

allele frequencies that are present in the original sample of DNA at each target loci (e.g., each polymorphic locus) from a set of target loci (e.g., polymorphic loci). While enrichment is particularly advantageous for methods for analyzing polymorphic loci, these enrichment methods can be readily adapted for nonpolymorphic loci if desired. In some embodiments the amplification and/or selective enrichment technique may involve PCR such as ligation mediated PCR, fragment capture by hybridization, Molecular Inversion Probes, or other circularizing probes. In some embodiments, methods for amplification or selective enrichment may involve using probes where, upon correct hybridization to the target sequence, the 3-prime end or 5-prime end of a nucleotide probe is separated from the polymorphic site of the allele by a small number of nucleotides. This separation reduces preferential amplification of one allele, termed allele bias. This is an improvement over methods that involve using probes where the 3-prime end or 5-prime end of a correctly hybridized probe are directly adjacent to or very near to the polymorphic site of an allele. In an embodiment, probes in which the hybridizing region may or certainly contains a polymorphic site are excluded. Polymorphic sites at the site of hybridization can cause unequal hybridization or inhibit hybridization altogether in some alleles, resulting in preferential amplification of certain alleles. These embodiments are improvements over other methods that involve targeted amplification and/or selective enrichment in that they better preserve the original allele frequencies of the sample at each polymorphic locus, whether the sample is pure genomic sample from a single individual or mixture of individuals.

[0467] The use of a technique to enrich a sample of DNA at a set of target loci followed by sequencing as part of a method for non-invasive prenatal allele calling or ploidy calling may confer a number of unexpected advantages. In some embodiments of the present disclosure, the method involves measuring genetic data for use with an informatics based method, such as PARENTAL SUPPORT™ (PS). The ultimate outcome of some of the embodiments is the actionable genetic data of an embryo or a fetus. There are many methods that may be used to measure the genetic data of the individual and/or the related individuals as part of embodied methods. In an embodiment, a method for enriching the concentration of a set of targeted alleles is disclosed herein, the method comprising one or more of the following steps: targeted amplification of genetic material, addition of loci specific oligonucleotide probes, ligation of specified DNA strands, isolation of sets of desired DNA, removal of unwanted components of a reaction, detection of certain sequences of DNA by hybridization, and detection of the sequence of one or a plurality of strands of DNA by DNA sequencing methods. In some cases the DNA strands may refer to target genetic material, in some cases they may refer to primers, in some cases they may refer to synthesized sequences, or combinations thereof. These steps may be carried out in a number of different orders.

[0468] For example, a universal amplification step of the DNA prior to targeted amplification may confer several advantages, such as removing the risk of bottlenecking and reducing allelic bias. The DNA may be mixed an oligonucleotide probe that can hybridize with two neighboring regions of the target sequence, one on either side. After hybridization, the ends of the probe may be connected by adding a polymerase, a means for ligation, and any neces-

sary reagents to allow the circularization of the probe. After circularization, an exonuclease may be added to digest to non-circularized genetic material, followed by detection of the circularized probe. The DNA may be mixed with PCR primers that can hybridize with two neighboring regions of the target sequence, one on either side. After hybridization, the ends of the probe may be connected by adding a polymerase, a means for ligation, and any necessary reagents to complete PCR amplification. Amplified or unamplified DNA may be targeted by hybrid capture probes that target a set of loci; after hybridization, the probe may be localized and separated from the mixture to provide a mixture of DNA that is enriched in target sequences.

[0469] The use of a method to target certain loci followed by sequencing as part of a method for allele calling or ploidy calling may confer a number of unexpected advantages. Some methods by which DNA may be targeted, or preferentially enriched, include using circularizing probes, linked inverted probes (LIPs, MIPs), capture by hybridization methods such as SURESELECT, and targeted PCR or ligation-mediated PCR amplification strategies.

[0470] In some embodiments, a method of the present disclosure involves measuring genetic data for use with an informatics based method, such as PARENTAL SUPPORT™ (PS), which is described further herein. PARENTAL SUPPORT™ is an informatics based approach to manipulating genetic data, aspects of which are described herein. The ultimate outcome of some of the embodiments is the actionable genetic data of an embryo or a fetus followed by a clinical decision based on the actionable data. The algorithms behind the PS method take the measured genetic data of the target individual, often an embryo or fetus, and the measured genetic data from related individuals, and are able to increase the accuracy with which the genetic state of the target individual is known. In an embodiment, the measured genetic data is used in the context of making ploidy determinations during prenatal genetic diagnosis. In an embodiment, the measured genetic data is used in the context of making ploidy determinations or allele calls on embryos during in vitro fertilization. There are many methods that may be used to measure the genetic data of the individual and/or the related individuals in the aforementioned contexts. The different methods comprise a number of steps, those steps often involving amplification of genetic material, addition of oligonucleotide probes, ligation of specified DNA strands, isolation of sets of desired DNA, removal of unwanted components of a reaction, detection of certain sequences of DNA by hybridization, detection of the sequence of one or a plurality of strands of DNA by DNA sequencing methods. In some cases the DNA strands may refer to target genetic material, in some cases they may refer to primers, in some cases they may refer to synthesized sequences, or combinations thereof. These steps may be carried out in a number of different orders.

[0471] Note that in theory it is possible to target any number loci in the genome, anywhere from one loci to well over one million loci. If a sample of DNA is subjected to targeting, and then sequenced, the percentage of the alleles that are read by the sequencer will be enriched with respect to their natural abundance in the sample. The degree of enrichment can be anywhere from one percent (or even less) to ten-fold, a hundred-fold, a thousand-fold or even many million-fold. In the human genome there are roughly 3 billion base pairs, and nucleotides, comprising approxi-

mately 75 million polymorphic loci. The more loci that are targeted, the smaller the degree of enrichment is possible. The fewer the number of loci that are targeted, the greater degree of enrichment is possible, and the greater depth of read may be achieved at those loci for a given number of sequence reads.

[0472] In an embodiment of the present disclosure, the targeting or preferential may focus entirely on SNPs. In an embodiment, the targeting or preferential may focus on any polymorphic site. A number of commercial targeting products are available to enrich exons. Surprisingly, targeting exclusively SNPs, or exclusively polymorphic loci, is particularly advantageous when using a method for NPD that relies on allele distributions. There are also published methods for NPD using sequencing, for example U.S. Pat. No. 7,888,017, involving a read count analysis where the read counting focuses on counting the number of reads that map to a given chromosome, where the analyzed sequence reads do not focused on regions of the genome that are polymorphic. Those types of methodology that do not focus on polymorphic alleles would not benefit as much from targeting or preferential enrichment of a set of alleles.

[0473] In an embodiment of the present disclosure, it is possible to use a targeting method that focuses on SNPs to enrich a genetic sample in polymorphic regions of the genome. In an embodiment, it is possible to focus on a small number of SNPs, for example between 1 and 100 SNPs, or a larger number, for example, between 100 and 1,000, between 1,000 and 10,000, between 10,000 and 100,000 or more than 100,000 SNPs. In an embodiment, it is possible to focus on one or a small number of chromosomes that are correlated with live trisomic births, for example chromosomes 13, 18, 21, X and Y, or some combination thereof. In an embodiment, it is possible to enrich the targeted SNPs by a small factor, for example between 1.01 fold and 100 fold, or by a larger factor, for example between 100 fold and 1,000,000 fold, or even by more than 1,000,000 fold. In an embodiment of the present disclosure, it is possible to use a targeting method to create a sample of DNA that is preferentially enriched in polymorphic regions of the genome. In an embodiment, it is possible to use this method to create a mixture of DNA with any of these characteristics where the mixture of DNA contains maternal DNA and also free floating fetal DNA. In an embodiment, it is possible to use this method to create a mixture of DNA that has any combination of these factors. For example, the method described herein may be used to produce a mixture of DNA that comprises maternal DNA and fetal DNA, and that is preferentially enriched in DNA that corresponds to 200 SNPs, all of which are located on either chromosome 18 or 21, and which are enriched an average of 1000 fold. In another example, it is possible to use the method to create a mixture of DNA that is preferentially enriched in 10,000 SNPs that are all or mostly located on chromosomes 13, 18, 21, X and Y, and the average enrichment per loci is greater than 500 fold. Any of the targeting methods described herein can be used to create mixtures of DNA that are preferentially enriched in certain loci.

[0474] In some embodiments, a method of the present disclosure further includes measuring the DNA in the mixed fraction using a high throughput DNA sequencer, where the DNA in the mixed fraction contains a disproportionate number of sequences from one or more chromosomes, wherein the one or more chromosomes are taken from the

group comprising chromosome 13, chromosome 18, chromosome 21, chromosome X, chromosome Y and combinations thereof.

[0475] Described herein are three methods: multiplex PCR, targeted capture by hybridization, and linked inverted probes (LIPs), which may be used to obtain and analyze measurements from a sufficient number of polymorphic loci from a maternal plasma sample in order to detect fetal aneuploidy; this is not meant to exclude other methods of selective enrichment of targeted loci. Other methods may equally well be used without changing the essence of the method. In each case the polymorphism assayed may include single nucleotide polymorphisms (SNPs), small indels, or STRs. A preferred method involves the use of SNPs. Each approach produces allele frequency data; allele frequency data for each targeted locus and/or the joint allele frequency distributions from these loci may be analyzed to determine the ploidy of the fetus. Each approach has its own considerations due to the limited source material and the fact that maternal plasma consists of mixture of maternal and fetal DNA. This method may be combined with other approaches to provide a more accurate determination. In an embodiment, this method may be combined with a sequence counting approach such as that described in U.S. Pat. No. 7,888,017. The approaches described could also be used to detect fetal paternity noninvasively from maternal plasma samples. In addition each approach may be applied to other mixtures of DNA or pure DNA samples to detect the presence or absence of aneuploid chromosomes, to genotype a large number of SNP from degraded DNA samples, to detect segmental copy number variations (CNVs), to detect other genotypic states of interest, or some combination thereof.

Accurately Measuring the Allelic Distributions in a Sample

[0476] Current sequencing approaches can be used to estimate the distribution of alleles in a sample. One such method involves randomly sampling sequences from a pool DNA, termed shotgun sequencing. The proportion of a particular allele in the sequencing data is typically very low and can be determined by simple statistics. The human genome contains approximately 3 billion base pairs. So, if the sequencing method used make 100 bp reads, a particular allele will be measured about once in every 30 million sequence reads.

[0477] In an embodiment, a method of the present disclosure is used to determine the presence or absence of two or more different haplotypes that contain the same set of loci in a sample of DNA from the measured allele distributions of loci from that chromosome. The different haplotypes could represent two different homologous chromosomes from one individual, three different homologous chromosomes from a trisomic individual, three different homologous haplotypes from a mother and a fetus where one of the haplotypes is shared between the mother and the fetus, three or four haplotypes from a mother and fetus where one or two of the haplotypes are shared between the mother and the fetus, or other combinations. Alleles that are polymorphic between the haplotypes tend to be more informative, however any alleles where the mother and father are not both homozygous for the same allele will yield useful information through measured allele distributions beyond the information that is available from simple read count analysis.

[0478] Shotgun sequencing of such a sample, however, is extremely inefficient as it results in many sequences for regions that are not polymorphic between the different haplotypes in the sample, or are for chromosomes that are not of interest, and therefore reveal no information about the proportion of the target haplotypes. Described herein are methods that specifically target and/or preferentially enrich segments of DNA in the sample that are more likely to be polymorphic in the genome to increase the yield of allelic information obtained by sequencing. Note that for the measured allele distributions in an enriched sample to be truly representative of the actual amounts present in the target individual, it is critical that there is little or no preferential enrichment of one allele as compared to the other allele at a given loci in the targeted segments. Current methods known in the art to target polymorphic alleles are designed to ensure that at least some of any alleles present are detected. However, these methods were not designed for the purpose of measuring the unbiased allelic distributions of polymorphic alleles present in the original mixture. It is non-obvious that any particular method of target enrichment would be able to produce an enriched sample wherein the measured allele distributions would accurately represent the allele distributions present in the original unamplified sample better than any other method. While many enrichment methods may be expected, in theory, to accomplish such an aim, an ordinary person skilled in the art is well aware that there is a great deal of stochastic or deterministic bias in current amplification, targeting and other preferential enrichment methods. One embodiment of a method described herein allows a plurality of alleles found in a mixture of DNA that correspond to a given locus in the genome to be amplified, or preferentially enriched in a way that the degree of enrichment of each of the alleles is nearly the same. Another way to say this is that the method allows the relative quantity of the alleles present in the mixture as a whole to be increased, while the ratio between the alleles that correspond to each locus remains essentially the same as they were in the original mixture of DNA. For some reported methods, preferential enrichment of loci can result in allelic biases of more than 1%, more than 2%, more than 5% and even more than 10%. This preferential enrichment may be due to capture bias when using a capture by hybridization approach, or amplification bias which may be small for each cycle, but can become large when compounded over 20, 30 or 40 cycles. For the purposes of this disclosure, for the ratio to remain essentially the same means that the ratio of the alleles in the original mixture divided by the ratio of the alleles in the resulting mixture is between 0.95 and 1.05, between 0.98 and 1.02, between 0.99 and 1.01, between 0.995 and 1.005, between 0.998 and 1.002, between 0.999 and 1.001, or between 0.9999 and 1.0001. Note that the calculation of the allele ratios presented here may not be used in the determination of the ploidy state of the target individual, and may only a metric to be used to measure allelic bias.

[0479] In an embodiment, once a mixture has been preferentially enriched at the set of target loci, it may be sequenced using any one of the previous, current, or next generation of sequencing instruments that sequences a clonal sample (a sample generated from a single molecule; examples include ILLUMINA GAIIx, ILLUMINA HISEQ, LIFE TECHNOLOGIES SOLiD, 5500XL). The ratios can be evaluated by sequencing through the specific alleles

within the targeted region. These sequencing reads can be analyzed and counted according the allele type and the ratios of different alleles determined accordingly. For variations that are one to a few bases in length, detection of the alleles will be performed by sequencing and it is essential that the sequencing read span the allele in question in order to evaluate the allelic composition of that captured molecule. The total number of captured molecules assayed for the genotype can be increased by increasing the length of the sequencing read. Full sequencing of all molecules would guarantee collection of the maximum amount of data available in the enriched pool. However, sequencing is currently expensive, and a method that can measure allele distributions using a lower number of sequence reads will have great value. In addition, there are technical limitations to the maximum possible length of read as well as accuracy limitations as read lengths increase. The alleles of greatest utility will be of one to a few bases in length, but theoretically any allele shorter than the length of the sequencing read can be used. While allele variations come in all types, the examples provided herein focus on SNPs or variants contained of just a few neighboring base pairs. Larger variants such as segmental copy number variants can be detected by aggregations of these smaller variations in many cases as whole collections of SNP internal to the segment are duplicated. Variants larger than a few bases, such as STRs require special consideration and some targeting approaches work while others will not.

[0480] There are multiple targeting approaches that can be used to specifically isolate and enrich a one or a plurality of variant positions in the genome. Typically, these rely on taking advantage of the invariant sequence flanking the variant sequence. There are reports by others related to targeting in the context of sequencing where the substrate is maternal plasma (see, e.g., Liao et al., Clin. Chem. 2011; 57(1): pp. 92-101). However, these approaches use targeting probes that target exons, and do not focus on targeting polymorphic regions of the genome. In an embodiment, a method of the present disclosure involves using targeting probes that focus exclusively or almost exclusively on polymorphic regions. In an embodiment, a method of the present disclosure involves using targeting probes that focus exclusively or almost exclusively on SNPs. In some embodiments of the present disclosure, the targeted polymorphic sites consist of at least 10% SNPs, at least 20% SNPs, at least 30% SNPs, at least 40% SNPs, at least 50% SNPs, at least 60% SNPs, at least 70% SNPs, at least 80% SNPs, at least 90% SNPs, at least 95% SNPs, at least 98% SNPs, at least 99% SNPs, at least 99.9% SNPs, or exclusively SNPs.

