarly, if (u, v) is made forbidden, there must not be any node w such that both (u, w) and (v, w) are in the final clique graph. Following these conditions, we can compute *induced costs* for each edge (u, v), which reflect the costs of obligatory insertion and deletion operations for making (u, v) permanent or forbidden. These costs are called $\mathrm{icp}(u, v)$ and $\mathrm{icf}(u, v)$ respectively and were originally defined in [31]. Once an edge becomes permanent (forbidden), its weight is set to ∞ ($-\infty$) and all induced costs of incident edges are updated accordingly.

To improve the running time, we ignore zero-edges in the heuristic and assume them to not be present in the solution, unless one of them is needed to complete a clique.

Haplotype threading

For the second part of the algorithm, we developed a novel approach called *haplotype threading*, which performs the actual phasing to k haplotypes. The cluster editing step results in a set C of read clusters with two properties: First, the number of clusters at a position $i \in \{0, \ldots, n-1\}$ can be larger than k, so that some clusters do not contribute to any computed haplotype. Second, the reads in a cluster $c \in C$ usually do not cover the whole chromosome, but only a part of the n variants, so in order to obtain whole-chromosome haplotypes, these must be assembled from multiple clusters. This is done by *threading* a haplotype through the clusters, meaning that for every haplotype, a path through C is assembled by choosing one cluster $c \in C$ for each haplotype at every variant position i.

In a genome of ploidy k, we seek for k haplotypes and thus assemble all k sequences simultaneously by choosing k-tuples of clusters at each position. Duplicate clusters within tuples are allowed since reads from one cluster can belong to multiple true haplotypes: For regions with high local similarity between the true haplotypes, the corresponding reads are likely placed into one cluster by the cluster editing step.

In the threading process, we aim at achieving three objectives: (i) genotype concordance, (ii) read coverage, and (iii) haplotype contiguity. The first, *genotype concordance*, captures the agreement between the known target genotype and the chosen clusters. For the true haplotypes $H_0, \ldots, H_{k-1}$ of length n, the corresponding genotype can be described as the component-wise sum $G = H_0 + H_1 + \cdots + H_{k-1}$ and is denoted by $G = g_0, g_1, \ldots, g_{n-1}$. Furthermore, for each cluster c and each position i, we can compute the consensus $\mathrm{cons}(c, i) \in \{0, 1, -\}$ as the most frequent allele among all reads in c at position i. Using this definition, we can compute a consensus genotype of a k-tuple $(c_0, \ldots, c_{k-1})$ at position i as $\sum_{j=0}^{k-1} \mathrm{cons}(c_j, i)$. For each position i, we then only take those cluster tuples into account whose consensus genotype at i is *concordant* with the target genotype, i.e., $\sum_{j=0}^{k-1} \mathrm{cons}(c_j, i) = g_i$. Reducing the search space of possible tuples this way increases efficiency as well as accuracy by filtering out non-promising combinations beforehand. In case there is no tuple with a concordant genotype at position i, we allow genotype deviations of 1; if this also fails, all possible tuples are considered.

To determine the best fit among the possible cluster tuples, we designed an objective function that takes the remaining two criteria into account as follows. The second criterion is *read coverage*. Since in locally identical regions multiple haplotypes can be

threaded through the same cluster—which leads to multiple appearances of this cluster in the k-tuple—this number of haplotypes has to correspond to the coverage of the chosen cluster. The relative coverage of a cluster c at position i describes the number of reads in c covering i divided by the total number of reads in all clusters that cover i. We denote this value by $\mathrm{cov}(c, i)$. Then, we can compute the expected copy number of a cluster c at i, i.e., the expected number of haplotypes that are threaded through c, by $cn_{\exp}(c, i) = \lceil k\mathrm{cov}(c, i) - \frac{1}{2k} \rceil$. The true copy number of c in a chosen cluster tuple $(c_0, \dots, c_{k-1})$ is given by $cn_{\mathrm{true}}((c_0, \dots, c_{k-1}), c, i) = |\{i \mid i \in \{0, \dots, k - 1\}, c = c_i\}|$. Deviations of the true number of occurrences from the expected ones are penalized by a constant factor p_{cov} per cluster, so that a cluster tuple $(c_0, \dots, c_{k-1})$ is evaluated by the cost function

