SYSTEMS AND METHODS FOR CORRELATED ERROR EVENT MITIGATION FOR VARIANT CALLING

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application claims the benefit of U.S. Provisional Application No. 62/710,348, filed on February 16, 2018, and titled "Methods, Devices, and Systems for Performing Foreign Read Detection and Burst Error Detection," the disclosure of which is incorporated herein by reference in its entirety.

SEQUENCE LISTING

This application contains a Sequence Listing that has been submitted electronically as an ASCII text file named Sequence_Listing.txt. The ASCII text file, created on March 10, 2022, is 877 bytes in size. The material in the ASCII text file is hereby incorporated by reference in its entirety.

BACKGROUND

[0002] A nucleic acid sequencer is an instrument that is configured to automate the process of sequencing nucleic acids such as deoxyribonucleic acid (DNA) or ribonucleic acid (RNA). Nucleic acid sequencing is a process of determining an order of nucleotides in a genetic sequence.

[0003] The nucleic acid sequencer is configured to receive a nucleic acid sample and generate output data, referred to as one or more “reads,” that represents an order of nucleotides in the nucleic acid sample. The nucleotides in a DNA sample can include one or more of guanine (G), cytosine (C), adenine (A), and thymine (T) in any combination. The nucleotides in an RNA sample can include one or more of G, C, A, and uracil (U) in any combination in any combination.

[0004] The DNA reads generated by the DNA sequencer can be mapped and aligned to a known DNA sequence of a reference genome. Once mapped and aligned to a reference genome, the mapped and aligned sequence of reads can be analyzed in view of the reference genome to identify potential variations that exist between the mapped and aligned sequence of reads and the reference genome.

SUMMARY

[0005] Aspects of the present disclosure are directed towards methods, systems, and apparatus, including computer programs, for accounting for correlated error events that are problematic for variant callers used to determine whether an alternative (hereinafter “alt”) allele identified in a pileup of mapped and aligned sequence reads is a true variant.

[0032] FIG. 9A is an example of summary of experimental results from execution of a process for correlated error mitigation for variant calling that has been performed on a pileup of sequencing reads whose result indicates an example of a true positive results (SEQ ID NO: 1).

[0033] FIG. 9B is an example of a modified set of probability results for the pileup of FIG. 9A.

[0034] FIG. 10A is an example of summary of experimental results from execution of a process for correlated error mitigation for variant calling that has been performed on a pileup of sequencing reads whose result indicates a low likelihood of an occurrence of a mapping error (SEQ ID NO: 2).

[0035] FIG. 10B is an example of a modified set of probability results for the pileup of FIG. 10A.

[0036] FIG. 11A is an example of summary of experimental results from execution of a process for correlated error mitigation for variant calling that has been performed on a pileup of sequencing reads whose result indicates a high likelihood that the candidate alt allele is unlikely to be a true variant due to a sequencing error (SEQ ID NO: 3).

[0037] FIG. 11B is an example of a modified set of probability results for the pileup of FIG. 11A.

[0038] FIG. 12A is an example of summary of experimental results from execution of a process for correlated error mitigation for variant calling that has been performed on a pileup of sequencing reads whose result indicates a high likelihood that the candidate alt allele is unlikely to be a true variant due to a sequencing error on both read orientations (SEQ ID NO: 4).

[0039] FIG. 12B is an example of a modified set of probability results for the pileup of FIG. 12.

DETAILED DESCRIPTION

account for this potential error, indicators in the probability model for determining a likelihood that an occurrence of one or more mapping errors exits, let θ_i indicate the strand direction for read *i*, with values 0 and 1 indicating the forward and reverse strand directions, respectively. With a strand-aware modified probability mapping model for determining a likelihood that an occurrence of one or more mapping errors exist, three hypotheses can be evaluated. These three hypothesis include that the foreign reads occur exclusively on the forward strand, exclusively on the reverse strand, or on both strands. A solution includes a hypothesis that maximizes Ψ_m .

Starting with equation ~~**Error! Reference source not found**~~ (5), this variation is as follows:

$$\Psi_m = P(\mathbf{G}_m) \max_{n,\beta,p_F} p_F \left(\prod_{i \notin \Lambda_n} P(r_i \mid G_m) \right) \prod_{i \in \Lambda_n} \left(\beta P(r_i \mid F_m) U\left(-10\log_{10}\left(\frac{p_F}{P_0(F_m)} \right) - \mu_i \right) + (1-\beta) P(r_i \mid G_m) \right), \quad \text{Equation (8)}$$

where $\Lambda_0 = \{i : \theta_i = 0\}$, $\Lambda_1 = \{i : \theta_i = 1\}$, $\Lambda_2 = \{i : \theta_i = 0 \text{ or } 1\}$.

[0133] FIG. 5 is a flowchart of an example of a process for sequencing error mitigation for variant calling. The process 500 is described below as being performed by a computer such as the variant calling unit of FIG. 1. In some implementations, the variant calling unit may be implemented using hardware such as one or more FPGAs, ASICs, CPUs, GPUs, or any combination thereof, using software, or any combination thereof.