[0481] In an embodiment, a method of the present disclosure can be used to determine genotypes (base composition of the DNA at specific loci) and relative proportions of those genotypes from a mixture of DNA molecules, where those DNA molecules may have originated from one or a number of genetically distinct individuals. In an embodiment, a method of the present disclosure can be used to determine the genotypes at a set of polymorphic loci, and the relative ratios of the amount of different alleles present at those loci. In an embodiment the polymorphic loci may consist entirely of SNPs. In an embodiment, the polymorphic loci can comprise SNPs, single tandem repeats, and other polymorphisms. In an embodiment, a method of the present disclosure can be used to determine the relative distributions of alleles at a set of polymorphic loci in a mixture of DNA,

where the mixture of DNA comprises DNA that originates from a mother, and DNA that originates from a fetus. In an embodiment, the joint allele distributions can be determined on a mixture of DNA isolated from blood from a pregnant woman. In an embodiment, the allele distributions at a set of loci can be used to determine the ploidy state of one or more chromosomes on a gestating fetus.

[0482] In an embodiment, the mixture of DNA molecules could be derived from DNA extracted from multiple cells of one individual. In an embodiment, the original collection of cells from which the DNA is derived may comprise a mixture of diploid or haploid cells of the same or of different genotypes, if that individual is mosaic (germline or somatic). In an embodiment, the mixture of DNA molecules could also be derived from DNA extracted from single cells. In an embodiment, the mixture of DNA molecules could also be derived from DNA extracted from mixture of two or more cells of the same individual, or of different individuals. In an embodiment, the mixture of DNA molecules could be derived from DNA isolated from biological material that has already liberated from cells such as blood plasma, which is known to contain cell free DNA. In an embodiment, the this biological material may be a mixture of DNA from one or more individuals, as is the case during pregnancy where it has been shown that fetal DNA is present in the mixture. In an embodiment, the biological material could be from a mixture of cells that were found in maternal blood, where some of the cells are fetal in origin. In an embodiment, the biological material could be cells from the blood of a pregnant which have been enriched in fetal cells.

Circularizing Probes

[0483] Some embodiments of the present disclosure involve the use of “Linked Inverted Probes” (LIPs), which have been previously described in the literature, to amplify the target loci before or after amplification using primers that are not LIPs in the multiplex PCR methods of the invention. LIPs is a generic term meant to encompass technologies that involve the creation of a circular molecule of DNA, where the probes are designed to hybridize to targeted region of DNA on either side of a targeted allele, such that addition of appropriate polymerases and/or ligases, and the appropriate conditions, buffers and other reagents, will complete the complementary, inverted region of DNA across the targeted allele to create a circular loop of DNA that captures the information found in the targeted allele. LIPs may also be called pre-circularized probes, pre-circularizing probes, or circularizing probes. The LIPs probe may be a linear DNA molecule between 50 and 500 nucleotides in length, and in an embodiment between 70 and 100 nucleotides in length; in some embodiments, it may be longer or shorter than described herein. Others embodiments of the present disclosure involve different incarnations, of the LIPs technology, such as Padlock Probes and Molecular Inversion Probes (MIPs).

[0484] One method to target specific locations for sequencing is to synthesize probes in which the 3' and 5' ends of the probes anneal to target DNA at locations adjacent to and on either side of the targeted region, in an inverted manner, such that the addition of DNA polymerase and DNA ligase results in extension from the 3' end, adding bases to single stranded probe that are complementary to the target molecule (gap-fill), followed by ligation of the new 3' end to the 5' end of the original probe resulting in a circular DNA

molecule that can be subsequently isolated from background DNA. The probe ends are designed to flank the targeted region of interest. One aspect of this approach is commonly called MIPS and has been used in conjunction with array technologies to determine the nature of the sequence filled in. One drawback to the use of MIPS in the context of measuring allele ratios is that the hybridization, circularization and amplification steps do not happen at equal rates for different alleles at the same loci. This results in measured allele ratios that are not representative of the actual allele ratios present in the original mixture.

[0485] In an embodiment, the circularizing probes are constructed such that the region of the probe that is designed to hybridize upstream of the targeted polymorphic locus and the region of the probe that is designed to hybridize downstream of the targeted polymorphic locus are covalently connected through a non-nucleic acid backbone. This backbone can be any biocompatible molecule or combination of biocompatible molecules. Some examples of possible biocompatible molecules are poly(ethylene glycol), polycarbonates, polyurethanes, polyethylenes, polypropylenes, sulfone polymers, silicone, cellulose, fluoropolymers, acrylic compounds, styrene block copolymers, and other block copolymers.

[0486] In an embodiment of the present disclosure, this approach has been modified to be easily amenable to sequencing as a means of interrogating the filled in sequence. In order to retain the original allelic proportions of the original sample at least one key consideration must be taken into account. The variable positions among different alleles in the gap-fill region must not be too close to the probe binding sites as there can be initiation bias by the DNA polymerase resulting in differential of the variants. Another consideration is that additional variations may be present in the probe binding sites that are correlated to the variants in the gap-fill region which can result unequal amplification from different alleles. In an embodiment of the present disclosure, the 3' ends and 5' ends of the pre-circularized probe are designed to hybridize to bases that are one or a few positions away from the variant positions (polymorphic sites) of the targeted allele. The number of bases between the polymorphic site (SNP or otherwise) and the base to which the 3' end and/or 5' of the pre-circularized probe is designed to hybridize may be one base, it may be two bases, it may be three bases, it may be four bases, it may be five bases, it may be six bases, it may be seven to ten bases, it may be eleven to fifteen bases, or it may be sixteen to twenty bases, twenty to thirty bases, or thirty to sixty bases. The forward and reverse primers may be designed to hybridize a different number of bases away from the polymorphic site. Circularizing probes can be generated in large numbers with current DNA synthesis technology allowing very large numbers of probes to be generated and potentially pooled, enabling interrogation of many loci simultaneously. It has been reported to work with more than 300,000 probes. Two papers that discuss a method involving circularizing probes that can be used to measure the genomic data of the target individual include: Porreca et al., *Nature Methods*, 2007 4(11), pp. 931-936.; and also Turner et al., *Nature Methods*, 2009, 6(5), pp. 315-316. The methods described in these papers may be used in combination with other methods described herein. Certain steps of the method from these two papers may be used in combination with other steps from other methods described herein.

[0487] In some embodiments of the methods disclosed herein, the genetic material of the target individual is optionally amplified, followed by hybridization of the pre-circularized probes, performing a gap fill to fill in the bases between the two ends of the hybridized probes, ligating the two ends to form a circularized probe, and amplifying the circularized probe, using, for example, rolling circle amplification. Once the desired target allelic genetic information is captured by circularizing appropriately designed oligonucleotide probes, such as in the LIPs system, the genetic sequence of the circularized probes may be being measured to give the desired sequence data. In an embodiment, the appropriately designed oligonucleotide probes may be circularized directly on unamplified genetic material of the target individual, and amplified afterwards. Note that a number of amplification procedures may be used to amplify the original genetic material, or the circularized LIPs, including rolling circle amplification, MDA, or other amplification protocols. Different methods may be used to measure the genetic information on the target genome, for example using high throughput sequencing, Sanger sequencing, other sequencing methods, capture-by-hybridization, capture-by-circularization, multiplex PCR, other hybridization methods, and combinations thereof.

[0488] Once the genetic material of the individual has been measured using one or a combination of the above methods, an informatics based method, such as the PARENTAL SUPPORT™ method, along with the appropriate genetic measurements, can then be used to determine the ploidy state of one or more chromosomes on the individual, and/or the genetic state of one or a set of alleles, specifically those alleles that are correlated with a disease or genetic state of interest. Note that the use of LIPs has been reported for multiplexed capture of genetic sequences, followed by genotyping with sequencing. However, the use of sequencing data resulting from a LIPs-based strategy for the amplification of the genetic material found in a single cell, a small number of cells, or extracellular DNA, has not been used for the purpose of determining the ploidy state of a target individual.

[0489] Applying an informatics based method to determine the ploidy state of an individual from genetic data as measured by hybridization arrays, such as the ILLUMINA INFINIUM array, or the AFFYMETRIX gene chip has been described in documents references elsewhere in this document. However, the method described herein shows improvements over methods described previously in the literature. For example, the LIPs based approach followed by high throughput sequencing unexpectedly provides better genotypic data due to the approach having better capacity for multiplexing, better capture specificity, better uniformity, and low allelic bias. Greater multiplexing allows more alleles to be targeted, giving more accurate results. Better uniformity results in more of the targeted alleles being measured, giving more accurate results. Lower rates of allelic bias result in lower rates of miscalls, giving more accurate results. More accurate results result in an improvement in clinical outcomes, and better medical care.

[0490] It is important to note that LIPs may be used as a method for targeting specific loci in a sample of DNA for genotyping by methods other than sequencing. For example, LIPs may be used to target DNA for genotyping using SNP arrays or other DNA or RNA based microarrays.

Ligation-Mediated PCR

[0491] Ligation-mediated PCR may be used to amplify the target loci before or after PCR amplification using primers that are not ligated. Ligation-mediated PCR is a method of PCR used to preferentially enrich a sample of DNA by amplifying one or a plurality of loci in a mixture of DNA, the method comprising: obtaining a set of primer pairs, where each primer in the pair contains a target specific sequence and a non-target sequence, where the target specific sequence is preferably designed to anneal to a target region, one upstream and one downstream from the polymorphic site, and which can be separated from the polymorphic site by 0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11-20, 21-30, 31-40, 41-50, 51-100, or more than 100; polymerization of the DNA from the 3-prime end of upstream primer to the fill the single strand region between it and the 5-prime end of the downstream primer with nucleotides complementary to the target molecule; ligation of the last polymerized base of the upstream primer to the adjacent 5-prime base of the downstream primer; and amplification of only polymerized and ligated molecules using the non-target sequences contained at the 5-prime end of the upstream primer and the 3-prime end of the downstream primer. Pairs of primers to distinct targets may be mixed in the same reaction. The non-target sequences serve as universal sequences such that of all pairs of primers that have been successfully polymerized and ligated may be amplified with a single pair of amplification primers.

Capture by Hybridization

[0492] In some embodiments, a method of the present disclosure may involve using any of the following capture by hybridization methods in addition to using multiplex PCR to amplify the target loci. Preferential enrichment of a specific set of sequences in a target genome can be accomplished in a number of ways. Elsewhere in this document is a description of how LIPs can be used to target a specific set of sequences, but in all of those applications, other targeting and/or preferential enrichment methods can be used equally well for the same ends. One example of another targeting method is the capture by hybridization approach. Some examples of commercial capture by hybridization technologies include AGILENT's SURE SELECT and ILLUMINA's TRUSEQ. In capture by hybridization, a set of oligonucleotides that is complimentary or mostly complimentary to the desired targeted sequences is allowed to hybridize to a mixture of DNA, and then physically separated from the mixture. Once the desired sequences have hybridized to the targeting oligonucleotides, the effect of physically removing the targeting oligonucleotides is to also remove the targeted sequences. Once the hybridized oligos are removed, they can be heated to above their melting temperature and they can be amplified. Some ways to physically remove the targeting oligonucleotides is by covalently bonding the targeting oligos to a solid support, for example a magnetic bead, or a chip. Another way to physically remove the targeting oligonucleotides is by covalently bonding them to a molecular moiety with a strong affinity for another molecular moiety. An example of such a molecular pair is biotin and streptavidin, such as is used in SURE SELECT. Thus that targeted sequences could be covalently attached to a biotin molecule, and after hybridization, a solid support with

streptavidin affixed can be used to pull down the biotinylated oligonucleotides, to which are hybridized to the targeted sequences.

[0493] Hybrid capture involves hybridizing probes that are complementary to the targets of interest to the target molecules. Hybrid capture probes were originally developed to target and enrich large fractions of the genome with relative uniformity between targets. In that application, it was important that all targets be amplified with enough uniformity that all regions could be detected by sequencing, however, no regard was paid to retaining the proportion of alleles in original sample. Following capture, the alleles present in the sample can be determined by direct sequencing of the captured molecules. These sequencing reads can be analyzed and counted according the allele type. However, using the current technology, the measured allele distributions the captured sequences are typically not representative of the original allele distributions.

[0494] In an embodiment, detection of the alleles is performed by sequencing. In order to capture the allele identity at the polymorphic site, it is essential that the sequencing read span the allele in question in order to evaluate the allelic composition of that captured molecule. Since the capture molecules are often of variable lengths upon sequencing cannot be guaranteed to overlap the variant positions unless the entire molecule is sequenced. However, cost considerations as well as technical limitations as to the maximum possible length and accuracy of sequencing reads make sequencing the entire molecule unfeasible. In an embodiment, the read length can be increased from about 30 to about 50 or about 70 bases can greatly increase the number of reads that overlap the variant positions within the targeted sequences.

[0495] Another way to increase the number of reads that interrogate the position of interest is to decrease the length of the probe, as long as it does not result in bias in the underlying enriched alleles. The length of the synthesized probe should be long enough such that two probes designed to hybridize to two different alleles found at one locus will hybridize with near equal affinity to the various alleles in the original sample. Currently, methods known in the art describe probes that are typically longer than 120 bases. In a current embodiment, if the allele is one or a few bases then the capture probes may be less than about 110 bases, less than about 100 bases, less than about 90 bases, less than about 80 bases, less than about 70 bases, less than about 60 bases, less than about 50 bases, less than about 40 bases, less than about 30 bases, and less than about 25 bases, and this is sufficient to ensure equal enrichment from all alleles. When the mixture of DNA that is to be enriched using the hybrid capture technology is a mixture comprising free floating DNA isolated from blood, for example maternal blood, the average length of DNA is quite short, typically less than 200 bases. The use of shorter probes results in a greater chance that the hybrid capture probes will capture desired DNA fragments. Larger variations may require longer probes. In an embodiment, the variations of interest are one (a SNP) to a few bases in length. In an embodiment, targeted regions in the genome can be preferentially enriched using hybrid capture probes wherein the hybrid capture probes are of a length below 90 bases, and can be less than 80 bases, less than 70 bases, less than 60 bases, less than 50 bases, less than 40 bases, less than 30 bases, or less than 25 bases. In an embodiment, to increase the chance that

the desired allele is sequenced, the length of the probe that is designed to hybridize to the regions flanking the polymorphic allele location can be decreased from above 90 bases, to about 80 bases, or to about 70 bases, or to about 60 bases, or to about 50 bases, or to about 40 bases, or to about 30 bases, or to about 25 bases.

[0496] There is a minimum overlap between the synthesized probe and the target molecule in order to enable capture. This synthesized probe can be made as short as possible while still being larger than this minimum required overlap. The effect of using a shorter probe length to target a polymorphic region is that there will be more molecules that overlap the target allele region. The state of fragmentation of the original DNA molecules also affects the number of reads that will overlap the targeted alleles. Some DNA samples such as plasma samples are already fragmented due to biological processes that take place in vivo. However, samples with longer fragments benefit from fragmentation prior to sequencing library preparation and enrichment. When both probes and fragments are short (~60-80 bp) maximum specificity may be achieved relatively few sequence reads failing to overlap the critical region of interest.

[0497] In an embodiment, the hybridization conditions can be adjusted to maximize uniformity in the capture of different alleles present in the original sample. In an embodiment, hybridization temperatures are decreased to minimize differences in hybridization bias between alleles. Methods known in the art avoid using lower temperatures for hybridization because lowering the temperature has the effect of increasing hybridization of probes to unintended targets. However, when the goal is to preserve allele ratios with maximum fidelity, the approach of using lower hybridization temperatures provides optimally accurate allele ratios, despite the fact that the current art teaches away from this approach. Hybridization temperature can also be increased to require greater overlap between the target and the synthesized probe so that only targets with substantial overlap of the targeted region are captured. In some embodiments of the present disclosure, the hybridization temperature is lowered from the normal hybridization temperature to about 40° C., to about 45° C., to about 50° C., to about 55° C., to about 60° C., to about 65, or to about 70° C.