$$\mathrm{costs}_{\mathrm{cov}}((c_0, \dots, c_{k-1}), i) = \sum_{j=0}^{k-1} p_{\mathrm{cov}}[[\, cn_{\exp}(c_j, i) \neq cn_{\mathrm{true}}((c_0, \dots, c_{k-1}), c_j, i)]]$$

where $[[\, x \neq y]]$ returns 1 if $x \neq y$ and 0 otherwise.

The third and last criterion, *haplotype contiguity*, encourages haplotypes to stay in the same cluster as long as possible, so that switching of haplotypes between clusters is penalized. For two consecutive cluster tuples $(c_0, \dots, c_{k-1})$ and $(c'_0, \dots, c'_{k-1})$ at positions i and $i + 1$, we denote the cost factor by p_{switch}, which results in the cost function
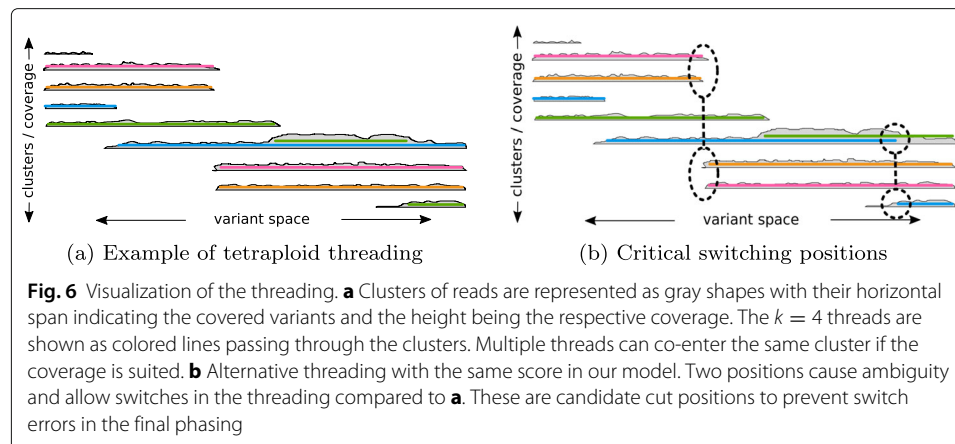
$$\mathrm{costs}_{\mathrm{switch}}\left((c_0, \dots, c_{k-1}), (c'_0, \dots, c'_{k-1})\right) = \sum_{i=0}^{k-1} p_{\mathrm{switch}}\left[[c_i \neq c'_i]\right]$$

We developed a dynamic programming approach to rapidly find the optimal sequence of tuples that minimizes all costs. We compute a two-dimensional matrix S with a column for every variant j from 0 to $n - 1$ and a row for every possible genotype-conform tuple of clusters. Since the number of eligible cluster tuples can differ between variant positions, the columns of S do not necessarily have the same lengths. We denote the length of a column j with l_j. Using the cost functions defined above, $S[\,i, j]$ is then computed as

$$S[\,i, 0] = \mathrm{costs}_{\mathrm{cov}}(c_i, 0) \ ,$$
$$S[\,i, j] = \mathrm{costs}_{\mathrm{cov}}(c_i, j) + $$
$$\min_{k \in \{0, \dots, l_{j-1}-1\}} \left(S[\,k, j - 1] + \mathrm{costs}_{\mathrm{switch}}(c_k, c_i)\right) \qquad \text{for } j > 0 \ ,$$

where c_i denotes the cluster tuple in row i. The optimal threading score is then given by the minimum value in the last column. Starting at this position, we assemble the sequence of clusters with minimum costs via backtracing.