[0134] A computer may begin performance of the process 500 by accessing a pileup of aligned sequence reads stored in one or more memory devices (step 510). The aligned sequence reads may have been generated using a mapping and aligning unit implemented using configurable logic gates of an FGPA device. In some implementations, the accessed reads are stored in the one or more memory devices and can include a first set of sequence reads that are forward aligned and a second set of sequence reads that are reverse aligned. Each respective set of sequence read corresponds to a particular read orientation or direction.

[0135] The computer may continue performance of the process 500 by obtaining information describing (i) a read orientation for each of the reads of the pileup of aligned sequence reads stored in the one or more memories, (ii) a position of each base at the reference location with

[0155] Let us define the extended candidate genotype $\mathbf{G}'_m = \begin{bmatrix} G_{m,1} & G_{m,2} & E_{m,0} & E_{m,1} \end{bmatrix}$, where $E_{m,\theta}$ is the error allele for strand direction θ .

[0156] Let $L_{E,\theta}$ be the length of the homopolymer matching base, $E_{m,\theta}$ immediately preceding the error in strand direction θ.

[0157] For each extended candidate genotype **G'$_m$** , we calculate:

$$\Psi_m = P(\mathbf{G}_m) \prod_{\theta=0,1} \max_{n_\theta,\alpha_\theta} f\left(\overline{q}_{\theta,n_\theta}, L_{E,\theta}\right) \left(\prod_{i>n_\theta} \frac{P\left(r_{\theta,i} \mid G_{m,1}\right) + P\left(r_{\theta,i} \mid G_{m,2}\right)}{2} \right) \prod_{i \le n_\theta} \left(\alpha_\theta P\left(r_{\theta,i} \mid E_{m,\theta}\right) + (1-\alpha_\theta) \frac{P\left(r_{\theta,i} \mid G_{m,1}\right) + P\left(r_{\theta,i} \mid G_{m,2}\right)}{2} \right), \quad \text{Equation (9)}$$

where $f(\overline{q}, L_E)$ indicates the prior probability of an error event affecting a subset of reads as a function of the mean base quality over the subset and the length of the homopolymer matching the error.

[0158] The quantity Ψ_m represents an estimate of the joint probability $P(G'_m, R)$ under the following assumptions: (1) the error events affect a contiguous subset of reads when ordered by decreasing distance from the 5' end, starting with the first, and do not affect the reads outside of that subset, (2) the error events happen independently for each strand direction, and (3) the prior probability of such error events is a function of the mean base quality over the subset of reads affected by the error event and the length of the homopolymer matching base, $E_{m,\theta}$ immediately preceding the error in strand direction θ.

[0159] In some implementations, the number of candidates can be reduced to only test evaluating cases where $G_{m,1} = G_{m,2} = Gm$. In this case the expression simplifies to

$$\Psi_m = P(\mathbf{G}_m) \prod_{\theta=0,1} \max_{n_\theta,\alpha_\theta} f\left(\overline{q}_{\theta,n_\theta}, L_{E,\theta}\right) \left(\prod_{i>n_\theta} P\left(r_{\theta,i} \mid G_m\right) \right) \prod_{i \le n_\theta} \left(\alpha_\theta P\left(r_{\theta,i} \mid E_{m,\theta}\right) + (1-\alpha_\theta) P\left(r_{\theta,i} \mid G_m\right) \right). \quad \text{Equation (10)}$$

[0160] In some implementations, the value of α_θ can be fixed as $\alpha_\theta = 0.5$. Thus, we can rewrite ~~**Error! Reference source not found**~~ (7) as:

$$\Psi_m = P(\mathbf{G}_m)\prod_{\theta=0,1}\max_{n_\theta} f\left(\overline{q}_{\theta,n_\theta}, L_{E,\theta}\right)\left(\prod_{i>n_\theta} P\left(r_{\theta,i} \mid G_m\right)\right)\prod_{i\le n_\theta}\left(\frac{P\left(r_{\theta,i} \mid E_{m,\theta}\right)+P\left(r_{\theta,i} \mid G_m\right)}{2}\right). \quad \text{Equation (11)}$$

[0161] In some implementations, the prior probability function $f(\overline{q}_{\theta,n_\theta,L_{E\theta}})$ in **~~Error! Reference source not found~~** (6) is expressed as a general function that could have a wide range of shapes. In theory, this function could be trained on real data. However, given limitations on the use of training this function on real data, some implementations of the present probability model for determining sequencing errors may use $f\left(\overline{q}, L_E\right) = 10^{-\max\left(0, 2.5\overline{q} - \min\left(20, 5L_E\right)\right)/10}$.

[0162] In some implementations, the equations above may accommodate the possibility of different error alleles for each strand direction. However, in some implementations, the probability model for determining a likelihood of one or more sequencing errors may only consider hypotheses where $E_{m,0} = E_{m,1}$. In such implementations, the θ subscript can be dropped and then the error allele can be denoted as E_m.