[0498] In an embodiment, the hybrid capture probes can be designed such that the region of the capture probe with DNA that is complementary to the DNA found in regions flanking the polymorphic allele is not immediately adjacent to the polymorphic site. Instead, the capture probe can be designed such that the region of the capture probe that is designed to hybridize to the DNA flanking the polymorphic site of the target is separated from the portion of the capture probe that will be in van der Waals contact with the polymorphic site by a small distance that is equivalent in length to one or a small number of bases. In an embodiment, the hybrid capture probe is designed to hybridize to a region that is flanking the polymorphic allele but does not cross it; this may be termed a flanking capture probe. The length of the flanking capture probe may be less than about 120 bases, less than about 110 bases, less than about 100 bases, less than about 90 bases, and can be less than about 80 bases, less than about 70 bases, less than about 60 bases, less than about 50 bases, less than about 40 bases, less than about 30 bases, or less than about 25 bases. The region of the genome that

is targeted by the flanking capture probe may be separated by the polymorphic locus by 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11-20, or more than 20 base pairs.

[0499] Description of a targeted capture based disease screening test using targeted sequence capture. Custom targeted sequence capture, like those currently offered by AGILENT (SURE SELECT), ROCHE-NIMBLEGEN, or ILLUMINA. Capture probes could be custom designed to ensure capture of various types of mutations. For point mutations, one or more probes that overlap the point mutation should be sufficient to capture and sequence the mutation.

[0500] For small insertions or deletions, one or more probes that overlap the mutation may be sufficient to capture and sequence fragments comprising the mutation. Hybridization may be less efficient between the probe-limiting capture efficiency, typically designed to the reference genome sequence. To ensure capture of fragments comprising the mutation one could design two probes, one matching the normal allele and one matching the mutant allele. A longer probe may enhance hybridization. Multiple overlapping probes may enhance capture. Finally, placing a probe immediately adjacent to, but not overlapping, the mutation may permit relatively similar capture efficiency of the normal and mutant alleles.

[0501] For Simple Tandem Repeats (STRs), a probe overlapping these highly variable sites is unlikely to capture the fragment well. To enhance capture a probe could be placed adjacent to, but not overlapping the variable site. The fragment could then be sequenced as normal to reveal the length and composition of the STR.

[0502] For large deletions, a series of overlapping probes, a common approach currently used in exon capture systems may work. However, with this approach it may be difficult to determine whether or not an individual is heterozygous. Targeting and evaluating SNPs within the captured region could potentially reveal loss of heterozygosity across the region indicating that an individual is a carrier. In an embodiment, it is possible to place non-overlapping or singleton probes across the potentially deleted region and use the number of fragments captured as a measure of heterozygosity. In the case where an individual carries a large deletion, one-half the number of fragments are expected to be available for capture relative to a non-deleted (diploid) reference locus. Consequently, the number of reads obtained from the deleted regions should be roughly half that obtained from a normal diploid locus. Aggregating and averaging the sequencing read depth from multiple singleton probes across the potentially deleted region may enhance the signal and improve confidence of the diagnosis. The two approaches, targeting SNPs to identify loss of heterozygosity and using multiple singleton probes to obtain a quantitative measure of the quantity of underlying fragments from that locus can also be combined. Either or both of these strategies may be combined with other strategies to better obtain the same end.

[0503] If during testing cfDNA detection of a male fetus, as indicated by the presence of the Y-chromosome fragments, captured and sequenced in the same test, and either an X-linked dominant mutation where mother and father are unaffected, or a dominant mutation where mother is not affected would indicate heightened risk to the fetus. Detection of two mutant recessive alleles within the same gene in an unaffected mother would imply the fetus had inherited a

mutant allele from father and potentially a second mutant allele from mother. In all cases, follow-up testing by amniocentesis or chorionic villus sampling may be indicated.

[0504] A targeted capture based disease screening test could be combined with a targeted capture based non-invasive prenatal diagnostic test for aneuploidy.

[0505] There are a number of ways to decrease depth of read (DOR) variability: for example, one could increase primer concentrations, one could use longer targeted amplification probes, or one could run more STA cycles (such as more than 25, more than 30, more than 35, or even more than 40)

Exemplary Methods of Determining the Number of DNA Molecules in a Sample.

[0506] A method is described herein to determine the number of DNA molecules in a sample by generating a uniquely identified molecule for each original DNA molecules in the sample during the first round of DNA amplification. Described here is a procedure to accomplish the above end followed by a single molecule or clonal sequencing method.

[0507] The approach entails targeting one or more specific loci and generating a tagged copy of the original molecules such manner that most or all of the tagged molecules from each targeted locus will have a unique tag and can be distinguished from one another upon sequencing of this barcode using clonal or single molecule sequencing. Each unique sequenced barcode represents a unique molecule in the original sample. Simultaneously, sequencing data is used to ascertain the locus from which the molecule originates. Using this information one can determine the number of unique molecules in the original sample for each locus.

[0508] This method can be used for any application in which quantitative evaluation of the number of molecules in an original sample is required. Furthermore, the number of unique molecules of one or more targets can be related to the number of unique molecules to one or more other targets to determine the relative copy number, allele distribution, or allele ratio. Alternatively, the number of copies detected from various targets can be modeled by a distribution in order to identify the mostly likely number of copies of the original targets. Applications include but are not limited to detection of insertions and deletions such as those found in carriers of Duchenne Muscular Dystrophy; quantitation of deletions or duplications segments of chromosomes such as those observed in copy number variants; chromosome copy number of samples from born individuals; chromosome copy number of samples from unborn individuals such as embryos or fetuses.

[0509] The method can be combined with simultaneous evaluation of variations contained in the targeted by sequence. This can be used to determine the number of molecules representing each allele in the original sample. This copy number method can be combined with the evaluation of SNPs or other sequence variations to determine the chromosome copy number of born and unborn individuals; the discrimination and quantification of copies from loci which have short sequence variations, but in which PCR may amplify from multiple target regions such as in carrier detection of Spinal Muscle Atrophy; determination of copy number of different sources of molecules from samples consisting of mixtures of different individual such as in

detection of fetal aneuploidy from free floating DNA obtained from maternal plasma.

[0510] In an embodiment, the method as it pertains to a single target locus may comprise one or more of the following steps: (1) Designing a standard pair of oligomers for PCR amplification of a specific locus. (2) Adding, during synthesis, a sequence of specified bases with no or minimal complementarity to the target locus or genome to the 5' end of the one of the target specific oligomer. This sequence, termed the tail, is a known sequence, to be used for subsequent amplification, followed by a sequence of random nucleotides. These random nucleotides comprise the random region. The random region comprises a randomly generated sequence of nucleic acids that probabilistically differ between each probe molecule. Consequently, following synthesis, the tailed oligomer pool will consist of a collection of oligomers beginning with a known sequence followed by unknown sequence that differs between molecules, followed by the target specific sequence. (3) Performing one round of amplification (denaturation, annealing, extension) using only the tailed oligomer. (4) Adding exonuclease to the reaction, effectively stopping the PCR reaction, and incubating the reaction at the appropriate temperature to remove forward single stranded oligos that did not anneal to template and extend to form a double stranded product. (5) Incubating the reaction at a high temperature to denature the exonuclease and eliminate its activity. (6) Adding to the reaction a new oligonucleotide that is complementary to tail of the oligomer used in the first reaction along with the other target specific oligomer to enable PCR amplification of the product generated in the first round of PCR. (7) Continuing amplification to generate enough product for downstream clonal sequencing. (8) Measuring the amplified PCR product by a multitude of methods, for example, clonal sequencing, to a sufficient number of bases to span the sequence.

[0511] In an embodiment, a method of the present disclosure involves targeting multiple loci in parallel or otherwise. Primers to different target loci can be generated independently and mixed to create multiplex PCR pools. In an embodiment, original samples can be divided into subpools and different loci can be targeted in each sub-pool before being recombined and sequenced. In an embodiment, the tagging step and a number of amplification cycles may be performed before the pool is subdivided to ensure efficient targeting of all targets before splitting, and improving subsequent amplification by continuing amplification using smaller sets of primers in subdivided pools.

[0512] One example of an application where this technology would be particularly useful is non-invasive prenatal aneuploidy diagnosis where the ratio of alleles at a given locus or a distribution of alleles at a number of loci can be used to help determine the number of copies of a chromosome present in a fetus. In this context, it is desirable to amplify the DNA present in the initial sample while maintaining the relative amounts of the various alleles. In some circumstances, especially in cases where there is a very small amount of DNA, for example, fewer than 5,000 copies of the genome, fewer than 1,000 copies of the genome, fewer than 500 copies of the genome, and fewer than 100 copies of the genome, one can encounter a phenomenon called bottlenecking. This is where there are a small number of copies of any given allele in the initial sample, and amplification biases can result in the amplified pool of DNA having significantly different ratios of those alleles than are

in the initial mixture of DNA. By applying a unique or nearly unique set of barcodes to each strand of DNA before standard PCR amplification, it is possible to exclude $n-1$ copies of DNA from a set of n identical molecules of sequenced DNA that originated from the same original molecule.

[0513] For example, imagine a heterozygous SNP in the genome of an individual, and a mixture of DNA from the individual where ten molecules of each allele are present in the original sample of DNA. After amplification there may be 100,000 molecules of DNA corresponding to that locus. Due to stochastic processes, the ratio of DNA could be anywhere from 1:2 to 2:1, however, since each of the original molecules was tagged with a unique tag, it would be possible to determine that the DNA in the amplified pool originated from exactly 10 molecules of DNA from each allele. This method would therefore give a more accurate measure of the relative amounts of each allele than a method not using this approach. For methods where it is desirable for the relative amount of allele bias to be minimized, this method will provide more accurate data.

[0514] Association of the sequenced fragment to the target locus can be achieved in a number of ways. In an embodiment, a sequence of sufficient length is obtained from the targeted fragment to span the molecule barcode as well as a sufficient number of unique bases corresponding to the target sequence to allow unambiguous identification of the target locus. In another embodiment, the molecular bar-coding primer that contains the randomly generated molecular barcode can also contain a locus specific barcode (locus barcode) that identifies the target to which it is to be associated. This locus barcode would be identical among all molecular bar-coding primers for each individual target and hence all resulting amplicons, but different from all other targets. In an embodiment, the tagging method described herein may be combined with a one-sided nesting protocol.

[0515] In an embodiment, the design and generation of molecular barcoding primers may be reduced to practice as follows: the molecular barcoding primers may consist of a sequence that is not complementary to the target sequence followed by random molecular barcode region followed by a target specific sequence. The sequence 5' of molecular barcode may be used for subsequent PCR amplification and may comprise sequences useful in the conversion of the amplicon to a library for sequencing. The random molecular barcode sequence could be generated in a multitude of ways. The preferred method synthesizes the molecule tagging primer in such a way as to include all four bases to the reaction during synthesis of the barcode region. All or various combinations of bases may be specified using the IUPAC DNA ambiguity codes. In this manner the synthesized collection of molecules will contain a random mixture of sequences in the molecular barcode region. The length of the barcode region will determine how many primers will contain unique barcodes. The number of unique sequences is related to the length of the barcode region as N^L where N is the number of bases, typically 4, and L is the length of the barcode. A barcode of five bases can yield up to 1024 unique sequences; a barcode of eight bases can yield 65536 unique barcodes. In an embodiment, the DNA can be measured by a sequencing method, where the sequence data represents the sequence of a single molecule. This can include methods in which single molecules are sequenced directly or methods in which single molecules are amplified to form clones

detectable by the sequence instrument, but that still represent single molecules, herein called clonal sequencing.

Exemplary Methods and Reagents for Quantification of Amplification Products

[0516] Quantitation of specific nucleic acid sequences of interest is typically done by quantitative real-time PCR techniques such as TAQMAN (LIFE TECHNOLOGIES), INVADER probes (THIRD WAVE TECHNOLOGIES), and the like. Such techniques suffer from numerous shortcomings such as limited ability to achieve the simultaneous analysis of multiple sequences in parallel (multiplexation) and the ability to provide accurate quantitative data for only a narrow range of possible amplification cycles (e.g., when the logarithm of PCR amplification production quantity versus the number of cycles is in the linear range). DNA sequencing techniques, particularly high throughput next-generation sequencing techniques (often referred to as massively parallel sequencing techniques) such as those employed in MYSEQ (ILLUMINA), HISEQ (ILLUMINA), ION TORRENT (LIFE TECHNOLOGIES), GENOME ANALYZER ILX (ILLUMINA), GS FLEX+(ROCHE 454) etc., can be used for by quantitative measurements of the number of copies of sequence of interest present in sample, thereby providing quantitative information about the starting materials, e.g., copy number or transcription levels. High throughput genetic sequencers are amenable to the use of bar coding (i.e., sample tagging with distinctive nucleic acid sequences) so as to identify specific samples from individuals thereby permitting the simultaneous analysis of multiple samples in a single run of the DNA sequencer. The number of times a given region of the genome in a library preparation (or other nucleic preparation of interest) is sequenced (number of reads) will be proportional to the number of copies of that sequence in the genome of interest (or expression level in the case of cDNA containing preparations). However, the preparation and sequencing of genetic libraries (and similar genome derived preparations) can introduce numerous biases that interfere with obtaining an accurate quantitative reading for the nucleic acid sequence of interest. For example, different nucleic acid sequences can amplify with different efficiencies during nucleic amplification steps that take place during the genetic library preparation or sample preparation.

[0517] The problem with differential amplification efficiencies can be mitigated by using certain embodiments of the subject invention. The subject invention includes various methods and compositions that relate to the use of standards for inclusion in amplification processes that can be used to improve the accuracy of quantitation. The invention is of use in, among other areas, the detection of aneuploidy in a fetus by analyzing free floating fetal DNA in maternal blood, as described herein and as described, among other places, U.S. Pat. Nos. 8,008,018; 7,332,277; PCT Published Application WO 2012/078792A2; and PCT Published Application WO 2011/146632 A1, which are each herein incorporated by reference in its entirety. Embodiments of the invention are also of use in the detection of aneuploidy in an in vitro generated embryos. Commercially significant aneuploidies that may be detected include aneuploidy of the human chromosomes 13, 18, 21, X and Y.

[0518] Embodiments of the invention may be used with either human or non-human nucleic acids, and may be applied to both animal and plant derived nucleic acids.

Embodiments of the invention may also be used to detect and/or quantitate alleles for other genetic disorders characterized by deletions or insertions. The deletion containing alleles can be detected in suspected carriers of the allele of interest.

[0519] One embodiment of the subject invention includes standards that are present in a known quantity (relative or absolute). For example, consider a genetic library made from a genetic source that is diploid for chromosome 8 (containing locus A) and triploid for chromosome 21 (containing locus B). A genetic library can be produced from this sample that will contain sequences in quantities that are a function of the number of chromosomes present in the sample, e.g., 200 copies of locus A and 300 copies of locus B. However, if locus A amplifies much more efficiently than locus B, after PCR there may be 60,000 copies of the A amplicon and 30,000 copies of the B amplicon, thus obscuring the true chromosomal copy number of the initial genomic sample when analysis by high throughput DNA sequencing (or other quantitative nucleic acid detection techniques). To mitigate this problem a standard sequence for locus A is employed, wherein the standard sequence amplifies with essentially the same efficiency as locus A. Similarly, a standard sequence for locus B is created, wherein the standard sequence amplifies with the essentially the same efficiency as locus B. A standard sequence of locus A and a standard sequence for locus B are added to the mixture prior to PCR (or other amplification techniques). These standard sequences are present in known quantities, either relative quantities or absolute quantities. Thus if a 1:1 mixture of standard sequence A and standard sequence B were added (prior to amplification) to the mixture in the previous example, 3000 copies of the standard A amplicon would be produced and 1000 copies of the standard B amplicon would be produced, showing that locus A is amplified 3 times more efficiently than locus B, under the same set of conditions.