The threading process is illustrated in Fig. 6a for $k = 4$. The clusters from the first step are drawn as gray shapes in a two-dimensional space, where the horizontal position refers to the variants covered by the reads inside a cluster and the height represents the relative coverage of a cluster at every position. The position on the y-axis has no numerical meaning and is just used for illustration purpose. Starting from the left, a 4-tuple of the five present clusters needs to be chosen. According to the coverage, the best choice is to thread one haplotype through each of the four clusters with highest coverage and ignoring the smallest one, as this is likely to contain noisy reads only. From thereon, the threads change clusters whenever a cluster ends or undergoes a drastic change in relative coverage.

(a) Example of tetraploid threading   (b) Critical switching positions

Fig. 6 Visualization of the threading. **a** Clusters of reads are represented as gray shapes with their horizontal span indicating the covered variants and the height being the respective coverage. The $k = 4$ threads are shown as colored lines passing through the clusters. Multiple threads can co-enter the same cluster if the coverage is suited. **b** Alternative threading with the same score in our model. Two positions cause ambiguity and allow switches in the threading compared to **a**. These are candidate cut positions to prevent switch errors in the final phasing

Block cuts

Phasing tools are able to divide phased haplotypes into blocks if there is not enough evidence in the data to connect these blocks. This is usually done when there are two variants with no connecting read in between. For polyploid organisms, however, even a single connecting read is not sufficient, as reads from $k - 1$ different haplotypes are needed to resolve the connection of k haplotypes on both sides. In general, block cuts are a trade-off between block length and accuracy, as one of these metrics can easily be optimized by giving up the other one. To offer more flexibility to the user, WHATSHAP POLYPHASE provides different modi of applying block cuts to either get short and accurate or long but less accurate blocks.

Since it is uncertain whether different read clusters represent the same or different haplotypes, the most conservative method is to cut the phasing whenever one thread switches to another cluster, which we call *single-switch-cuts*. While this yields the lowest block-wise error rate, many of these cuts can be avoided. If only one thread switches the cluster, while the other $k - 1$ threads stay, one could conclude that the old cluster is linked to the new one by process of elimination. If two or more threads switch, the continuation is ambiguous and a cut can be placed here to prevent switch errors, which we call a *multi-switch-cut*. In principle, only the switching threads need to be cut, while the rest can stay connected. To the best of our knowledge, however, there is no established method to express such selective block cuts in a VCF file. Therefore, all haplotypes are interrupted in case of a multi-switch-cut.

The last type of cuts is the *separation-cut*, which is necessary to handle collapsed regions. Assume a cluster contains multiple threads at some position and the number of threads has to be decreased by 1 for the next position due to a decrease in coverage. Even though this is not covered by the multi-switch-cuts, there is still a choice which of the contained threads should leave the cluster. If all threads have been in the cluster since the start of the current block, the leaving thread can be chosen arbitrarily. However, if they entered the current cluster on different positions or from different predecessor clusters, the choice affects the resulting haplotype sequences and we need to insert a separation-cut here to avoid potential switch errors. Figure 6b shows an example, where two threads (green and blue) share the same cluster before one of them has to leave. Either of them switching would lead to a different result, for which we do not know the correct one.

Preprocessing and phasing the potato genome

We ran a recently developed error correction pipeline [25] to reduce the typically high number of sequencing errors in the Oxford Nanopore MinION reads, in order to use them for variant calling and phasing. Illumina reads were first self-corrected using Lighter [32], the corrected reads were used to build a de Bruijn graph with bcalm2 [33], and the MinION reads were aligned to the graph with GraphAligner [25]. We used the default parameters for the error correction pipeline (Lighter $k = 21$, bcalm2 $k = 61$ and abundance $= 3$, GraphAligner default alignment parameters). The corrected read sequence of each mapped MinION read was obtained from the path of its respective alignment in the graph. The corrected reads where aligned to the reference genome using minimap2 [26] and converted to BAM-format using samtools [34]. In the next step, we ran FreeBayes [27] (with additional parameters: `-p 4 -no-indels -no-mnsp -no-complex`) inside of all gene regions to call SNPs from the corrected Nanopore alignments. As base qualities are not produced during error correction and FreeBayes seems to need them in order to compute genotypes, we added a constant quality of 40 for all bases to the BAM file before calling SNPs. Finally, we ran WHATSHAP POLYPHASE in order to phase the variants with option `-verify-genotypes`. This option invokes an additional step prior to phasing, which re-genotypes all variants and only keeps those positions for which the computed genotype matches the input genotype. For determining the genotype of a position, we implemented a simple algorithm that calculates the fraction of reference and alternative alleles that cover a variant and compare it to the fractions that we would expect for all possible genotypes. We then assign the genotype whose expected fractions of reference and alternative alleles best match the ones observed in the data.