[0163] FIG. 6 is another flowchart of an example of a process 600 for correlated error mitigation for variant calling. The process 600 is described below as being performed by a computer such as the variant calling unit of FIG. 1. In some implementations, the variant calling unit may be implemented using hardware such as one or more FPGAs, ASICs, CPUs, GPUs, or any combination thereof, using software, or any combination thereof.

[0164] A computer may begin performance of the process 600 by accessing a pileup of aligned sequence reads stored in one or more memory devices (step 610). The aligned sequence reads may have been generated using a mapping and aligning unit implemented using configurable logic gates of an FGPA device. In some implementations, the accessed reads are stored in the one or more memory devices and can include a first set of sequence reads that are forward aligned and a second set of sequence reads that are reverse aligned. Each respective set of sequence read corresponds to a particular read orientation or direction.

CLAIMS

What is claimed is:

[Claim 1]. A method for improving the accuracy of a variant call by accounting for indications of correlated error events, the method comprising:

accessing, by one or more computers and from one or more memory devices, a pileup of a plurality of sequence reads aligned to a first region of a reference genome;

obtaining, by the one or more computers, information describing ~~one or more~~ characteristics of each of the plurality of reads of the pileup corresponding to a first position of the reference genome, wherein the information describing the characteristics of the respective reads include information describing:

(i) a mapping quality score for each sequence read in the pileup at the first position,

(ii) a read-allele score for each sequence read in the pileup at the first position for each candidate allele at the first position,

(iii) a read orientation for each sequence read of the pileup at the first position,

(iv) a position of each base at the first position within each sequence read of the pileup at the first position with reference to the 5′ end of the sequence read,

(v) a read-allele score for each sequence read of the plurality of sequence reads for each candidate allele at the reference position, and

(vi) a base quality score for each read for the base at the first position;

providing, by the one or more computers and based on the obtained information, one or more inputs to a probability model describing the ~~one or more~~ characteristics of the plurality of reads of the pileup, wherein the probability model is configured to determine a score, for each hypothesis of one or more hypotheses selected based on the one or more inputs, that indicates whether the hypothesis is true;

obtaining, by the one or more computers, output information for each of the one or more hypotheses, wherein the output information for each of the one or more hypotheses is (i) generated by the probability model based on the probability model's processing of the one or more inputs to the probability model describing the one or more characteristics of the respective reads of the pileup and (ii) indicative of a score that indicates whether the hypothesis is true, wherein the output information includes:

first output information for a first hypothesis of the one or more hypotheses that includes a likelihood that the sequence reads at the first position indicate the occurrence of a homozygous reference with a foreign allele that matches an alternative allele (alt);

second output information for a second hypothesis of the one or more hypotheses that includes a likelihood that the sequence reads at the first position indicate the occurrence of a homozygous alt with a foreign allele that matches a reference allele,

third output information for a third hypothesis of the one or more hypotheses that includes a likelihood that the sequence reads at the first position indicate the occurrence of a homozygous reference with a sequencing error that matches the alt allele, and

fourth output information for a fourth hypothesis of the one or more hypotheses that includes a likelihood that the sequence reads at the first position indicate the occurrence of a homozygous alt with a sequencing error that matches the reference allele; and

determining, by the one or more computers and based on the obtained output information generated by the probability model for each of the plurality of hypotheses, a likelihood that a true variant exists at the first position.

[Claim 2]. The method of claim 1, wherein determining, by the one or more computers and based on the obtained output information generated by the probability model for each of the plurality of hypotheses, a likelihood that a true variant exists at the first position comprises:

determining, by the one or more computers, a collective score that is based on the output information generated by the probability model for each of the plurality of hypothesis, wherein the collective score indicates a likelihood that the true variant exists;

determining, by the one or more computers, whether the score generated by the collective score satisfies a predetermined threshold; and

based on determining, by the one or more computers, that the collective score satisfies the predetermined threshold, adding information to a VCF file indicating that a true variant exists at the first position, optionally wherein the information indicating that a true variant exists at the first position includes information identifying (i) the first position, (ii) a candidate alt allele at the first position, and (iii) the collective score.

~~3. The method of claim 2, wherein the information indicating that a true variant exists at the first position includes information identifying (i) the first position, (ii) a candidate alt allele at the first position, (iii) the collective score.~~

~~4. The method of claim 1, wherein the information describing the one or more characteristics of the respective reads includes information describing (i) a mapping quality score for each sequence read in the pileup at the first position and (ii) a read-allele score for each sequence read in the pileup at the first position for each candidate allele at the first position.~~

[Claim 3~~5~~]. The method of claim 41, wherein the read-allele score for each sequence read of the pileup at the first position is based on an output produced by a pair Hidden Markov Model (P-HMM) model, for each of the sequence reads at the first position, which indicates a probability of observing a sequence read r_i given a particular candidate allele $G_{m,\phi}$.