[0520] In various embodiments one or more selected regions of a genome containing a SNP (or other polymorphism) of interest can be specifically amplified and subsequently sequenced. This target specific amplification can take place during the formation of a genetic library for sequencing. The library can contain numerous targeted regions for amplification. In some embodiments at least 10; 100; 300; 500; 750; 1,000; 2,000; 5,000; 7,500; 10,000; 15,000; 19,000; 20,000; 25,000; 27,000; 28,000; 30,000; 40,000; 50,000; 75,000; or 100,000 regions of interest. Examples of such libraries are described herein and can be found in U.S. Patent Application No. 2012/0270212, filed Nov. 18, 2011, which is herein incorporated by reference in its entirety.

[0521] Many high throughput DNA sequencing techniques require the modification of the genetic starting material, e.g., the ligation of universal priming sites and/or barcodes, so as to form libraries to facilitate the clonal amplification of small nucleic acid fragments prior to performing subsequent sequencing reactions. In some embodiments, one or more standard sequences are added during genetic library formation or added to a precursor component of a genetic library prior to amplification of the library. The standard sequences can be selected so as to mimic (yet be distinguishable based on nucleotide base sequence) target genomic fragments to be prepared for sequencing by a high throughput genetic sequencing technique. In one embodi-

ment, the standard sequence can be identical to the target genomic fragment excepting one, two, three, four to ten, or eleven to twenty nucleotides. In some embodiments, when the target genetic sequence contains a SNP, the standard sequence can be identical to the SNP excepting the nucleotide at the polymorphic base, which may be chosen to be one of the four nucleotides that is not observed at that location in nature. The standard sequences can be used in a highly multiplexed analysis of multiple target loci (such as polymorphic loci). Standard sequences can be added during the process of library formation (prior to amplification) in known quantities (relative or absolute) so as to provide a standard metric for greater accuracy in determining the amount of target sequence of interest in the sample of analysis. The combination of knowledge of the known quantities of the standard sequences used in conjunction with the knowledge of the ploidy level formation of library for sequencing formed from a genome of previously characterized ploidy level, e.g., known to be diploid for all autosomal chromosomes, can be used to calibrate the amplification properties of each standard sequence with respect to its corresponding target sequence and account for variations between batches of mixtures comprising multiple standard sequences. Given that it is often necessary to simultaneously analyze a large number of loci, it is useful to produce a mixture comprising a large set standard sequences. Embodiments of the invention include mixtures comprising multiple standard sequences. Ideally the amount of each standard sequence in the mixture is known with high precision. However, it is extremely difficult to achieve this ideal because as a practical matter there is a significant amount of variation in the quantity of each standard sequence in the mixture, particularly for mixtures comprising a large number of different synthetic oligonucleotides. This variation has numerous sources, e.g., variations in in vitro oligonucleotide synthesis reaction efficiencies between batch, inaccuracies in volume measurement, variations in pipetting. Furthermore, this variation can occur between different batches of that theoretically contain the exact same set of standard sequences in the exact same amounts. Accordingly, it is of interest to calibrate each batch of standard sequences independently. Batches of standard sequences can be calibrated against reference genomes of known chromosomal composition. Batched of standard sequences can be calibrated by sequencing the batch of standard sequences with minimal or no amplifications steps included in the sequencing protocol. Embodiments of the invention include calibrated mixtures of different standard sequences. Other embodiments of the invention include methods of calibrating mixtures of different standard sequences and calibrated mixtures of different standard sequences made by the subject methods.

[0522] Various embodiments of the subject mixtures of standard sequences and methods for using them can comprise at least 10; 100, 300; 500; 750; 1,000; 2,000; 5,000; 7,500; 10,000; 15,000; 19,000; 20,000; 25,000; 27,000; 28,000; 30,000; 40,000; 50,000; 75,000; or 100,000 or more standards sequences, as well as various intermediate amounts. The number of the standard sequences can be the same as the number of target sequences selected for analysis during the generation of a targeted library for DNA sequencing. However, in some embodiments, it may be advantageous to use a lower number of standard sequences than the number of targeted regions in the library being constructed. It may be advantageous to use the lower number so as avoid

coming up against the limits of the sequencing capacity of the high throughput DNA sequencer being employed. The number of standard sequences can be 50% or less than the number of targeted regions, 40% or less than the number of targeted regions, 30% or less than the number of targeted regions, 20% or less than the number of targeted regions, 10% or less than the number of targeted regions, 5% or less than the number of targeted regions, 1% or less than the number of targeted regions, as well as various intermediate values. For example, if a genetic library is created using 15,000 pairs of primers targeted to specific SNP containing loci, a suitable a mixture containing 1500 standard sequences corresponding to 1500 of the 15,000 targeted loci can be added prior to the amplification step of library constructions.

[0523] The amount of standard sequences added during library construction can vary considerably among different embodiments. In some embodiments, the amount of each standard sequence can be approximately the same as the predicted amount of the target sequence present in the genomic material sample used for library preparation. In other embodiments, the amount of each standard sequence can be greater or less than the predicted amount of the target sequence present in the genomic material sample used for library preparation. While the initial relative amounts of the target sequence and the standard sequence are not critical for the function of the invention, it is preferable that the amount be within the range 100 times greater to 100 times less than the amount of the target sequence present in the genomic material sample used for library preparation. Excessive amounts of standard may use too much sequencing capacity of the DNA sequencer in a given run of the instrument. Using too low an amount of standard sequences will produce insufficient data to aid in the analysis of variation in amplification efficiency.

[0524] The standard sequences may be selected to be very similar in nucleotide base sequence to the amplified regions of interest; preferably the standard sequence has the exact same primer-binding sites as the analyzed genomic region, i.e., the "target sequence." The standard sequence must be distinguishable from the corresponding target sequence at a given locus. For the sake of convenience, this distinguishable region of the standard sequence will be referred to as a "marker sequence." In some embodiments, the marker sequence region of the target sequences contains the polymorphic region, e.g., a SNP, and can be flanked on both sides by primer binding regions. The standard sequence may be selected to closely match the GC content of the corresponding target sequence. In some embodiments, the primer binding regions of the standard sequence are flanked by universal priming sites. These universal priming sites are selected to match universal priming sites used in a genomic library for analysis. In other embodiments, the standard sequences do not have universal priming sites and the universal priming sites are added during the creation of a library. Standard sequences are typically provided in single stranded form. A standard sequence is defined with respect to a corresponding target sequence and the sequence specific reagents used to amplify the target sequence. In some embodiments, the target sequence contains the polymorphism of interest, e.g., a SNP, a deletion, or insertion, present in the nucleic acid sample for analysis. The standard sequence is a synthetic polynucleotide that is similar in nucleotide base sequence to the target sequence, but is

nonetheless distinguishable from the target sequence by virtue of at least one nucleotide base difference, thereby providing a mechanism for distinguishing amplicon sequences derived from the standard sequence from amplicon sequences derived from the target sequence. Standard sequences are selected so as to have essentially the same amplification properties as the corresponding target sequence when amplified with the same set of amplification reagents, e.g., PCR primers. In some embodiments, the standard sequences can have the same primer sequence binding sites than the corresponding target sequences. In other embodiments, the standard sequences can have a different primer sequence binding sites than the corresponding target sequences. In some embodiments, the standard sequences can be selected to produce amplicons that have the same length as the length of amplicons produced from the corresponding target sequences. In other embodiments, the standard sequences can be selected to produce amplicons that have the slightly different lengths than the length of amplicons produced from the corresponding target sequences.

[0525] After the amplification reactions have been completed, the library is sequenced on a high throughput DNA sequencer where individual molecule are clonally amplified and sequenced. The number of sequence reads for each allele of the target sequence is counted, also counted are the number of sequence reads for the standard sequence corresponding to the target sequence. The process is also carried out for at least one other pair of target sequences and corresponding standard sequences. Consider for example, locus A, X_{A1} reads for allele 1 of locus A are produced; X_{A2} reads for allele 2 of locus A are produced, and X_{AC} reads for standard sequence A are produced. The ratio of (X_{A1} plus X_{A2}) to X_{AC} is determined for each locus of interest. As discussed earlier, the process can be performed on a reference genome, e.g., a genome that is known to be diploid for all chromosomes. The process can be repeated many times in order to provide a large number of read values so as to determine a mean number of reads and the standard deviation in the number of reads. The process is performed with a mixture comprising a large number of different standard sequences corresponding to different loci. By assuming that (1) X_{A1} plus X_{A2} corresponds to the known number of chromosome, e.g., 2 for the normal human female genome and (2) the standard sequences have similar amplification (and detectability) properties as their corresponding natural loci, the relative amounts of the different standard sequences in the multiplex standard mixture can be determined. The calibrated multiplex standard sequence mixture can then be used to adjust for the variability in amplification efficiency between the different loci in a multiplex amplification reaction.

[0526] Other embodiments of the invention include methods and compositions for measuring the copy number of specific genes of interest, including duplications and mutant genes characterized by large deletions that would interfere with quantitation by sequencing. Sequencing would have problems detecting alleles having such deletions. Standard sequences included the amplification process can be used to reduce this problem.

[0527] In one embodiment of the invention the target sequence for analysis is a gene having a wild type (i.e. functional) form and a mutant form characterized by a deletion. Exemplary of such genes is SMN1, an allele

having deletion being responsible for the genetic disease spinal muscular atrophy (SMA). It is of interest to detect an individual carrying the mutant form of the gene by means of high throughput genetic sequencing techniques. The application of such techniques to the detection of deletion mutations can be problematic because, among other reasons, the lack of sequences observed in sequencing (as opposed to detecting a simple point mutation or SNP). Such embodiments employ (1) a pair of amplification primers specific for the gene of interest, where in the amplification primers will amplify the gene of interest (or a portion thereof) and will not significantly amplify the mutant allele, (2) a standard sequence corresponding to the wild type allele of the gene of interest (i.e., a target sequence), but differing by at least one detectable nucleotide base, (3) a pair of amplification primers specific for a second target sequence that serves as a reference sequence, and (4) a standard sequence corresponding to the reference sequence.

[0528] In one embodiment of the invention is provided a method for measuring the number of copies of the gene of interest, where in the gene of interest has one meaning allele that comprises a deletion. The method can employ amplification reagent specific for the gene of interest, e.g., PCR primers, that are specific for the gene of interest by amplifying at least a portion of the gene of interest, or the entire gene of interest, or a region adjacent to the gene of interest, while not amplifying the deletion comprising allele of the gene of interest. Additionally the subject method employs a standard sequence corresponding to the gene of interest, wherein the standard sequence differs by at least one nucleotide base from the gene of interest (so that the sequence of the standard sequence can be readily distinguished from the naturally occurring gene of interest). Typically, the standard sequence will contain the same primer binding sites as the gene of interest so as to minimize any amplification discrimination between the gene of interest and the standard sequence corresponding to the gene of interest. The reaction will also comprises amplification reagents specific for a reference sequence. The reference sequence is a sequence of known (or at least assumed to be known) copy number in the genome to be analyzed. The reaction further comprises a standard sequence corresponding to the reference sequence. Typically, the standard sequence corresponding to the reference sequence will contain the same primer binding sites as the reference sequence so as to minimize any amplification discrimination between the reference sequence and the standard sequence corresponding to the reference sequence.

Exemplary PCR Conditions

[0529] If desired, any of the PCR conditions disclosed herein or any standard PCR conditions can be used to test a primer library to determine, e.g., the percent of primer dimers, percent of target amplicons, and percent of target loci that are amplified. If desired, standard methods can be used to optimize the reaction conditions to improve the performance of a primer library. Any of these PCR conditions may also be used in any of the methods of the invention to amplify target loci. It was determined that high ionic strength solutions can surprisingly be used for multiplex PCR. In some embodiments, monovalent cations are used to increase the ionic strength to, e.g., help the primers bind the template.

[0530] In some embodiments, the reaction volume includes ethylenediaminetetraacetic acid (EDTA), magne-

sium, tetramethyl ammonium chloride (TMAC), or any combination thereof. In some embodiments, the concentration of TMAC is between 20 and 80 mM, such as between 25 and 70 mM, 30 and 60 mM, 30 and 40 mM, 40 and 50 mM, 50 and 60 mM, or 60 and 70 mM, inclusive. While not meant to be bound to any particular theory, it is believed that TMAC binds to DNA, stabilizes duplexes, increases primer specificity, and/or equalizes the melting temperatures of different primers. In some embodiments, TMAC increases the uniformity in the amount of amplified products for the different targets. In some embodiments, the concentration of magnesium (such as magnesium from magnesium chloride) is between 1 and 10 mM, such as between 1 and 8 mM, 1 and 5 mM, 1 and 3 mM, 3 and 5 mM, 3 and 6 mM, or 5 and 8 mM, inclusive.

[0531] In some embodiments, the concentration of available magnesium (the concentration of magnesium that is assumed to be available for binding the polymerase and not bound to molecules other than the polymerase), such as the magnesium that is not bound by phosphate groups on dNTPs, primers, or nucleic acid templates, or carboxylic acid groups on magnetic or other beads, if present) is between 0.5 to 10 mM, such as between 1 and 8 mM, 1 and 5 mM, 1 and 3 mM, 3 and 5 mM, 3 and 6 mM, 4 and 6 mM, or 5 and 8 mM, inclusive. The large number of primers used for multiplex PCR of a large number of targets may chelate a lot of the magnesium (2 phosphates in the primers chelate 1 magnesium). For example, if enough primers are used such that the concentration of phosphate from the primers is ~9 mM, then the primers may reduce the effective magnesium concentration by ~4.5 mM. In some embodiments, EDTA is used to decrease the amount of magnesium available as a cofactor for the polymerase since high concentrations of magnesium can result in PCR errors, such as amplification of non-target loci. In some embodiments, the concentration of EDTA reduces the amount of available magnesium to between 1 and 5 mM (such as between 3 and 5 mM).

[0532] In some embodiments, the pH is between 7.5 and 8.5, such as between 7.5 and 8, 8 and 8.3, or 8.3 and 8.5, inclusive. In some embodiments, Tris is used at, for example, a concentration of between 10 and 100 mM, such as between 10 and 25 mM, 25 and 50 mM, 50 and 75 mM, or 25 and 75 mM, inclusive. In some embodiments, any of these concentrations of Tris are used at a pH between 7.5 and 8.5. In some embodiments, a combination of KCl and $(\text{NH}_4)_2\text{SO}_4$ is used, such as between 50 and 150 mM KCl and between 10 and 90 mM $(\text{NH}_4)_2\text{SO}_4$, inclusive. In some embodiments, the concentration of KCl is between 0 and 30 mM, between 50 and 100 mM, or between 100 and 150 mM, inclusive. In some embodiments, the concentration of $(\text{NH}_4)_2\text{SO}_4$ is between 10 and 50 mM, 50 and 90 mM, 10 and 20 mM, 20 and 40 mM, 40 mM and 60, or 60 mM and 80 mM $(\text{NH}_4)_2\text{SO}_4$, inclusive. In some embodiments, the ammonium $[\text{NH}_4^+]$ concentration is between 0 and 160 mM, such as between 0 to 50, 50 to 100, or 100 to 160 mM, inclusive. In some embodiments, the sum of the potassium and ammonium concentration ($[\text{K}^+] + [\text{NH}_4^+]$) is between 0 and 160 mM, such as between 0 to 25, 25 to 50, 50 to 150, 50 to 75, 75 to 100, 100 to 125, or 125 to 160 mM, inclusive. An exemplary buffer with $[\text{K}^+] + [\text{NH}_4^+] = 120$ mM is 20 mM KCl and 50 mM $(\text{NH}_4)_2\text{SO}_4$. In some embodiments, the buffer includes 25 to 75 mM Tris, pH 7.2 to 8, 0 to 50 mM KCl, 10 to 80 mM ammonium sulfate, and 3 to 6 mM

magnesium, inclusive. In some embodiments, the buffer includes 25 to 75 mM Tris pH 7 to 8.5, 3 to 6 mM MgCl_2 , 10 to 50 mM KCl, and 20 to 80 mM $(\text{NH}_4)_2\text{SO}_4$, inclusive. In some embodiments, 100 to 200 Units/mL of polymerase are used. In some embodiments, 100 mM KCl, 50 mM $(\text{NH}_4)_2\text{SO}_4$, 3 mM MgCl_2 , 7.5 nM of each primer in the library, 50 mM TMAC, and 7 μL DNA template in a 20 μL final volume at pH 8.1 is used.