We focused on the FRIGIDA-like protein 5 isoform X2 gene to demonstrate a use case of polyploid phasing. We first extracted all phased variants that are part of the longest phasing block reported by WHATSHAP POLYPHASE for this gene. In order to assign reads to the haplotypes computed by WHATSHAP POLYPHASE, we extended the command `whatshap haplotag`, which was previously implemented for the diploid version of `whatshap`, to the polyploid case. Given a phased VCF with predicted haplotypes and BAM file with sequencing reads, we assign each read to the haplotype it is most similar to in terms of the alleles observed at variant positions in the read. This assignment is stored by tagging the respective sequences in the BAM file, which enables visualizing the haplotype clusters by programs like IGV [23] (see Fig. 4c). Furthermore, we extended the subcommand `whatshap split` to higher ploidies, which can be used to split tagged reads by haplotype and store them in separate files. For each haplotype, we produced a BAM file with reads in this way.

In the next step, we ran wtdbg2 [28] (with options `-x ccs -g 1m`) separately for reads corresponding to each haplotype to generate haplotype-resolved assemblies for the Frigida gene. Those were further analyzed with NCBI's ORFfinder and COBALT algorithms [29, 30] using their web interfaces (https://www.ncbi.nlm.nih.gov/orffinder/, https://www.ncbi.nlm.nih.gov/tools/cobalt/re_cobalt.cgi).

Supplementary information
Supplementary information accompanies this paper at https://doi.org/10.1186/s13059-020-02158-1.

Additional file 1: We provide the switch errors and Hamming rates of WHATSHAP POLYPHASE and H-PoPG, HapCompass, HapTree and Ranbow on an exemplary region of chromosome 1. Further, we provide figures showing

the N50 block lengths and the corresponding Hamming rates for different block cut strategies of WHATSHAP POLYPHASE on the artificial tetraploid human dataset, including real and simulated reads on different coverages. We also include a comparison between WHATSHAP POLYPHASE and H-PoPG in collapsing regions, using long blocks similar to H-PoPG. For the potato data, we show the relation between the fraction of phased variants and heterozygosity level in the potato genes and also include a figure showing the haplotype assemblies for the FRIGIDA-like protein 5 isoform X2 gene. Additionally, we show an example for the phasing behavior of WHATSHAP POLYPHASE and H-PoPG over weakly connected variants. Last, we provide a description of the haploid N50 block length and the comparison between the regular and the haploid N50 for several simulated datasets.

Additional file 2: Review history.

Authors' contributions
SDS, RSM, JE, GWK, and TM developed the algorithmic concepts and designed the study. RSM designed the haplotype threading algorithm and implemented a prototype. SDS designed and implemented the cluster editing algorithm, designed the block cut strategies, and optimized the threading implementation. JE performed the evaluation and analyzed the potato dataset. MR ran the error correction on the potato reads. LS, JJR, and BU performed potato sequencing, and BU helped with the interpretation of phasing results. SDS, RSM, and JE integrated all software components into WhatsHap and tested the workflow. SDS, RSM, JE, GWK, and TM wrote the paper. All authors read and approved the final manuscript.

Authors' information
Twitter handles: @guwekl (Gunnar W. Klau), @tobiasmarschal (Tobias Marschall), @rserramari (Rebecca Serra Mari), @usadellab (Björn Usadel).