~~6. The method of claim 4, wherein the output information includes:~~

~~first output information for a first hypothesis of the one or more hypotheses that includes a likelihood that the sequence reads at the first position indicate the occurrence of a homozygous reference with a foreign allele that matches the alt; and~~

~~second output information for a second hypothesis of the one or more hypotheses that includes a likelihood that the sequence reads at the first position indicate the occurrence of a homozygous alt with a foreign allele that matches a reference allele.~~

~~7. The method of claim 1, wherein the information describing the one or more characteristics of the respective sequence reads includes information describing (i) a read orientation for each sequence read of the pileup at the first position, (ii) a position of each base at the first position within each sequence read of the pileup at the first position with reference to the 5′ end of the sequence read, (iii) a read-allele score for each sequence read of the plurality of~~

sequence reads for each candidate allele at the reference position, and (iv) a base quality score for each read for the base at the first position.

8. The method of claim 7, wherein the read-allele score for each sequence read of the pileup at the first position is based on an output produced by a P-HMM model, for each of the sequence reads at the first position, that indicates a probability of observing a sequence read r_i given a particular candidate allele $G_{m,\phi}$.

9. The method of claim 1, wherein the information describing the one or more characteristics of the respective sequence reads includes information describing (i) a read orientation for each sequence read of the pileup at the first position, (ii) a position of each base at the first position within each sequence read of the pileup at the first position with reference to the 5′ end of the sequence read, and (iii) a read-allele score for each sequence read of the plurality of sequence reads for each candidate allele at the first position.

10. The method of claim 7, wherein the output information includes:

first output information for a first hypothesis of the one or more hypotheses that includes a likelihood that the sequence reads at the first position indicate the occurrence of a homozygous reference with a sequencing error that matches the alt allele, and

second output information for a second hypothesis of the one or more hypotheses that includes a likelihood that the sequence reads at the first position indicate the occurrence of a homozygous alt with a sequencing error that matches the reference allele.

11. The method of claim 1, wherein the information describing the one or more characteristics of the respective sequence reads includes information describing (i) a read orientation for each sequence read of the pileup at the first position, (ii) a position of each base at the first position such as position "0" 142 within each sequence read with reference to the 5′ end of the sequence read, (iii) a mapping quality score for each sequence read of the pileup at the first position, (iv) a read-allele score for each sequence read of the plurality of reads for each

~~candidate allele at the reference position, and (v) a base quality score for each read for the base aligned at the first position.~~

~~12. The method of claim 11, wherein the read-allele score for each sequence read of the pileup at the first position is based on an output produced by a P-HMM model, for each of the sequence reads at the first position, that indicates a probability of observing a read r_i given a particular candidate allele $G_{m,\phi}$.~~

~~13. The method of claim 11, wherein the output information includes:~~

~~a first likelihood that the sequence reads at the first position indicate the occurrence of a homozygous reference with a foreign allele that matches the alt,~~

~~a second likelihood that the sequence reads at the first position indicate the occurrence of a homozygous alt with a foreign allele that matches a reference allele,~~

~~a third output information for a first hypothesis of the one or more hypotheses that includes a likelihood that the sequence reads at the first position indicate the occurrence of a homozygous reference with a sequencing error that matches the alt allele, and~~

~~a fourth output information for a second hypothesis of the one or more hypotheses that includes a likelihood that the sequence reads at the first position indicate the occurrence of a homozygous alt with a sequencing error that matches the reference allele.~~

[Claim ~~1~~4]. The method of claim 1, wherein the one or more memory devices received the pileup of aligned sequence reads from a Field Programmable Gate Array (FPGA) device, wherein the FPGA includes one or more configurable digital logic gates that have been configured as a mapping and alignment unit to perform read mapping and alignment.

[Claim ~~1~~5]. The method of claim ~~1~~4,

wherein the computer is configured to access the one or more memory devices using one or more wired or wireless networks,

wherein a Field Programmable Gate Array (FPGA) device and the one or more memory devices are housed in an expansion card that has been coupled to a circuit board of a sequencer,

wherein the sequencer is configured to generate sequence reads based on an input sample and store the generated sequence reads in the one or more memory devices, and

wherein the mapping and alignment unit of the FPGA is configured to access the one or more memory devices to obtain the generated sequence reads; or

wherein the computer and the sequencer are each configured to access the one or more memory devices using one or more wired or wireless networks,

wherein the Field Programmable Gate Array (FPGA) device and the one or more memory devices are housed in an expansion card that has been coupled to a circuit board of server that is located remotely from the computer and the sequencer,

wherein the sequencer is configured to generate sequence reads based on an input sample, provide the generated sequence reads to the server using the one or more wired or wireless networks for storage, of the generated sequence reads, in the one or more memory devices, and

wherein the mapping and alignment unit of the FPGA is configured to access the one or more memory devices to obtain the generated sequence reads.