[0533] In some embodiments, a crowding agent is used, such as polyethylene glycol (PEG, such as PEG 8,000) or glycerol. In some embodiments, the amount of PEG (such as PEG 8,000) is between 0.1 to 20%, such as between 0.5 to 15%, 1 to 10%, 2 to 8%, or 4 to 8%, inclusive. In some embodiments, the amount of glycerol is between 0.1 to 20%, such as between 0.5 to 15%, 1 to 10%, 2 to 8%, or 4 to 8%, inclusive. In some embodiments, a crowding agent allows either a low polymerase concentration and/or a shorter annealing time to be used. In some embodiments, a crowding agent improves the uniformity of the DOR and/or reduces dropouts (undetected alleles). For example, at 8% PEG, and 50 U/mL polymerase, the uniformity was as good as 150 U/mL polymerase and no PEG. If the error rate increases when PEG is included, a higher magnesium chloride concentration (such greater than or about 4, 5, 6, 7, 8, 9, or 10 MgCl_2) can be used to reduce or prevent the increase in error rate. Inclusion of 8% PEG 8,000 allowed successful multiplexing with an annealing time of only 1 minute at an annealing temperature of 63° C.

[0534] In some embodiments, a polymerase with proof-reading activity, a polymerase without (or with negligible) proof-reading activity, or a mixture of a polymerase with proof-reading activity and a polymerase without (or with negligible) proof-reading activity is used. In some embodiments, a hot start polymerase, a non-hot start polymerase, or a mixture of a hot start polymerase and a non-hot start polymerase is used. In some embodiments, a HotStarTaq DNA polymerase is used (see, for example, QIAGEN catalog No. 203203, see, e.g., information available at the World Wide Web at qiagen.com/us/products/catalog/assay-technologies/end-point-pcr-and-rt-pcr-reagents/hotstartaq-dna-polymerase/, which is hereby incorporated by reference in its entirety). In some embodiments, AmpliTaq Gold® DNA Polymerase is used; it is a chemically modified form of AmpliTaq® DNA Polymerase requiring thermal activation (see, for example, Applied Biosystems catalog No. N8080241 see, e.g., information available at the World Wide Web at lifetechnologies.com/order/catalog/product/N8080241, which is hereby incorporated by reference in its entirety). In some embodiments, KAPA Taq DNA Polymerase or KAPA Taq HotStart DNA Polymerase is used; they are based on the single-subunit, wild-type Taq DNA polymerase of the thermophilic bacterium *Thermus aquaticus*. KAPA Taq and KAPA Taq HotStart DNA Polymerase have 5'-3' polymerase and 5'-3' exonuclease activities, but no 3' to 5' exonuclease (proofreading) activity (see, for example, KAPA BIOSYSTEMS catalog No. BK1000 see, e.g., information available at the World Wide Web at kapabiosystems.com/product-applications/products/pcr-2/kapa-taq-pcr-kits/, which is hereby incorporated by reference in its entirety). In some embodiments, Pfu DNA polymerase is used; it is a highly thermostable DNA polymerase from the hyperthermophilic archaeum *Pyrococcus furiosus*. The enzyme catalyzes the template-dependent polymerization of nucleotides into duplex DNA in the 5'-3' direction. Pfu DNA

Polymerase also exhibits 3'-5' exonuclease (proofreading) activity that enables the polymerase to correct nucleotide incorporation errors. It has no 5'-3' exonuclease activity (see, for example, Thermo Scientific catalog No. EP0501 see, e.g., information available at the World Wide Web at thermoscientificbio.com/pcr-enzymes-master-mixes-and-reagents/pfu-dna-polymerase/, which is hereby incorporated by reference in its entirety). In some embodiments KlenTaqI is used; it is a Klenow-fragment analog of Taq DNA polymerase, it has no exonuclease or endonuclease activity (see, for example, DNA POLYMERASE TECHNOLOGY, Inc. St. Louis, Mo., catalog No. 100 see, e.g., information available at the World Wide Web at klenTaq.com/products/klenTaq, which is hereby incorporated by reference in its entirety). In some embodiments, the polymerase is a PUSHION DNA polymerase, such as PHUSION High Fidelity DNA polymerase (M0530S, New England BioLabs, Inc.) or PHUSION Hot Start Flex DNA polymerase (M0535S, New England BioLabs, Inc.; Frey and Suppmann *BioChemica*. 2:34-35, 1995; Chester and Marshak *Analytical Biochemistry*. 209:284-290, 1993, which are each hereby incorporated by reference in its entirety). In some embodiments, the polymerase is a Q5® DNA Polymerase, such as Q5® High-Fidelity DNA Polymerase (M0491S, New England BioLabs, Inc.) or Q5® Hot Start High-Fidelity DNA Polymerase (M0493S, New England BioLabs, Inc.). In some embodiments, the polymerase is a T4 DNA polymerase (M0203S, New England BioLabs, Inc.; Tabor and Struh. (1989). "DNA-Dependent DNA Polymerases," In Ausubel et al. (Ed.), *Current Protocols in Molecular Biology*. 3.5. 10-3.5.12. New York: John Wiley & Sons, Inc., 1989; Sambrook et al. *Molecular Cloning: A Laboratory Manual*. (2nd ed.), 5.44-5.47. Cold Spring Harbor: Cold Spring Harbor Laboratory Press, 1989, which are each hereby incorporated by reference in its entirety).

[0535] In some embodiment, between 5 and 600 Units/mL (Units per 1 mL of reaction volume) of polymerase is used, such as between 5 to 100, 100 to 200, 200 to 300, 300 to 400, 400 to 500, or 500 to 600 Units/mL, inclusive. One unit is commonly defined as the amount of enzyme that will incorporate 15 nmol of dNTP into acid-insoluble material in 30 minutes at 75° C. Exemplary assay conditions for measuring unit activity include 1× THERMOPOL Reaction Buffer, 200 µM dNTPs including [³H]-dTTP and 200 µg/ml activated Calf Thymus DNA (see, e.g., information available at the world wide web at neb.com/products/m0267-taq-dna-polymerase-with-thermopol-buffer, which is hereby incorporated by reference in its entirety). 1× THERMOPOL® Reaction Buffer contains 20 mM Tris-HCl, 10 mM (NH₄)₂SO₄, 10 mM KCl, 2 mM MgSO₄, and 0.1% TRITON® X-100, pH 8.8.

[0536] In some embodiments, hot-start PCR is used to reduce or prevent polymerization prior to PCR thermocycling. Exemplary hot-start PCR methods include initial inhibition of the DNA polymerase, or physical separation of reaction components reaction until the reaction mixture reaches the higher temperatures. In some embodiments, the enzyme is spatially separated from the reaction mixture by wax that melts when the reaction reaches high temperature. In some embodiments, slow release of magnesium is used. DNA polymerase requires magnesium ions for activity, so the magnesium is chemically separated from the reaction by binding to a chemical compound, and is released into the solution only at high temperature. In some embodiments,

non-covalent binding of an inhibitor is used. In this method a peptide, antibody, or aptamer are non-covalently bound to the enzyme at low temperature and inhibit its activity. After incubation at elevated temperature, the inhibitor is released and the reaction starts. In some embodiments, a cold-sensitive Taq polymerase is used, such as a modified DNA polymerase with almost no activity at low temperature. In some embodiments, chemical modification is used. In this method, a molecule is covalently bound to the side chain of an amino acid in the active site of the DNA polymerase. The molecule is released from the enzyme by incubation of the reaction mixture at elevated temperature. Once the molecule is released, the enzyme is activated.

[0537] In some embodiments, the amount to template nucleic acids (such as an RNA or DNA sample) is between 20 and 5,000 ng, such as between 20 to 200, 200 to 400, 400 to 600, 600 to 1,000; 1,000 to 1,500; or 2,000 to 3,000 ng, inclusive.

[0538] In some embodiments QIAGEN Multiplex PCR Kit is used (QIAGEN catalog No. 206143; see, e.g., information available at the World Wide Web at qiagen.com/products/catalog/assay-technologies/end-point-pcr-and-rt-pcr-reagents/qiagen-multiplex-pcr-kit, which is hereby incorporated by reference in its entirety). For 100×50 µl multiplex PCR reactions, the kit includes 2× QIAGEN Multiplex PCR Master Mix (providing a final concentration of 3 mM MgCl₂, 3×0.85 ml), 5× Q-Solution (1×2.0 ml), and RNase-Free Water (2×1.7 ml). The QIAGEN Multiplex PCR Master Mix (MM) contains a combination of KCl and (NH₄)₂SO₄ as well as the PCR additive, Factor MP, which increases the local concentration of primers at the template. Factor MP stabilizes specifically bound primers, allowing efficient primer extension by HotStarTaq DNA Polymerase. HotStarTaq DNA Polymerase is a modified form of Taq DNA polymerase and has no polymerase activity at ambient temperatures. In some embodiments, HotStarTaq DNA Polymerase is activated by a 15-minute incubation at 95° C. which can be incorporated into any existing thermal-cycler program.

[0539] In some embodiments, 1× QIAGEN MM final concentration (the recommended concentration), 7.5 nM of each primer in the library, 50 mM TMAC, and 7 µl DNA template in a 20 µl final volume is used. In some embodiments, the PCR thermocycling conditions include 95° C. for 10 minutes (hot start); 20 cycles of 96° C. for 30 seconds; 65° C. for 15 minutes; and 72° C. for 30 seconds; followed by 72° C. for 2 minutes (final extension); and then a 4° C. hold.

[0540] In some embodiments, 2× QIAGEN MM final concentration (twice the recommended concentration), 2 nM of each primer in the library, 70 mM TMAC, and 7 µl DNA template in a 20 µl total volume is used. In some embodiments, up to 4 mM EDTA is also included. In some embodiments, the PCR thermocycling conditions include 95° C. for 10 minutes (hot start); 25 cycles of 96° C. for 30 seconds; 65° C. for 20 minutes; and 72° C. for 30 seconds; followed by 72° C. for 2 minutes (final extension); and then a 4° C. hold.

[0541] Another exemplary set of PCR thermocycling conditions includes 95° C. for 10 minutes, 15 cycles of 95° C. for 30 seconds, 65° C. for 1 minute, 60° C. for 5 minutes, 65° C. for 5 minutes and 72° C. for 30 seconds; and then 72° C. for 2 minutes. In some embodiments, this set of PCR thermocycling conditions is used with the following reaction

conditions: 100 mM KCl, 50 mM $(\text{NH}_4)_2\text{SO}_4$, 3 mM MgCl_2 , 7.5 nM of each primer in the library, 50 mM TMAC, and 7 μl DNA template in a 20 μl final volume at pH 8.1.

[0542] Another exemplary set of conditions includes a semi-nested PCR approach. The first PCR reaction uses 20 μl a reaction volume with 2 $\times$ QIAGEN MM final concentration, 1.875 nM of each primer in the library (outer forward and reverse primers), and DNA template.

[0543] Thermocycling parameters include 95° C. for 10 minutes; 25 cycles of 96° C. for 30 seconds, 65° C. for 1 minute, 58° C. for 6 minutes, 60° C. for 8 minutes, 65° C. for 4 minutes, and 72° C. for 30 seconds; and then 72° C. for 2 minutes, and then a 4° C. hold. Next, 2 μl of the resulting product, diluted 1:200, is as input in a second PCR reaction. This reaction uses a 10 μl reaction volume with 1 $\times$ QIAGEN MM final concentration, 20 nM of each inner forward primer, and 1 μM of reverse primer tag. Thermocycling parameters include 95° C. for 10 minutes; 15 cycles of 95° C. for 30 seconds, 65° C. for 1 minute, 60° C. for 5 minutes, 65° C. for 5 minutes, and 72° C. for 30 seconds; and then 72° C. for 2 minutes, and then a 4° C. hold.

[0544] Any of the methods disclosed herein or any standard methods can be used to test a primer library to determine, e.g., the percent of primer dimers, percent of target amplicons, and percent of target loci that are amplified. In some embodiments, the PCR products are sequenced as described in Example 15 or using standard sequencing methods. In some embodiments, the percentage of primer dimers can be determined by measuring the number of sequencing reads from primer dimers, the percentage of amplified products that are target amplicons can be determined by measuring the number of sequencing reads that map to target loci; the percent of target loci that are amplified can be determined by measuring the number of target loci for which there are sequencing reads that map to the target loci; the number of copies of a particular amplified target loci can be determined based on the number of sequencing reads that map to that target loci (such as by comparing the number of sequencing reads compared to the sequences reads from a standard of known concentration or amount).

[0545] FIG. 45 contains data (such as percent mapped reads and error rate) from multiplex PCR with various buffers. In this figure, "1 $\times$ MM" denotes 1 $\times$ QIAGEN Master Mix (the recommended concentration) discussed above, and "2 $\times$ MM" denotes 2 $\times$ QIAGEN Master Mix (twice the recommended concentration). FIG. 45 also lists the components of buffer F-A (also called F-A Gold), F-B (also called F-B Gold), F-D, and F-J (also called F-B Qiagen or F-B Qia) as well as the amount and type of polymerase used to generate the data. FIG. 46 is a graph illustrating the uniformity in DOR for multiplex PCR with buffers from FIG. 45. FIG. 47 is a graph illustrating the normalized depth of read (DOR) for multiplex PCR with buffers from FIG. 45 with the DOR normalized to that of buffer 2 $\times$ MM.

Limit of Detection

[0546] As demonstrated by experiments provided in the Examples section, methods provided herein are capable of detecting an average allelic imbalance in a sample with a limit of detection or sensitivity of 0.45% AAI, which is the limit of detection for aneuploidy of an illustrative method of the present invention. Similarly, in certain embodiments, methods provided herein are capable of detecting an average allelic imbalance in a sample of 0.45, 0.5, 0.6, 0.8, 0.8, 0.9,

or 1.0%. That is, the test method is capable of detecting chromosomal aneuploidy in a sample down to an AAI of 0.45, 0.5, 0.6, 0.8, 0.8, 0.9, or 1.0%. As demonstrated by experiments provided in the Examples section, methods provided herein are capable of detecting the presence of an SNV in a sample for at least some SNVs, with a limit of detection or sensitivity of 0.2%, which is the limit of detection for at least some SNVs in one illustrative embodiment. Similarly, in certain embodiments, the method is capable of detecting an SNV with a frequency or SNV AAI of 0.2, 0.3, 0.4, 0.5, 0.6, 0.8, 0.8, 0.9, or 1.0%. That is, the test method is capable of detecting an SNV in a sample down to a limit of detection of 0.2, 0.3, 0.4, 0.5, 0.6, 0.8, 0.9, or 1.0% of the total allele counts at the chromosomal locus of the SNV.

[0547] In some embodiments, a limit of detection of a mutation (such as an SNV or CNV) of a method of the invention is less than or equal to 10, 5, 2, 1, 0.5, 0.1, 0.05, 0.01, or 0.005%. In some embodiments, a limit of detection of a mutation (such as an SNV or CNV) of a method of the invention is between 15 to 0.005%, such as between 10 to 0.005%, 10 to 0.01%, 10 to 0.1%, 5 to 0.005%, 5 to 0.01%, 5 to 0.1%, 1 to 0.005%, 1 to 0.01%, 1 to 0.1%, 0.5 to 0.005%, 0.5 to 0.01%, 0.5 to 0.1%, or 0.1 to 0.01, inclusive.