Availability of data and materials
WHATSHAP POLYPHASE is available as open source code under the MIT License [35]. The version to generate the results in this paper is archived at [36]. All scripts needed to handle the input data and to reproduce the analyses can be found at https://github.com/eblerjana/whatshap-polyphase-experiments. The input files for the data generation are alignments of the diploid samples and gold standard phasings of these samples, which can be found at ftp://ftp.1000genomes.ebi.ac.uk/vol1/ftp/data_collections/hgsv_sv_discovery/working/20180102_pacbio_blasr_reheader/ and on https://zenodo.org/record/3999218, respectively. For read simulation, the pipeline uses the file SRR3658380.fastq as example file, which is available at https://www.ncbi.nlm.nih.gov/sra/SRX1837675. *Solanum tuberosum* sequencing data has been made available at the NCBI with accession number PRJNA587397 [37].

Ethics approval and consent to participate
Not applicable.

Consent for publication
Not applicable.

Competing interests
The authors declare that they have no competing interests.

Author details
[1] Algorithmic Bioinformatics, Heinrich Heine University Düsseldorf, Universitätsstr. 1, 40225 Düsseldorf, Germany. [2] Institute for Medical Biometry and Bioinformatics, Medical Faculty, Heinrich Heine University Düsseldorf, Moorenstraße 5, 40225 Düsseldorf, Germany. [3] Center for Bioinformatics, Saarland University, Saarland Informatics Campus E2.1, 66123 Saarbrücken, Germany. [4] Graduate School of Computer Science, Saarland Informatics Campus E1.3, 66123 Saarbrücken, Germany. [5] Max Planck Institute for Informatics, 66123 Saarbrücken, Germany. [6] Forschungszentrum Jülich IBG-4, Wilhelm-Johnen-Str., 52428 Jülich, Germany. [7] Institute for Biology I, RWTH Aachen, Worringer Weg 3, 52074 Aachen, Germany. [8] Cluster of Excellence on Plant Sciences (CEPLAS), Heinrich Heine University Düsseldorf, Universitätsstr. 1, 40225 Düsseldorf, Germany.