~~16. The method of claim 14,~~

~~wherein the computer and the sequencer are each configured to access the one or more memory devices using one or more wired or wireless networks,~~

~~wherein the Field Programmable Gate Array (FPGA) device and the one or more memory devices are housed in an expansion card that has been coupled to a circuit board of server that is located remotely from the computer and the sequencer,~~

~~wherein the sequencer is configured to generate sequence reads based on an input sample, provide the generated sequence reads to the server using the one or more wired or wireless networks for storage, of the generated sequence reads, in the one or more memory devices, and~~

~~wherein the mapping and alignment unit of the FPGA is configured to access the one or more memory devices to obtain the generated sequence reads.~~

[Claim 6~~17~~]. A system comprising:

one or more computers and one or more storage devices storing instructions that are operable, when executed by the one or more computers, to cause the one or more computers to perform operations comprising:

accessing, by one or more computers and from one or more memory devices, a pileup of a plurality of sequence reads aligned to a first region of a reference genome;

obtaining, by the one or more computers, information describing one or more characteristics of each of the plurality of reads of the pileup corresponding to a first position of the reference genome, wherein the information describing the characteristics of the respective reads include information describing:

(i) a mapping quality score for each sequence read in the pileup at the first position and

(ii) a read-allele score for each sequence read in the pileup at the first position for each candidate allele at the first position,

(i) a read orientation for each sequence read of the pileup at the first position,

(ii) a position of each base at the first position within each sequence read of the pileup at the first position with reference to the 5′ end of the sequence read,

(iii) a read-allele score for each sequence read of the plurality of sequence reads for each candidate allele at the reference position, and

(iv) a base quality score for each read for the base at the first position;

providing, by the one or more computers and based on the obtained information, one or more inputs to a probability model describing the ~~one or more~~ characteristics of the plurality of reads of the pileup, wherein the probability model is configured to determine a score, for each hypothesis of one or more hypotheses selected based on the one or more inputs, that indicates whether the hypothesis is true;

obtaining, by the one or more computers, output information for each of the one or more hypotheses, wherein the output information for each of the one or more hypotheses is (i) -generated by the probability model based on the probability model's processing of the one or more inputs to the probability model describing the one or more characteristics of the respective reads of the pileup and (ii) indicative of a score that indicates whether the hypothesis is true, wherein the output information includes:

first output information for a first hypothesis of the one or more hypotheses that includes a likelihood that the sequence reads at the first position indicate the occurrence of a homozygous reference with a foreign allele that matches an alternative allele (alt);

second output information for a second hypothesis of the one or more hypotheses that includes a likelihood that the sequence reads at the first position indicate the occurrence of a homozygous alt with a foreign allele that matches a reference allele,

third output information for a third hypothesis of the one or more hypotheses that includes a likelihood that the sequence reads at the first position indicate the occurrence of a homozygous reference with a sequencing error that matches the alt allele, and

fourth output information for a fourth hypothesis of the one or more hypotheses that includes a likelihood that the sequence reads at the first position indicate the occurrence of a homozygous alt with a sequencing error that matches the reference allele; and

determining, by the one or more computers and based on the obtained output information generated by the probability model for each of the plurality of hypotheses, a likelihood that a true variant exists at the first position.

[Claim 7~~18~~]. The system of claim 6~~19~~, wherein determining, by the one or more computers and based on the obtained output information generated by the probability model for each of the plurality of hypotheses, a likelihood that a true variant exists at the first position comprises:

determining, by the one or more computers, a collective score that is based on the output information generated by the probability model for each of the plurality of hypothesis, wherein the collective score indicates a likelihood that the true variant exists;

determining, by the one or more computers, whether the score generated by the collective score satisfies a predetermined threshold; and

based on determining, by the one or more computers, that the collective score satisfies the predetermined threshold, adding information to a Variant Call Format (VCF) file indicating that a true variant exists at the first position, optionally wherein the information indicating that a true variant exists at the first position includes information identifying (i) the first position, (ii) a candidate alt allele at the first position, (iii) the collective score.

~~19. The system of claim 18, wherein the information indicating that a true variant exists at the first position includes information identifying (i) the first position, (ii) a candidate alt allele at the first position, (iii) the collective score.~~

~~20. The system of claim 17, wherein the information describing the one or more characteristics of the respective reads includes information describing (i) a mapping quality score for each sequence read in the pileup at the first position and (ii) a read-allele score for each sequence read in the pileup at the first position for each candidate allele at the first position.~~

[Claim ~~821~~]. The system of claim ~~206~~, wherein the read-allele score for each sequence read of the pileup at the first position is based on an output produced by a pair Hidden Markov Model (P-HMM) model, for each of the sequence reads at the first position, which indicates a probability of observing a sequence read r_i given a particular candidate allele $G_{m,\phi}$.