[0548] In some embodiments, a limit of detection is such that a mutation (such as an SNV or CNV) that is present in less than or equal to 10, 5, 2, 1, 0.5, 0.1, 0.05, 0.01, or 0.005% of the DNA or RNA molecules with that locus in a sample (such as a sample of cfDNA or cfRNA) is detected (or is capable of being detected). For example, the mutation can be detected even if less than or equal to 10, 5, 2, 1, 0.5, 0.1, 0.05, 0.01, or 0.005% of the DNA or RNA molecules that have that locus have that mutation in the locus (instead of, for example, a wild-type or non-mutated version of the locus or a different mutation at that locus). In some embodiments, a limit of detection is such that a mutation (such as an SNV or CNV) that is present in less than or equal to 10, 5, 2, 1, 0.5, 0.1, 0.05, 0.01, or 0.005% of the DNA or RNA molecules in a sample (such as a sample of cfDNA or cfRNA) is detected (or is capable of being detected). In some embodiments in which the CNV is a deletion, the deletion can be detected even if it is only present in less than or equal to 10, 5, 2, 1, 0.5, 0.1, 0.05, 0.01, or 0.005% of the DNA or RNA molecules that have a region of interest that may or may not contain the deletion in a sample. In some embodiments in which the CNV is a deletion, the deletion can be detected even if it is only present in less than or equal to 10, 5, 2, 1, 0.5, 0.1, 0.05, 0.01, or 0.005% of the DNA or RNA molecules in a sample. In some embodiments in which the CNV is a duplication, the duplication can be detected even if the extra duplicated DNA or RNA that is present is less than or equal to 10, 5, 2, 1, 0.5, 0.1, 0.05, 0.01, or 0.005% of the DNA or RNA molecules that have a region of interest that may or may not be duplicated in a sample in a sample. In some embodiments in which the CNV is a duplication, the duplication can be detected even if the extra duplicated DNA or RNA that is present is less than or equal to 10, 5, 2, 1, 0.5, 0.1, 0.05, 0.01, or 0.005% of the DNA or RNA molecules in a sample. Experiment 23 provides exemplary methods for calculating the limit of detection. In some embodiments, the "LOD-zs5.0-mr5" method of Example 23 is used.

Exemplary Nucleic Acid Samples

[0549] In some embodiments, the genetic sample may be prepared and/or purified. There are a number of standard procedures known in the art to accomplish such an end. In some embodiments, the sample may be centrifuged to separate various layers. In some embodiments, the DNA may be isolated using filtration. In some embodiments, the preparation of the DNA may involve amplification, separation, purification by chromatography, liquid separation, isolation, preferential enrichment, preferential amplification, targeted amplification, or any of a number of other techniques either known in the art or described herein.

[0550] In some embodiments, a method disclosed herein could be used in situations where there is a very small amount of DNA present, such as in *in vitro* fertilization, or in forensic situations, where one or a few cells are available (typically less than ten cells, less than twenty cells or less than 40 cells.) In these embodiments, a method disclosed herein serves to make ploidy calls from a small amount of DNA that is not contaminated by other DNA, but where the ploidy calling very difficult the small amount of DNA. In some embodiments, a method disclosed herein could be used in situations where the target DNA is contaminated with DNA of another individual, for example in maternal blood in the context of prenatal diagnosis, paternity testing, or products of conception testing. Some other situations where these methods would be particularly advantageous would be in the case of cancer testing where only one or a small number of cells were present among a larger amount of normal cells. The genetic measurements used as part of these methods could be made on any sample comprising DNA or RNA, for example but not limited to: blood, plasma, body fluids, urine, hair, tears, saliva, tissue, skin, fingernails, blastomeres, embryos, fetal cells, amniotic fluid, chorionic villus samples, feces, bile, lymph, cervical mucus, semen, or other cells or materials comprising nucleic acids. In an embodiment, a method disclosed herein could be run with nucleic acid detection methods such as sequencing, microarrays, qPCR, digital PCR, or other methods used to measure nucleic acids. If for some reason it were found to be desirable, the ratios of the allele count probabilities at a locus could be calculated, and the allele ratios could be used to determine ploidy state in combination with some of the methods described herein, provided the methods are compatible. In some embodiments, a method disclosed herein involves calculating, on a computer, allele ratios at the plurality of polymorphic loci from the DNA measurements made on the processed samples. In some embodiments, a method disclosed herein involves calculating, on a computer, allele ratios at the plurality of polymorphic loci from the DNA measurements made on the processed samples along with any combination of other improvements described in this disclosure. Exemplary methods for isolating fetal cells, such as a single fetal cell are disclosed in U.S. Ser. No. 61/978,648, filed Apr. 11, 2014 and U.S. Ser. No. 61/984,546, filed Apr. 25, 2014. Fetal cells or fetal nucleic acids can be isolated from a pregnant mother using invasive (such as CVS or amniocentesis) or noninvasive methods (such as from a maternal blood sample).

[0551] In some embodiments, this method may be used to genotype a single cell, a small number of cells, two to five cells, six to ten cells, ten to twenty cells, twenty to fifty cells, fifty to one hundred cells, one hundred to one thousand cells, or a small amount of extracellular DNA, for example from

one to ten picograms, from ten to one hundred picograms, from one hundred picograms to one nanogram, from one to ten nanograms, from ten to one hundred nanograms, from 30 to 500 nanograms, or from one hundred nanograms to one microgram. In some embodiments, nucleic acids (such as DNA and/or RNA) from less than 100, 75, 50, 40, 30, 20, 10, 8, 6, 4, 2, or 1 cell is amplified with any of the methods of the invention. In some embodiments, the nucleic acid sample includes less than 80, 60, 40, 20, or 10% of the nucleic acids (such as DNA and/or RNA) from a single cell. In some embodiments, in which a small number of cells (such as one cell) or a small amount of nucleic acids is used, nested PCR such as hemi-nested or semi-nested PCR is used and/or the number of PCR cycles is increased compare to that used for samples with a larger amount of cells or nucleic acids. In some embodiments, a large amount of cells or nucleic acids are used (such as in cases in which a larger amount is desired to improve performance of any of the methods of the invention. In some embodiments, a sample with at least 2, 5, 10, 15, 20, 30, 50, 100, or more cells (or DNA or RNA from such cells) is used in any of the methods of the invention. In some embodiments, at least 0.5, 1, 10, 25, 50, 100, 500, 1,000; or 5,000 ng of DNA or RNA is used.

[0552] In some embodiments, the cells in the sample are lysed prior to PCR. In some embodiments, the Arcturus PicoPure DNA extraction kit from Applied Biosystems is used. (Applied Biosystems cat. No. KIT0103, see, e.g., information available at the World Wide Web at lifetech-nologies.com/order/catalog/product/KIT0103?ICID=search-product, which is hereby incorporated by reference in its entirety). This kit contains Arcturus reconstitution buffer and Protease K. In some embodiments, the following cell lysis thermocycling protocol is used: 56° C. for 1 hour, 95° C. for 10 minutes, 25° C. for 15 minutes, and then a 4° C. hold.

[0553] In some embodiments, the nucleic acids are processed using the consecutive steps of end-repairing, dA-tailing, and adaptor ligating the nucleic acids. The consecutive steps exclude purifying the end-repaired products prior to the dA-tailing step and exclude purifying the dA-tailing products prior to the adaptor ligating step. The resulting products are amplified in any of the multiplex PCR methods of the invention. In some embodiments, the amplified products are then sequenced.

Exemplary Nucleic Acid Studies

[0554] The multiplex PCR methods of the invention can be used to increase the number of target loci that can be evaluated to measure the amount of one or more specific nucleic acid molecules of interest or of one or more types of nucleic acids. In some embodiments, there is a change in the total amount or concentration of one or more types of DNA (such as cfDNA cf mDNA, cf nDNA, cellular DNA, or mitochondrial DNA) or RNA (cfRNA, cellular RNA, cytoplasmic RNA, coding cytoplasmic RNA, non-coding cytoplasmic RNA, mRNA, miRNA, mitochondrial RNA, rRNA, or tRNA). In some embodiments, there is a change in the amount or concentration of one or more specific DNA (such as cfDNA cf mDNA, cf nDNA, cellular DNA, or mitochondrial DNA) or RNA (cfRNA, cellular RNA, cytoplasmic RNA, coding cytoplasmic RNA, non-coding cytoplasmic RNA, mRNA, miRNA, mitochondrial RNA, rRNA, or tRNA) molecules. In some embodiments, one allele is expressed more than another allele of a locus of interest.

Exemplary miRNAs are short 20-22 nucleotide RNA molecules that regulate the expression of a gene. In some embodiments, there is a change in the transcriptome, such as a change in the identity or amount of one or more RNA molecules.

[0555] In some embodiments, an increase in the total amount or concentration of cfDNA or cfRNA is associated with a disease such as cancer, or an increased risk for a disease such as cancer. In some embodiments, the total concentration of a type of DNA (such as cfDNA cf mDNA, cf nDNA, cellular DNA, or mitochondrial DNA) or RNA (cfRNA, cellular RNA, cytoplasmic RNA, coding cytoplasmic RNA, non-coding cytoplasmic RNA, mRNA, miRNA, mitochondrial RNA, rRNA, or tRNA) increases by at least 2, 3, 4, 5, 6, 7, 8, 9, 10-fold, or more compared to the total concentration of that type of DNA or RNA in healthy (such as non-cancerous) subjects. In some embodiments, a total concentration of cfDNA between 75 to 100 ng/mL, 100 to 150 ng/mL, 150 to 200 ng/mL, 200 to 300 ng/mL, 300 to 400 ng/mL, 400 to 600 ng/mL, 600 to 800 ng/mL, 800 to 1,000 ng/mL, inclusive, or a total concentration of cfDNA of more than 100 ng, mL, such as more than 200, 300, 400, 500, 600, 700, 800, 900, or 1,000 ng/mL is indicative of cancer, an increased risk for cancer, an increased risk of a tumor being malignant rather than benign, a decreased probability of the cancer going into remission, or a worse prognosis for the cancer. In some embodiments, the amount of a type of DNA (such as cfDNA cf mDNA, cf nDNA, cellular DNA, or mitochondrial DNA) or RNA (cfRNA, cellular RNA, cytoplasmic RNA, coding cytoplasmic RNA, non-coding cytoplasmic RNA, mRNA, miRNA, mitochondrial RNA, rRNA, or tRNA) having one or more polymorphisms/mutations (such as deletions or duplications) associated with a disease such as cancer or an increased risk for a disease such as cancer is at least 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 14, 16, 18, 20, or 25% of the total amount of that type of DNA or RNA. In some embodiments, at least 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 14, 16, 18, 20, or 25% of the total amount of a type of DNA (such as cfDNA cf mDNA, cf nDNA, cellular DNA, or mitochondrial DNA) or RNA (cfRNA, cellular RNA, cytoplasmic RNA, coding cytoplasmic RNA, non-coding cytoplasmic RNA, mRNA, miRNA, mitochondrial RNA, rRNA, or tRNA) has a particular polymorphism or mutation (such as a deletion or duplication) associated with a disease such as cancer or an increased risk for a disease such as cancer.

Exemplary RNA Expression Studies

[0556] The multiplex PCR methods of the invention can be used to increase the number of target loci that can be evaluated during gene expression profiling experiments. For example, the expression levels of thousands of genes can be simultaneously monitored to determine whether a person has a sequence (such as a polymorphism or other mutation) associated with a disease (such as cancer) or an increased risk of a disease. These methods can be used to identify sequences (such as polymorphisms or other mutations) associated with an increased or decreased risk for a disease such as cancer by comparing gene expression (such as the expression of particular mRNA alleles) in samples from patients with and without the disease. Additionally, the effect of particular treatments, diseases, or developmental stages on gene expression can be determined. Similarly, these methods can be used to identify genes whose expression is

changed in response to pathogens or other organisms by comparing gene expression in infected and uninfected cells or tissues. In these methods the number of sequencing reads can be adjusted based on the frequency of the polymorphisms that are being analyzed such that sufficient reads are performed for the polymorphisms to be detected if they are present. In some embodiments, the polymorphisms or mutation is present at a higher frequency in subjects with a disease or disorder (such as cancer) than subjects without the disease or disorder (such as cancer). In some embodiments, the polymorphisms or mutation is indicative of cancer, such as a causative mutation.

[0557] In some embodiments, a sample containing RNA (such as mRNA) is amplified using a reverse transcriptase (RT) and the resulting DNA (such as cDNA) is then amplified using a DNA polymerase (PCR). The RT and PCR steps may be carried out sequentially in the same reaction volume or separately. Any of the primer libraries of the invention can be used in this reverse transcription polymerase chain reaction (RT-PCR) method. In various embodiments, the reverse transcription is performed using oligo-dT, random primers, a mixture of oligo-dT and random primers, or primers specific to the target loci. To avoid amplification of contaminating genomic DNA, primers for RT-PCR can be designed so that part of one primer hybridizes to the 3' end of one exon and the other part of the primer hybridizes to the 5' end of the adjacent exon. Such primers anneal to cDNA synthesized from spliced mRNAs, but not to genomic DNA. To detect amplification of contaminating DNA, RT-PCR primer pairs may be designed to flank a region that contains at least one intron. Products amplified from cDNA (no introns) are smaller than those amplified from genomic DNA (containing introns). Size difference in products is used to detect the presence of contaminating DNA. In some embodiments when only the mRNA sequence is known, primer annealing sites are chosen that are at least 300-400 base pairs apart since it is likely that fragments of this size from eukaryotic DNA contain splice junctions. Alternatively, the sample can be treated with DNase to degrade contaminating DNA.

Exemplary Methods for Paternity Testing

[0558] The multiplex PCR methods of the invention can be used to improve the accuracy of paternity testing since so many target loci can be analyzed at once (see, e.g., U.S. Publication No. 2012/0122701, filed Dec. 22, 2011, in which is hereby incorporated by reference in its entirety). For example, the multiplex PCR method can allow thousands of polymorphic loci (such as SNPs) to be analyzed for use in the PARENTAL SUPPORT algorithm described herein to determine whether an alleged father is the biological father of a fetus. In some embodiments the method involves (i) simultaneously amplifying a plurality of polymorphic loci that includes at least 25; 50; 75; 100; 300; 500; 750; 1,000; 2,000; 5,000; 7,500; 10,000; 15,000; 19,000; 20,000; 25,000; 27,000; 28,000; 30,000; 40,000; 50,000; 75,000; or 100,000 different polymorphic loci on genetic material from the alleged father to produce a first set of amplified products; (ii) simultaneously amplifying the corresponding plurality of polymorphic loci on a mixed sample of DNA originating from a blood sample from the pregnant mother to produce a second set of amplified products; wherein the mixed sample of DNA comprises fetal DNA and maternal DNA; (iii) determining on a computer the probability that the alleged

father is the biological father of the fetus using genotypic measurements based on the first and second sets of amplified products; and (iv) establishing whether the alleged father is the biological father of the fetus using the determined probability that the alleged father is the biological father of the fetus. In various embodiments, the method further includes simultaneously amplifying the corresponding plurality of polymorphic loci on genetic material from the mother to produce a third set of amplified products; wherein the probability that the alleged father is the biological father of the fetus is determined using genotypic measurements based on the first, second, and third sets of amplified products.