References

1. Yang J, Moeinzadeh M-H, Kuhl H, Helmuth J, Xiao P, Haas S, Liu G, Zheng J, Sun Z, Fan W, Deng G, Wang H, Hu F, Zhao S, Fernie AR, Boerno S, Timmermann B, Zhang P, Vingron M. Haplotype-resolved sweet potato genome traces back its hexaploidization history. Nat Plants. 2017;3(9):696–703. https://www.nature.com/articles/s41477-017-0002-z.
2. Visser RGF, Bachem CWB, Borm T, de Boer J, van Eck HJ, Finkers R, van der Linden G, Maliepaard CA, J G A M, Voorrips R, Vos P, Wolters AMA. Possibilities and challenges of the potato genome sequence. Potato Res. 2014;57(3-4):327–30.
3. Li K-T, Moulin M, Mangel N, Albersen M, Verhoeven-Duif NM, Ma Q, Zhang P, Fitzpatrick TB, Gruissem W, Vanderschuren H. Increased bioavailable vitamin B6 in field-grown transgenic cassava for dietary sufficiency. Nat Biotechnol. 2015;33:1029–32.
4. Klau GW, Marschall T. A guided tour to computational haplotyping. In: Unveiling dynamics and complexity. Lecture Notes in Computer Science. Cham: Springer; 2017. p. 50–63.
5. Lippert R, Schwartz R, Lancia G, Istrail S. Algorithmic strategies for the single nucleotide polymorphism haplotype assembly problem. Brief Bioinform. 2002;3(1):23–31.
6. Patterson M, Marschall T, Pisanti N, van Iersel L, Stougie L, Klau GW, Schönhuth A. WhatsHap: weighted haplotype assembly for future-generation sequencing reads. J Comput Biol. 2015;22(6):498–509.
7. Aguiar D, Istrail S. Haplotype assembly in polyploid genomes and identical by descent shared tracts. Bioinformatics. 2013;29(13):352–60.
8. Aguiar D, Istrail S. HapCompass: a fast cycle basis algorithm for accurate haplotype assembly of sequence data. J Comput Biol. 2012;19(6):577–90.
9. Berger E, Yorukoglu D, Peng J, Berger B. HapTree: a novel Bayesian framework for single individual polyplotyping using NGS data. PLoS Comput Biol. 2014;10(3):1003502.
10. Motazedi E, Finkers R, Maliepaard C, de Ridder D. Exploiting next-generation sequencing to solve the haplotyping puzzle in polyploids: a simulation study. Brief Bioinform. 2017;19(3):387–403. https://academic.oup.com/bib/article/19/3/387/2870504.
11. Das S, Vikalo H. SDhaP: haplotype assembly for diploids and polyploids via semi-definite programming. BMC Genomics. 2015;16:260.
12. Xie M, Wu Q, Wang J, Jiang T. H-PoP and H-PoPG: heuristic partitioning algorithms for single individual haplotyping of polyploids. Bioinformatics. 2016;32(24):3735–44.
13. He D, Saha S, Finkers R, Parida L. Efficient algorithms for polyploid haplotype phasing. BMC Genomics. 2018;19(Suppl 2):110.
14. Motazedi E, de Ridder D, Finkers R, Baldwin S, Thomson S, Monaghan K, Maliepaard C. Tripoly: haplotype estimation for polyploids using sequencing data of related individuals. Bioinformatics. 2018;34(22):3864–72. https://doi.org/10.1093/bioinformatics/bty442.
15. Chaisson MJP, Mukherjee S, Kannan S, Eichler EE. Resolving multicopy duplications de novo using polyploid phasing. Res Comput Mol Biol. 2017;10229:117–33.
16. Cai C, Sanghavi S, Vikalo H. Structured Low-Rank matrix factorization for haplotype assembly. IEEE J Sel Top Signal Process. 2016;10(4):647–57.
17. Hashemi A, Zhu B, Vikalo H. Sparse tensor decomposition for haplotype assembly of diploids and polyploids. BMC Genomics. 2018;19(Suppl 4):191.
18. Siragusa E, Haiminen N, Finkers R, Visser R, Parida L. Haplotype assembly of autotetraploid potato using integer linear programing. Bioinformatics. 2019;35(18):3279–86. https://doi.org/10.1093/bioinformatics/btz060.
19. Zahn CTJ. Approximating symmetric relations by equivalence relations. J Soc Ind Appl Math. 1964;12. https://doi.org/10.1137/0112071.
20. Töpfer A, Marschall T, Bull RA, Luciani F, Schönhuth A, Beerenwinkel N. Viral quasispecies assembly via maximal clique enumeration. PLoS Comput Biol. 2014;10(3):1–10. https://doi.org/10.1371/journal.pcbi.1003515.
21. Chaisson MJP, Sanders AD, Zhao X, Malhotra A, Porubsky D, Rausch T, Gardner EJ, Rodriguez O, Guo L, Collins RL, Fan X, Wen J, Handsaker RE, Fairley S, Kronenberg ZN, Kong X, Hormozdiari F, Lee D, Wenger AM, Hastie A, Antaki D, Audano P, Brand H, Cantsilieris S, Cao H, Cerveira E, Chen C, Chen X, Chin C-S, Chong Z, Chuang NT, Lambert CC, Church DM, Clarke L, Farrell A, Flores J, Galeev T, Gorkin D, Gujral M, Guryev V, Heaton WH, Korlach J, Kumar S, Kwon JY, Lee JE, Lee J, Lee W-P, Lee SP, Li S, Marks P, Viaud-Martinez K, Meiers S, Munson KM, Navarro F, Nelson BJ, Nodzak C, Noor A, Kyriazopoulou-Panagiotopoulou S, Pang A, Qiu Y, Rosanio G, Ryan M, Stütz A, Spierings DCJ, Ward A, Welch AE, Xiao M, Xu W, Zhang C, Zhu Q, Zheng-Bradley X, Lowy E, Yakneen S, McCarroll S, Jun G, Ding L, Koh CL, Ren B, Flicek P, Chen K, Gerstein MB, Kwok P-Y, Lansdorp PM, Marth G, Sebat J, Shi X, Bashir A, Ye K, Devine SE, Talkowski M, Mills RE, Marschall T, Korbel JO, Eichler EE, Lee C. Multi-platform discovery of haplotype-resolved structural variation in human genomes. Nat Commun. 2019;10(1):1784. https://doi.org/10.1038/s41467-018-08148-z.
22. Ono Y, Asai K, Hamada M. PBSIM: PacBio reads simulator—toward accurate genome assembly. Bioinformatics. 2012;29(1):119–21. https://doi.org/10.1093/bioinformatics/bts649.
23. Robinson JT, Thorvaldsdóttir H, Winckler W, Guttman M, Lander ES, Getz G, Mesirov JP. Integrative genomics viewer. Nat Biotechnol. 2011;29(1):24.
24. Hardigan MA, Crisovan E, Hamilton JP, Kim J, Laimbeer P, Leisner CP, Manrique-Carpintero NC, Newton L, Pham GM, Vaillancourt B, Yang X, Zeng Z, Douches DS, Jiang J, Veilleux RE, Buell CR. Genome reduction uncovers a large dispensable genome and adaptive role for copy number variation in asexually propagated Solanum tuberosum. Plant Cell. 2016;28(2):388–405. https://doi.org/10.1105/tpc.15.00538.
25. Rautiainen M, Marschall T. Graphaligner: rapid and versatile sequence-to-graph alignment. BioRxiv. 2019810812. https://doi.org/10.1101/810812.
26. Li H. Minimap2: pairwise alignment for nucleotide sequences. Bioinformatics. 2018;34(18):3094–100. https://doi.org/10.1093/bioinformatics/bty191.