~~22. The system of claim 20, wherein the output information includes:~~

~~first output information for a first hypothesis of the one or more hypotheses that includes a likelihood that the sequence reads at the first position indicate the occurrence of a homozygous reference with a foreign allele that matches the alt; and~~

~~second output information for a second hypothesis of the one or more hypotheses that includes a likelihood that the sequence reads at the first position indicate the occurrence of a homozygous alt with a foreign allele that matches a reference allele.~~

~~23. The system of claim 17, wherein the information describing the one or more characteristics of the respective sequence reads includes information describing (i) a read orientation for each sequence read of the pileup at the first position, (ii) a position of each base at the first position within each sequence read of the pileup at the first position with reference to the 5′ end of the sequence read, (iii) a read-allele score for each sequence read of the plurality of sequence reads for each candidate allele at the reference position, and (iv) a base quality score for each read for the base at the first position.~~

~~24. The system of claim 23, wherein the read-allele score for each sequence read of the pileup at the first position is based on an output produced by a P-HMM model, for each of the sequence reads at the first position, that indicates a probability of observing a sequence read r_i given a particular candidate allele $G_{m,\phi}$.~~

~~25. The system of claim 17, wherein the information describing the one or more characteristics of the respective sequence reads includes information describing (i) a read orientation for each sequence read of the pileup at the first position, (ii) a position of each base at the first position within each sequence read of the pileup at the first position with reference to the~~

5′ end of the sequence read, and (iii) a read allele score for each sequence read of the plurality of sequence reads for each candidate allele at the first position.

26. The system of claim 25, wherein the output information includes:

first output information for a first hypothesis of the one or more hypotheses that includes a likelihood that the sequence reads at the first position indicate the occurrence of a homozygous reference with a sequencing error that matches the alt allele, and

second output information for a second hypothesis of the one or more hypotheses that includes a likelihood that the sequence reads at the first position indicate the occurrence of a homozygous alt with a sequencing error that matches the reference allele.

27. The system of claim 17, wherein the information describing the one or more characteristics of the respective sequence reads includes information describing (i) a read orientation for each sequence read of the pileup at the first position, (ii) a position of each base at the first position such as position "0" 142 within each sequence read with reference to the 5′ end of the sequence read, (iii) a mapping quality score for each sequence read of the pileup at the first position, (iv) a read allele score for each sequence read of the plurality of reads for each candidate allele at the reference position, and (v) a base quality score for each read for the base aligned at the first position.

28. The system of claim 27, wherein the read allele score for each sequence read of the pileup at the first position is based on an output produced by a P HMM model, for each of the sequence reads at the first position, that indicates a probability of observing a read r_i given a particular candidate allele $G_{m,\phi}$.

29. The system of claim 27, wherein the output information includes:

a first likelihood that the sequence reads at the first position indicate the occurrence of a homozygous reference with a foreign allele that matches the alt,

a second likelihood that the sequence reads at the first position indicate the occurrence of a homozygous alt with a foreign allele that matches a reference allele,

~~a third output information for a first hypothesis of the one or more hypotheses that includes a likelihood that the sequence reads at the first position indicate the occurrence of a homozygous reference with a sequencing error that matches the alt allele, and~~

~~a fourth output information for a second hypothesis of the one or more hypotheses that includes a likelihood that the sequence reads at the first position indicate the occurrence of a homozygous alt with a sequencing error that matches the reference allele.~~

[Claim 9~~30~~]. The system of claim 6~~17~~, wherein the one or more memory devices received the pileup of aligned sequence reads from a Field Programmable Gate Array (FPGA) device, wherein the FPGA includes one or more configurable digital logic gates that have been configured as a mapping and alignment unit to perform read mapping and alignment.

[Claim ~~3~~10]. The system of claim 9~~30~~,

wherein the computer is configured to access the one or more memory devices using one or more wired or wireless networks,

wherein a Field Programmable Gate Array (FPGA) device and the one or more memory devices are housed in an expansion card that has been coupled to a circuit board of a sequencer,

wherein the sequencer is configured to generate sequence reads based on an input sample and store the generated sequence reads in the one or more memory devices, and

wherein the mapping and alignment unit of the FPGA is configured to access the one or more memory devices to obtain the generated sequence reads; or

wherein the computer and the sequencer are each configured to access the one or more memory devices using one or more wired or wireless networks,

wherein the Field Programmable Gate Array (FPGA) device and the one or more memory devices are housed in an expansion card that has been coupled to a circuit board of server that is located remotely from the computer and the sequencer,

wherein the sequencer is configured to generate sequence reads based on an input sample, provide the generated sequence reads to the server using the one or more wired or wireless networks for storage, of the generated sequence reads, in the one or more memory devices, and

wherein the mapping and alignment unit of the FPGA is configured to access the one or more memory devices to obtain the generated sequence reads.