Exemplary Methods for Embryo Characterization and Selection

[0559] The multiplex PCR methods of the invention can be used to improve the selection of embryos for in vitro fertilization by allowing thousands of target loci to be analyzed at once (see, e.g., U.S. Pub. No. 2011/0092763, filed May 27, 2008, filed Dec. 22, 2011, is which is hereby incorporated by reference in its entirety). For example, the multiplex PCR method can allow thousands of polymorphic loci (such as SNPs) to be analyzed for use in the PARENTAL SUPPORT algorithm described herein to select an embryo out of a set of embryos for in vitro fertilization

[0560] In some embodiments, the invention provides methods of estimating relative likelihoods that each embryo from a set of embryos will develop as desired. In some embodiments, the method involves contacting a sample from each embryo with a library of primers that simultaneously hybridize to at least 25; 50; 75; 100; 300; 500; 750; 1,000; 2,000; 5,000; 7,500; 10,000; 15,000; 19,000; 20,000; 25,000; 27,000; 28,000; 30,000; 40,000; 50,000; 75,000; or 100,000 different target loci to produce a reaction mixture for each embryo, wherein the samples are each derived from one or more cells from an embryo. In some embodiments, each reaction mixture is subjected to primer extension reaction conditions to produce amplified products. In some embodiments, the method includes determining on a computer one or more characteristics of at least one cell from each embryo based on the amplified products; and estimating on a computer the relative likelihoods that each embryo will develop as desired, based on the one or more characteristics of the at least one cell for each embryo. In some embodiments, the method includes using an informatics based method to determine the at least one characteristic, such as the PARENTAL SUPPORT algorithm described herein. In some embodiments, the characteristic includes a ploidy state. In some embodiments, the characteristic is selected from the group consisting of aneuploid, euploid, mosaic, nullsomy, monosomy, uniparental disomy, trisomy, tetrasomy, a type of aneuploidy, unmatched copy error trisomy, matched copy error trisomy, maternal origin of aneuploidy, paternal origin of aneuploidy, a presence or absence of a disease-linked gene, a chromosomal identity of any aneuploid chromosome, an abnormal genetic condition, a deletion or duplication, a likelihood of a characteristic, and combinations thereof. The characteristic may be associated with a chromosome taken from the group consisting of chromosome one, chromosome two, chromosome three, chromosome four, chromosome five, chromosome six, chromosome seven, chromosome eight, chromosome nine, chromosome ten, chromosome eleven, chromosome twelve,

chromosome thirteen, chromosome fourteen, chromosome fifteen, chromosome sixteen, chromosome seventeen, chromosome eighteen, chromosome nineteen, chromosome twenty, chromosome twenty-one, chromosome twenty-two, X chromosome or Y chromosome, and combinations thereof.

Exemplary Prenatal Diagnostic Methods

[0561] The multiplex PCR methods of the present invention can be used to improve prenatal diagnostic methods, such as the determination of the ploidy status of fetal chromosomes. Given that the large number of target loci that can be simultaneously amplified, more accurate determinations can be made.

[0562] In an embodiment, the present disclosure provides ex vivo methods for determining the ploidy status of a chromosome in a gestating fetus from genotypic data measured from a mixed sample of DNA (i.e., DNA from the mother of the fetus, and DNA from the fetus) and optionally from genotypic data measured from a sample of genetic material from the mother and possibly also from the father, wherein the determining is done by using a joint distribution model to create a set of expected allele distributions for different possible fetal ploidy states given the parental genotypic data, and comparing the expected allelic distributions to the actual allelic distributions measured in the mixed sample, and choosing the ploidy state whose expected allelic distribution pattern most closely matches the observed allelic distribution pattern. In an embodiment, the mixed sample is derived from maternal blood, or maternal serum or plasma. In an embodiment, the mixed sample of DNA may be preferentially enriched at a target loci (e.g., plurality of polymorphic loci). In an embodiment, the preferential enrichment is done in a way that minimizes the allelic bias. In an embodiment, the present disclosure relates to a composition of DNA that has been preferentially enriched at a plurality of loci such that the allelic bias is low. In an embodiment, the allelic distribution(s) are measured by sequencing the DNA from the mixed sample. In an embodiment, the joint distribution model assumes that the alleles will be distributed in a binomial fashion. In an embodiment, the set of expected joint allele distributions are created for genetically linked loci while considering the extant recombination frequencies from various sources, for example, using data from the International HapMap Consortium.

[0563] In an embodiment, the present disclosure provides methods for non-invasive prenatal diagnosis (NPD), specifically, determining the aneuploidy status of a fetus by observing allele measurements at a plurality of polymorphic loci in genotypic data measured on DNA mixtures, where certain allele measurements are indicative of an aneuploid fetus, while other allele measurements are indicative of a euploid fetus. In an embodiment, the genotypic data is measured by sequencing DNA mixtures that were derived from maternal plasma. In an embodiment, the DNA sample may be preferentially enriched in molecules of DNA that correspond to the plurality of loci whose allele distributions are being calculated. In an embodiment a sample of DNA comprising only or almost only genetic material from the mother and possibly also a sample of DNA comprising only or almost only genetic material from the father are measured. In an embodiment, the genetic measurements of one or both parents along with the estimated fetal fraction are used to

create a plurality of expected allele distributions corresponding to different possible underlying genetic states of the fetus; the expected allele distributions may be termed hypotheses. In an embodiment, the maternal genetic data is not determined by measuring genetic material that is exclusively or almost exclusively maternal in nature, rather, it is estimated from the genetic measurements made on maternal plasma that comprises a mixture of maternal and fetal DNA. In some embodiments the hypotheses may comprise the ploidy of the fetus at one or more chromosomes, which segments of which chromosomes in the fetus were inherited from which parents, and combinations thereof. In some embodiments, the ploidy state of the fetus is determined by comparing the observed allele measurements to the different hypotheses where at least some of the hypotheses correspond to different ploidy states, and selecting the ploidy state that corresponds to the hypothesis that is most likely to be true given the observed allele measurements. In an embodiment, this method involves using allele measurement data from some or all measured SNPs, regardless of whether the loci are homozygous or heterozygous, and therefore does not involve using alleles at loci that are only heterozygous. This method may not be appropriate for situations where the genetic data pertains to only one polymorphic locus. This method is particularly advantageous when the genetic data comprises data for more than ten polymorphic loci for a target chromosome or more than twenty polymorphic loci. This method is especially advantageous when the genetic data comprises data for more than 50 polymorphic loci for a target chromosome, more than 100 polymorphic loci or more than 200 polymorphic loci for a target chromosome. In some embodiments, the genetic data may comprise data for more than 500 polymorphic loci for a target chromosome, more than 1,000 polymorphic loci, more than 2,000 polymorphic loci, or more than 5,000 polymorphic loci for a target chromosome.

[0564] In an embodiment, a method disclosed herein yields a quantitative measure of the number of independent observations of each allele at a polymorphic locus. This is unlike most methods such as microarrays or qualitative PCR which provide information about the ratio of two alleles but do not quantify the number of independent observations of either allele. With methods that provide quantitative information regarding the number of independent observations, only the ratio is utilized in ploidy calculations, while the quantitative information by itself is not useful. To illustrate the importance of retaining information about the number of independent observations consider the sample locus with two alleles, A and B. In a first experiment twenty A alleles and twenty B alleles are observed, in a second experiment 200 A alleles and 200 B alleles are observed. In both experiments the ratio ($A/(A+B)$) is equal to 0.5, however the second experiment conveys more information than the first about the certainty of the frequency of the A or B allele. Some methods by others involve averaging or summing allele ratios (channel ratios) (i.e. x_i/y_i) from individual allele and analyzes this ratio, either comparing it to a reference chromosome or using a rule pertaining to how this ratio is expected to behave in particular situations. No allele weighting is implied in such methods, where it is assumed that one can ensure about the same amount of PCR product for each allele and that all the alleles should behave the same way. Such a method has a number of disadvantages, and more

importantly, precludes the use a number of improvements that are described elsewhere in this disclosure.

[0565] In an embodiment, a method disclosed herein explicitly models the allele frequency distributions expected in disomy as well as a plurality of allele frequency distributions that may be expected in cases of trisomy resulting from nondisjunction during meiosis I, nondisjunction during meiosis II, and/or nondisjunction during mitosis early in fetal development. To illustrate why this is important, imagine a case where there were no crossovers: nondisjunction during meiosis I would result a trisomy in which two different homologs were inherited from one parent; in contrast, nondisjunction during meiosis II or during mitosis early in fetal development would result in two copies of the same homolog from one parent. Each scenario would result in different expected allele frequencies at each polymorphic locus and also at all loci considered jointly, due to genetic linkage. Crossovers, which result in the exchange of genetic material between homologs, make the inheritance pattern more complex; in an embodiment, the instant method accommodates for this by using recombination rate information in addition to the physical distance between loci. In an embodiment, to enable improved distinction between meiosis I nondisjunction and meiosis II or mitotic nondisjunction the instant method incorporate into the model an increasing probability of crossover as the distance from the centromere increases. Meiosis II and mitotic nondisjunction can distinguished by the fact that mitotic nondisjunction typically results in identical or nearly identical copies of one homolog while the two homologs present following a meiosis II nondisjunction event often differ due to one or more crossovers during gametogenesis.

[0566] In some embodiments, a method disclosed herein involves comparing the observed allele measurements to theoretical hypotheses corresponding to possible fetal genetic aneuploidy, and does not involve a step of quantitating a ratio of alleles at a heterozygous locus. Where the number of loci is lower than about 20, the ploidy determination made using a method comprising quantitating a ratio of alleles at a heterozygous locus and a ploidy determination made using a method comprising comparing the observed allele measurements to theoretical allele distribution hypotheses corresponding to possible fetal genetic states may give a similar result. However, where the number of loci is above 50 these two methods is likely to give significantly different results; where the number of loci is above 400, above, 1,000 or above 2,000 these two methods are very likely to give results that are increasingly significantly different. These differences are due to the fact that a method that comprises quantitating a ratio of alleles at a heterozygous locus without measuring the magnitude of each allele independently and aggregating or averaging the ratios precludes the use of techniques including using a joint distribution model, performing a linkage analysis, using a binomial distribution model, and/or other advanced statistical techniques, whereas using a method comprising comparing the observed allele measurements to theoretical allele distribution hypotheses corresponding to possible fetal genetic states may use these techniques which can substantially increase the accuracy of the determination.

[0567] In an embodiment, a method disclosed herein involves determining whether the distribution of observed allele measurements is indicative of a euploid or an aneuploid fetus using a joint distribution model. The use of a joint

distribution model is a different from and a significant improvement over methods that determine heterozygosity rates by treating polymorphic loci independently in that the resultant determinations are of significantly higher accuracy. Without being bound by any particular theory, it is believed that one reason they are of higher accuracy is that the joint distribution model takes into account the linkage between SNPs, and likelihood of crossovers having occurred during the meiosis that gave rise to the gametes that formed the embryo that grew into the fetus. The purpose of using the concept of linkage when creating the expected distribution of allele measurements for one or more hypotheses is that it allows the creation of expected allele measurements distributions that correspond to reality considerably better than when linkage is not used. For example, imagine that there are two SNPs, 1 and 2 located nearby one another, and the mother is A at SNP 1 and A at SNP 2 on one homolog, and B at SNP 1 and B at SNP 2 on homolog two. If the father is A for both SNPs on both homologs, and a B is measured for the fetus SNP 1, this indicates that homolog two has been inherited by the fetus, and therefore that there is a much higher likelihood of a B being present on the fetus at SNP 2. A model that takes into account linkage would predict this, while a model that does not take linkage into account would not. Alternately, if a mother was AB at SNP 1 and AB at nearby SNP 2, then two hypotheses corresponding to maternal trisomy at that location could be used—one involving a matching copy error (nondisjunction in meiosis II or mitosis in early fetal development), and one involving an unmatched copy error (nondisjunction in meiosis I). In the case of a matching copy error trisomy, if the fetus inherited an AA from the mother at SNP 1, then the fetus is much more likely to inherit either an AA or BB from the mother at SNP 2, but not AB. In the case of an unmatched copy error, the fetus would inherit an AB from the mother at both SNPs. The allele distribution hypotheses made by a ploidy calling method that takes into account linkage would make these predictions, and therefore correspond to the actual allele measurements to a considerably greater extent than a ploidy calling method that did not take into account linkage. Note that a linkage approach is not possible when using a method that relies on calculating allele ratios and aggregating those allele ratios.

[0568] One reason that it is believed that ploidy determinations that use a method that comprises comparing the observed allele measurements to theoretical hypotheses corresponding to possible fetal genetic states are of higher accuracy is that when sequencing is used to measure the alleles, this method can glean more information from data from alleles where the total number of reads is low than other methods; for example, a method that relies on calculating and aggregating allele ratios would produce disproportionately weighted stochastic noise. For example, imagine a case that involved measuring the alleles using sequencing, and where there was a set of loci where only five sequence reads were detected for each locus. In an embodiment, for each of the alleles, the data may be compared to the hypothesized allele distribution, and weighted according to the number of sequence reads; therefore the data from these measurements would be appropriately weighted and incorporated into the overall determination. This is in contrast to a method that involved quantitating a ratio of alleles at a heterozygous locus, as this method could only calculate ratios of 0%, 20%, 40%, 60%, 80% or 100% as the possible

allele ratios; none of these may be close to expected allele ratios. In this latter case, the calculated allele ratios would either have to be discarded due to insufficient reads or else would have disproportionate weighting and introduce stochastic noise into the determination, thereby decreasing the accuracy of the determination. In an embodiment, the individual allele measurements may be treated as independent measurements, where the relationship between measurements made on alleles at the same locus is no different from the relationship between measurements made on alleles at different loci.

[0569] In an embodiment, a method disclosed herein involves determining whether the distribution of observed allele measurements is indicative of a euploid or an aneuploid fetus without comparing any metrics to observed allele measurements on a reference chromosome that is expected to be disomic (termed the RC method). This is a significant improvement over methods, such as methods using shotgun sequencing which detect aneuploidy by evaluating the proportion of randomly sequenced fragments from a suspect chromosomes relative to one or more presumed disomic reference chromosome. This RC method yields incorrect results if the presumed disomic reference chromosome is not actually disomic. This can occur in cases where aneuploidy is more substantial than trisomy of a single chromosome or where the fetus is triploid and all autosomes are trisomic. In the case of a female triploid (69, XXX) fetus there are in fact no disomic chromosomes at all. The method described herein does not require a reference chromosome and would be able to correctly identify trisomic chromosomes in a female triploid fetus. For each chromosome, hypothesis, child fraction and noise level, a joint distribution model may be fit, without any of: reference chromosome data, an overall child fraction estimate, or a fixed reference hypothesis.

[0570] In an embodiment, a method disclosed herein demonstrates how observing allele distributions at polymorphic loci can be used to determine the ploidy state of a fetus with greater accuracy than methods in the prior art. In an embodiment, the method uses the targeted sequencing to obtain mixed maternal-fetal genotypes and optionally mother and/or father genotypes at a plurality of SNPs to first establish the various expected allele frequency distributions under the different hypotheses, and then observing the quantitative allele information obtained on the maternal-fetal mixture and evaluating which hypothesis fits the data best, where the genetic state corresponding to the hypothesis with the best fit to the data is called as the correct genetic state. In an embodiment, a method disclosed herein also uses the degree of fit to generate a confidence that the called genetic state is the correct genetic state. In an embodiment, a method disclosed herein involves using algorithms that analyze the distribution of alleles found for loci that have different parental contexts, and comparing the observed allele distributions to the expected allele distributions for different ploidy states for the different parental contexts (different parental genotypic patterns). This is different from and an improvement over methods that do not use methods that enable the estimation of the number of independent instances of each allele at each locus in a mixed maternal-fetal sample. In an embodiment, a method disclosed herein involves determining whether the distribution of observed allele measurements is indicative of a euploid or an aneuploid fetus using observed allelic distributions measured at loci where the mother is heterozygous. This is different from

and an improvement over methods that do not use observed allelic distributions at loci where the mother is heterozygous because, in cases where the DNA is not preferentially enriched or is preferentially enriched for loci that are not known to be highly informative for that particular target individual, it allows the use of about twice as much genetic measurement data from a set of sequence data in the ploidy determination, resulting in a more accurate determination.