27. Garrison E, Marth G. Haplotype-based variant detection from short-read sequencing. arXiv preprint arXiv:1207.3907. 2012.
28. Ruan J, Li H. Fast and accurate long-read assembly with wtdbg2. Nat Methods. 2020;17:155–58. https://www.nature.com/articles/s41592-019-0669-3.
29. Wheeler DL, Church DM, Federhen S, Lash AE, Madden TL, Pontius JU, Schuler GD, Schriml LM, Sequeira E, Tatusova TA, et al. Database resources of the national center for biotechnology. Nucleic Acids Res. 2003;31(1):28–33.
30. Papadopoulos JS, Agarwala R. Cobalt: constraint-based alignment tool for multiple protein sequences. Bioinformatics. 2007;23(9):1073–9.
31. Böcker S, Briesemeister S, Klau GW. Exact algorithms for cluster editing: evaluation and experiments. Algorithmica. 2011;60(2):316–34. https://doi.org/10.1007/s00453-009-9339-7.
32. Song L, Florea L, Langmead B. Lighter: fast and memory-efficient sequencing error correction without counting. Genome Biol. 2014;15(11):509.
33. Chikhi R, Limasset A, Medvedev P. Compacting de Bruijn graphs from sequencing data quickly and in low memory. Bioinformatics. 2016;32(12):201–8. https://doi.org/10.1093/bioinformatics/btw279.
34. Li H, Handsaker B, Wysoker A, Fennell T, Ruan J, Homer N, Marth G, Abecasis G, Durbin R. The sequence alignment/map format and samtools. Bioinformatics. 2009;25(16):2078–9.
35. Schrinner S, Serra Mari R, Ebler J, Marschall T, Klau GW. WHATSHAP POLYPHASE source code. 2020. https://github.com/whatshap/whatshap. Accessed 25 Aug 2020.
36. Schrinner S, Serra Mari R, Ebler J, Marschall T, Klau GW. Version of WHATSHAP POLYPHASE used to produce the results in this manuscript. 2020. https://zenodo.org/record/3999208. Accessed 25 Aug 2020.
37. Seillier L, Usadel B, Reimer J. Solanum tuberosum genome sequencing. Oxford Nanopore and Illumina Data. NCBI Short Read archive. 2019. https://www.ncbi.nlm.nih.gov/bioproject/PRJNA587397. Accessed 04 Nov 2019.

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.