~~32. The system of claim 30,~~

~~wherein the computer and the sequencer are each configured to access the one or more memory devices using one or more wired or wireless networks,~~

~~wherein the Field Programmable Gate Array (FPGA) device and the one or more memory devices are housed in an expansion card that has been coupled to a circuit board of server that is located remotely from the computer and the sequencer,~~

~~wherein the sequencer is configured to generate sequence reads based on an input sample, provide the generated sequence reads to the server using the one or more wired or wireless networks for storage, of the generated sequence reads, in the one or more memory devices, and~~

~~wherein the mapping and alignment unit of the FPGA is configured to access the one or more memory devices to obtain the generated sequence reads.~~

[Claim 113~~3~~]. A non-transitory computer-readable medium storing software comprising instructions executable by one or more computers which, upon such execution, cause the one or more computers to perform operations comprising:

accessing, by the one or more computers and from one or more memory devices, a pileup of a plurality of sequence reads aligned to a first region of a reference genome;

obtaining, by the one or more computers, information describing one or more characteristics of each of the plurality of reads of the pileup corresponding to a first position of the reference genome, wherein the information describing the characteristics of the respective reads include information describing:

(i) a mapping quality score for each sequence read in the pileup at the first position,

(ii) a read-allele score for each sequence read in the pileup at the first position for each candidate allele at the first position,

(iii) a read orientation for each sequence read of the pileup at the first position,

(iv) a position of each base at the first position within each sequence read of the pileup at the first position with reference to the 5′ end of the sequence read,

(v) a read-allele score for each sequence read of the plurality of sequence reads for each candidate allele at the reference position, and

(vi) a base quality score for each read for the base at the first position;

providing, by the one or more computers and based on the obtained information, one or more inputs to a probability model describing the ~~one or more~~ characteristics of the plurality of reads of the pileup, wherein the probability model is configured to determine a score, for each hypothesis of one or more hypotheses selected based on the one or more inputs, that indicates whether the hypothesis is true;

obtaining, by the one or more computers, output information for each of the one or more hypotheses, wherein the output information for each of the one or more hypotheses is (i) generated by the probability model based on the probability model's processing of the one or more inputs to the probability model describing the one or more characteristics of the respective reads of the pileup and (ii) indicative of a score that indicates whether the hypothesis is true wherein the output information includes:

first output information for a first hypothesis of the one or more hypotheses that includes a likelihood that the sequence reads at the first position indicate the occurrence of a homozygous reference with a foreign allele that matches an alternative allele (alt);

second output information for a second hypothesis of the one or more hypotheses that includes a likelihood that the sequence reads at the first position indicate the occurrence of a homozygous alt with a foreign allele that matches a reference allele,

third output information for a third hypothesis of the one or more hypotheses that includes a likelihood that the sequence reads at the first position indicate the occurrence of a homozygous reference with a sequencing error that matches the alt allele, and

fourth output information for a fourth hypothesis of the one or more hypotheses that includes a likelihood that the sequence reads at the first position indicate the occurrence of a homozygous alt with a sequencing error that matches the reference allele; and

determining, by the one or more computers and based on the obtained output information generated by the probability model for each of the plurality of hypotheses, a likelihood that a true variant exists at the first position.

[Claim 1234]. The computer-readable medium of claim 1133, wherein determining, by the one or more computers and based on the obtained output information generated by the probability model for each of the plurality of hypotheses, a likelihood that a true variant exists at the first position comprises:

determining, by the one or more computers, a collective score that is based on the output information generated by the probability model for each of the plurality of hypothesis, wherein the collective score indicates a likelihood that the true variant exists;

determining, by the one or more computers, whether the score generated by the collective score satisfies a predetermined threshold; and

based on determining, by the one or more computers, that the collective score satisfies the predetermined threshold, adding information to a VCF file indicating that a true variant exists at the first position, optionally wherein the information indicating that a true variant exists at the first position includes information identifying (i) the first position, (ii) a candidate alt allele at the first position, (iii) the collective score.

~~35. The computer-readable medium of claim 34, wherein the information indicating that a true variant exists at the first position includes information identifying (i) the first position, (ii) a candidate alt allele at the first position, (iii) the collective score.~~

~~36. The computer-readable medium of claim 33, wherein the information describing the one or more characteristics of the respective reads includes information describing (i) a mapping quality score for each sequence read in the pileup at the first position and (ii) a read-allele score for each sequence read in the pileup at the first position for each candidate allele at the first position.~~

[Claim 13~~7~~]. The computer-readable medium of claim 11~~36~~, wherein the read-allele score for each sequence read of the pileup at the first position is based on an output produced by a pair Hidden Markov Model (P-HMM) model, for each of the sequence reads at the first position, which indicates a probability of observing a sequence read r_i given a particular candidate allele $G_{m,\phi}$.

~~38. The computer-readable medium of claim 36, wherein the output information includes:~~

~~first output information for a first hypothesis of the one or more hypotheses that includes a likelihood that the sequence reads at the first position indicate the occurrence of a homozygous reference with a foreign allele that matches the alt; and~~

~~second output information for a second hypothesis of the one or more hypotheses that includes a likelihood that the sequence reads at the first position indicate the occurrence of a homozygous alt with a foreign allele that matches a reference allele.~~

39. The computer-readable medium of claim 33, wherein the information describing the one or more characteristics of the respective sequence reads includes information describing (i) a read orientation for each sequence read of the pileup at the first position, (ii) a position of each base at the first position within each sequence read of the pileup at the first position with reference to the 5′ end of the sequence read, (iii) a read-allele score for each sequence read of the plurality of sequence reads for each candidate allele at the reference position, and (iv) a base quality score for each read for the base at the first position.

40. The computer-readable medium of claim 39, wherein the read-allele score for each sequence read of the pileup at the first position is based on an output produced by a P-HMM model, for each of the sequence reads at the first position, that indicates a probability of observing a sequence read r_i given a particular candidate allele $G_{m,\phi}$.

41. The computer-readable medium of claim 39, wherein the information describing the one or more characteristics of the respective sequence reads includes information describing (i) a read orientation for each sequence read of the pileup at the first position, (ii) a position of each base at the first position within each sequence read of the pileup at the first position with reference to the 5′ end of the sequence read, and (iii) a read-allele score for each sequence read of the plurality of sequence reads for each candidate allele at the first position.

42. The computer-readable medium of claim 39, wherein the output information includes:

first output information for a first hypothesis of the one or more hypotheses that includes a likelihood that the sequence reads at the first position indicate the occurrence of a homozygous reference with a sequencing error that matches the alt allele, and

second output information for a second hypothesis of the one or more hypotheses that includes a likelihood that the sequence reads at the first position indicate the occurrence of a homozygous alt with a sequencing error that matches the reference allele.

43. The computer-readable medium of claim 33, wherein the information describing the one or more characteristics of the respective sequence reads includes information describing (i) a

~~read orientation for each sequence read of the pileup at the first position, (ii) a position of each base at the first position such as position "0" 142 within each sequence read with reference to the 5′ end of the sequence read, (iii) a mapping quality score for each sequence read of the pileup at the first position, (iv) a read-allele score for each sequence read of the plurality of reads for each candidate allele at the reference position, and (v) a base quality score for each read for the base aligned at the first position.~~

~~44. The computer-readable medium of claim 43, wherein the read-allele score for each sequence read of the pileup at the first position is based on an output produced by a P-HMM model, for each of the sequence reads at the first position, that indicates a probability of observing a read r_i given a particular candidate allele $G_{m,\phi}$.~~

~~45. The computer-readable medium of claim 43, wherein the output information includes:~~

~~a first likelihood that the sequence reads at the first position indicate the occurrence of a homozygous reference with a foreign allele that matches the alt,~~

~~a second likelihood that the sequence reads at the first position indicate the occurrence of a homozygous alt with a foreign allele that matches a reference allele,~~

~~a third output information for a first hypothesis of the one or more hypotheses that includes a likelihood that the sequence reads at the first position indicate the occurrence of a homozygous reference with a sequencing error that matches the alt allele, and~~

~~a fourth output information for a second hypothesis of the one or more hypotheses that includes a likelihood that the sequence reads at the first position indicate the occurrence of a homozygous alt with a sequencing error that matches the reference allele.~~

[Claim 1~~4~~6]. The computer-readable medium of claim 11~~33~~, wherein the one or more memory devices received the pileup of aligned sequence reads from a Field Programmable Gate Array (FPGA) device, wherein the FPGA includes one or more configurable digital logic gates that have been configured as a mapping and alignment unit to perform read mapping and alignment.

[Claim 15~~47~~]. The computer-readable medium of claim 14~~36~~,

wherein the computer is configured to access the one or more memory devices using one or more wired or wireless networks,

wherein a Field Programmable Gate Array (FPGA) device and the one or more memory devices are housed in an expansion card that has been coupled to a circuit board of a sequencer,

wherein the sequencer is configured to generate sequence reads based on an input sample and store the generated sequence reads in the one or more memory devices, and

wherein the mapping and alignment unit of the FPGA is configured to access the one or more memory devices to obtain the generated sequence reads; or

wherein the computer and the sequencer are each configured to access the one or more memory devices using one or more wired or wireless networks,

wherein the Field Programmable Gate Array (FPGA) device and the one or more memory devices are housed in an expansion card that has been coupled to a circuit board of server that is located remotely from the computer and the sequencer,

wherein the sequencer is configured to generate sequence reads based on an input sample, provide the generated sequence reads to the server using the one or more wired or wireless networks for storage, of the generated sequence reads, in the one or more memory devices, and

wherein the mapping and alignment unit of the FPGA is configured to access the one or more memory devices to obtain the generated sequence reads.

~~48. The computer-readable medium of claim 36,~~

~~wherein the computer and the sequencer are each configured to access the one or more memory devices using one or more wired or wireless networks,~~

~~wherein the Field Programmable Gate Array (FPGA) device and the one or more memory devices are housed in an expansion card that has been coupled to a circuit board of server that is located remotely from the computer and the sequencer,~~

~~wherein the sequencer is configured to generate sequence reads based on an input sample, provide the generated sequence reads to the server using the one or more wired or wireless networks for storage, of the generated sequence reads, in the one or more memory devices, and~~

~~wherein the mapping and alignment unit of the FPGA is configured to access the one or more memory devices to obtain the generated sequence reads.~~