[0571] In an embodiment, a method disclosed herein uses a joint distribution model that assumes that the allele frequencies at each locus are multinomial (and thus binomial when SNPs are biallelic) in nature. In some embodiments the joint distribution model uses beta-binomial distributions. When using a measuring technique, such as sequencing, provides a quantitative measure for each allele present at each locus, binomial model can be applied to each locus and the degree underlying allele frequencies and the confidence in that frequency can be ascertained. With methods known in the art that generate ploidy calls from allele ratios, or methods in which quantitative allele information is discarded, the certainty in the observed ratio cannot be ascertained. The instant method is different from and an improvement over methods that calculate allele ratios and aggregate those ratios to make a ploidy call, since any method that involves calculating an allele ratio at a particular locus, and then aggregating those ratios, necessarily assumes that the measured intensities or counts that are indicative of the amount of DNA from any given allele or locus will be distributed in a Gaussian fashion. The method disclosed herein does not involve calculating allele ratios. In some embodiments, a method disclosed herein may involve incorporating the number of observations of each allele at a plurality of loci into a model. In some embodiments, a method disclosed herein may involve calculating the expected distributions themselves, allowing the use of a joint binomial distribution model which may be more accurate than any model that assumes a Gaussian distribution of allele measurements. The likelihood that the binomial distribution model is significantly more accurate than the Gaussian distribution increases as the number of loci increases. For example, when fewer than 20 loci are interrogated, the likelihood that the binomial distribution model is significantly better is low. However, when more than 100, or especially more than 400, or especially more than 1,000, or especially more than 2,000 loci are used, the binomial distribution model will have a very high likelihood of being significantly more accurate than the Gaussian distribution model, thereby resulting in a more accurate ploidy determination. The likelihood that the binomial distribution model is significantly more accurate than the Gaussian distribution also increases as the number of observations at each locus increases. For example, when fewer than 10 distinct sequences are observed at each locus are observed, the likelihood that the binomial distribution model is significantly better is low. However, when more than 50 sequence reads, or especially more than 100 sequence reads, or especially more than 200 sequence reads, or especially more than 300 sequence reads are used for each locus, the binomial distribution model will have a very high likelihood of being significantly more accurate than the Gaussian distribution model, thereby resulting in a more accurate ploidy determination.

[0572] In an embodiment, a method disclosed herein uses sequencing to measure the number of instances of each

allele at each locus in a DNA sample. Each sequencing read may be mapped to a specific locus and treated as a binary sequence read; alternately, the probability of the identity of the read and/or the mapping may be incorporated as part of the sequence read, resulting in a probabilistic sequence read, that is, the probable whole or fractional number of sequence reads that map to a given loci. Using the binary counts or probability of counts it is possible to use a binomial distribution for each set of measurements, allowing a confidence interval to be calculated around the number of counts. This ability to use the binomial distribution allows for more accurate ploidy estimations and more precise confidence intervals to be calculated. This is different from and an improvement over methods that use intensities to measure the amount of an allele present, for example methods that use microarrays, or methods that make measurements using fluorescence readers to measure the intensity of fluorescently tagged DNA in electrophoretic bands.

[0573] In an embodiment, a method disclosed herein uses aspects of the present set of data to determine parameters for the estimated allele frequency distribution for that set of data. This is an improvement over methods that utilize training set of data or prior sets of data to set parameters for the present expected allele frequency distributions, or possibly expected allele ratios. This is because there are different sets of conditions involved in the collection and measurement of every genetic sample, and thus a method that uses data from the instant set of data to determine the parameters for the joint distribution model that is to be used in the ploidy determination for that sample will tend to be more accurate.

[0574] In an embodiment, a method disclosed herein involves determining whether the distribution of observed allele measurements is indicative of a euploid or an aneuploid fetus using a maximum likelihood technique. The use of a maximum likelihood technique is different from and a significant improvement over methods that use single hypothesis rejection technique in that the resultant determinations will be made with significantly higher accuracy. One reason is that single hypothesis rejection techniques set cut off thresholds based on only one measurement distribution rather than two, meaning that the thresholds are usually not optimal. Another reason is that the maximum likelihood technique allows the optimization of the cut off threshold for each individual sample instead of determining a cut off threshold to be used for all samples regardless of the particular characteristics of each individual sample. Another reason is that the use of a maximum likelihood technique allows the calculation of a confidence for each ploidy call. The ability to make a confidence calculation for each call allows a practitioner to know which calls are accurate, and which are more likely to be wrong. In some embodiments, a wide variety of methods may be combined with a maximum likelihood estimation technique to enhance the accuracy of the ploidy calls. In an embodiment, the maximum likelihood technique may be used in combination with the method described in U.S. Pat. No. 7,888,017. In an embodiment, the maximum likelihood technique may be used in combination with the method of using targeted PCR amplification to amplify the DNA in the mixed sample followed by sequencing and analysis using a read counting method such as used by TANDEM DIAGNOSTICS, as presented at the International Congress of Human Genetics 2011, in Montreal in October 2011. In an embodiment, a method

disclosed herein involves estimating the fetal fraction of DNA in the mixed sample and using that estimation to calculate both the ploidy call and the confidence of the ploidy call. Note that this is both different and distinct from methods that use estimated fetal fraction as a screen for sufficient fetal fraction, followed by a ploidy call made using a single hypothesis rejection technique that does not take into account the fetal fraction nor does it produce a confidence calculation for the call.

[0575] In an embodiment, a method disclosed herein takes into account the tendency for the data to be noisy and contain errors by attaching a probability to each measurement. The use of maximum likelihood techniques to choose the correct hypothesis from the set of hypotheses that were made using the measurement data with attached probabilistic estimates makes it more likely that the incorrect measurements will be discounted, and the correct measurements will be used in the calculations that lead to the ploidy call. To be more precise, this method systematically reduces the influence of data that is incorrectly measured on the ploidy determination. This is an improvement over methods where all data is assumed to be equally correct or methods where outlying data is arbitrarily excluded from calculations leading to a ploidy call. Existing methods using channel ratio measurements claim to extend the method to multiple SNPs by averaging individual SNP channel ratios. Not weighting individual SNPs by expected measurement variance based on the SNP quality and observed depth of read reduces the accuracy of the resulting statistic, resulting in a reduction of the accuracy of the ploidy call significantly, especially in borderline cases.

[0576] In an embodiment, a method disclosed herein does not presuppose the knowledge of which SNPs or other polymorphic loci are heterozygous on the fetus. This method allows a ploidy call to be made in cases where paternal genotypic information is not available. This is an improvement over methods where the knowledge of which SNPs are heterozygous must be known ahead of time in order to appropriately select loci to target, or to interpret the genetic measurements made on the mixed fetal/maternal DNA sample.

[0577] The methods described herein are particularly advantageous when used on samples where a small amount of DNA is available, or where the percent of fetal DNA is low. This is due to the correspondingly higher allele dropout rate that occurs when only a small amount of DNA is available and/or the correspondingly higher fetal allele dropout rate when the percent of fetal DNA is low in a mixed sample of fetal and maternal DNA. A high allele dropout rate, meaning that a large percentage of the alleles were not measured for the target individual, results in poorly accurate fetal fractions calculations, and poorly accurate ploidy determinations. Since methods disclosed herein may use a joint distribution model that takes into account the linkage in inheritance patterns between SNPs, significantly more accurate ploidy determinations may be made. The methods described herein allow for an accurate ploidy determination to be made when the percent of molecules of DNA that are fetal in the mixture is less than 40%, less than 30%, less than 20%, less than 10%, less than 8%, and even less than 6%.

[0578] In an embodiment, it is possible to determine the ploidy state of an individual based on measurements when that individual's DNA is mixed with DNA of a related individual. In an embodiment, the mixture of DNA is the

free floating DNA found in maternal plasma, which may include DNA from the mother, with known karyotype and known genotype, and which may be mixed with DNA of the fetus, with unknown karyotype and unknown genotype. It is possible to use the known genotypic information from one or both parents to predict a plurality of potential genetic states of the DNA in the mixed sample for different ploidy states, different chromosome contributions from each parent to the fetus, and optionally, different fetal DNA fractions in the mixture. Each potential composition may be referred to as a hypothesis. The ploidy state of the fetus can then be determined by looking at the actual measurements, and determining which potential compositions are most likely given the observed data.

[0579] Further discussion of the points above may be found elsewhere in this document.

Non-Invasive Prenatal Diagnosis (NPD)

[0580] The process of non-invasive prenatal diagnosis involves a number of steps. Some of the steps may include: (1) obtaining the genetic material from the fetus; (2) enriching the genetic material of the fetus that may be in a mixed sample, ex vivo; (3) amplifying the genetic material, ex vivo; (4) preferentially enriching specific loci in the genetic material, ex vivo; (5) measuring the genetic material, ex vivo; and (6) analyzing the genotypic data, on a computer, and ex vivo. Methods to reduce to practice these six and other relevant steps are described herein. At least some of the method steps are not directly applied on the body. In an embodiment, the present disclosure relates to methods of treatment and diagnosis applied to tissue and other biological materials isolated and separated from the body. At least some of the method steps are executed on a computer.

[0581] Some embodiments of the present disclosure allow a clinician to determine the genetic state of a fetus that is gestating in a mother in a non-invasive manner such that the health of the baby is not put at risk by the collection of the genetic material of the fetus, and that the mother is not required to undergo an invasive procedure. Moreover, in certain aspects, the present disclosure allows the fetal genetic state to be determined with high accuracy, significantly greater accuracy than, for example, the non-invasive maternal serum analyte based screens, such as the triple test, that are in wide use in prenatal care.

[0582] The high accuracy of the methods disclosed herein is a result of an informatics approach to analysis of the genotype data, as described herein. Modern technological advances have resulted in the ability to measure large amounts of genetic information from a genetic sample using such methods as high throughput sequencing and genotyping arrays. The methods disclosed herein allow a clinician to take greater advantage of the large amounts of data available, and make a more accurate diagnosis of the fetal genetic state. The details of a number of embodiments are given below. Different embodiments may involve different combinations of the aforementioned steps. Various combinations of the different embodiments of the different steps may be used interchangeably.

[0583] In an embodiment, a blood sample is taken from a pregnant mother, and the free floating DNA in the plasma of the mother's blood, which contains a mixture of both DNA of maternal origin, and DNA of fetal origin, is isolated and used to determine the ploidy status of the fetus. In an embodiment, a method disclosed herein involves preferen-

tial enrichment of those DNA sequences in a mixture of DNA that correspond to polymorphic alleles in a way that the allele ratios and/or allele distributions remain mostly consistent upon enrichment. In an embodiment, a method disclosed herein involves the highly efficient targeted PCR based amplification such that a very high percentage of the resulting molecules correspond to targeted loci. In an embodiment, a method disclosed herein involves sequencing a mixture of DNA that contains both DNA of maternal origin, and DNA of fetal origin. In an embodiment, a method disclosed herein involves using measured allele distributions to determine the ploidy state of a fetus that is gestating in a mother. In an embodiment, a method disclosed herein involves reporting the determined ploidy state to a clinician. In an embodiment, a method disclosed herein involves taking a clinical action, for example, performing follow up invasive testing such as chorionic villus sampling or amniocentesis, preparing for the birth of a trisomic individual or an elective termination of a trisomic fetus.

[0584] This application makes reference to U.S. Utility application Ser. No. 11/603,406, filed Nov. 28, 2006 (US Publication No.: 20070184467); U.S. Utility application Ser. No. 12/076,348, filed Mar. 17, 2008 (US Publication No.: 20080243398); PCT Application Serial No. PCT/US09/52730, filed Aug. 4, 2009 (PCT Publication No.: WO/2010/017214); PCT Application Serial No. PCT/US10/050824, filed Sep. 30, 2010 (PCT Publication No.: WO/2011/041485), U.S. Utility application Ser. No. 13/110,685, filed May 18, 2011, and PCT Application Serial No. PCT/12/58578, filed Oct. 3, 2012, which are each herein incorporated by reference in its entirety. Some of the vocabulary used in this filing may have its antecedents in these references. Some of the concepts described herein may be better understood in light of the concepts found in these references.

Screening Maternal Blood Comprising Free Floating Fetal DNA

[0585] The methods described herein may be used to help determine the genotype of a child, fetus, or other target individual where the genetic material of the target is found in the presence of a quantity of other genetic material. In some embodiments the genotype may refer to the ploidy state of one or a plurality of chromosomes, it may refer to one or a plurality of disease linked alleles, or some combination thereof. In this disclosure, the discussion focuses on determining the genetic state of a fetus where the fetal DNA is found in maternal blood, but this example is not meant to limit to possible contexts that this method may be applied to. In addition, the method may be applicable in cases where the amount of target DNA is in any proportion with the non-target DNA; for example, the target DNA could make up anywhere between 0.000001 and 99.999999% of the DNA present. In addition, the non-target DNA does not necessarily need to be from one individual, or even from a related individual, as long as genetic data from some or all of the relevant non-target individual(s) is known. In an embodiment, a method disclosed herein can be used to determine genotypic data of a fetus from maternal blood that contains fetal DNA. It may also be used in a case where there are multiple fetuses in the uterus of a pregnant woman, or where other contaminating DNA may be present in the sample, for example from other already born siblings.

[0586] This technique may make use of the phenomenon of fetal blood cells gaining access to maternal circulation

through the placental villi. Ordinarily, only a very small number of fetal cells enter the maternal circulation in this fashion (not enough to produce a positive Kleihauer-Betke test for fetal-maternal hemorrhage). The fetal cells can be sorted out and analyzed by a variety of techniques to look for particular DNA sequences, but without the risks that invasive procedures inherently have. This technique may also make use of the phenomenon of free floating fetal DNA gaining access to maternal circulation by DNA release following apoptosis of placental tissue where the placental tissue in question contains DNA of the same genotype as the fetus. The free floating DNA found in maternal plasma has been shown to contain fetal DNA in proportions as high as 30-40% fetal DNA.

[0587] In an embodiment, blood may be drawn from a pregnant woman. Research has shown that maternal blood may contain a small amount of free floating DNA from the fetus, in addition to free floating DNA of maternal origin. In addition, there also may be enucleated fetal blood cells comprising DNA of fetal origin, in addition to many blood cells of maternal origin, which typically do not contain nuclear DNA. There are many methods known in the art to isolate fetal DNA, or create fractions enriched in fetal DNA. For example, chromatography has been shown to create certain fractions that are enriched in fetal DNA.

[0588] Once the sample of maternal blood, plasma, or other fluid, drawn in a relatively non-invasive manner, and that contains an amount of fetal DNA, either cellular or free floating, either enriched in its proportion to the maternal DNA, or in its original ratio, is in hand, one may genotype the DNA found in said sample. In some embodiments, the blood may be drawn using a needle to withdraw blood from a vein, for example, the basilica vein. The method described herein can be used to determine genotypic data of the fetus. For example, it can be used to determine the ploidy state at one or more chromosomes, it can be used to determine the identity of one or a set of SNPs, including insertions, deletions, and translocations. It can be used to determine one or more haplotypes, including the parent of origin of one or more genotypic features.

[0589] Note that this method will work with any nucleic acids that can be used for any genotyping and/or sequencing methods, such as the ILLUMINA INFINITIUM ARRAY platform, AFFYMETRIX GENECHIP, ILLUMINA GENOME ANALYZER, or LIFE TECHNOLOGIES' SOLID SYSTEM. This includes extracted free-floating DNA from plasma or amplifications (e.g. whole genome amplification, PCR) of the same; genomic DNA from other cell types (e.g. human lymphocytes from whole blood) or amplifications of the same. For preparation of the DNA, any extraction or purification method that generates genomic DNA suitable for the one of these platforms will work as well. This method could work equally well with samples of RNA. In an embodiment, storage of the samples may be done in a way that will minimize degradation (e.g. below freezing, at about -20 C, or at a lower temperature).

Parental Support

[0590] Some embodiments may be used in combination with the PARENTAL SUPPORT™ (PS) method, embodiments of which are described in U.S. application Ser. No. 11/603,406 (US Publication No.: 20070184467), U.S. application Ser. No. 12/076,348 (US Publication No.: 20080243398), U.S. application Ser. No. 13/110,685, PCT

Application PCT/US09/52730 (PCT Publication No.: WO/2010/017214), and PCT Application No. PCT/US10/050824 (PCT Publication No.: WO/2011/041485) which are incorporated herein by reference in their entirety. PARENTAL SUPPORT™ is an informatics based approach that can be used to analyze genetic data. In some embodiments, the methods disclosed herein may be considered as part of the PARENTAL SUPPORT™ method. In some embodiments, The PARENTAL SUPPORT™ method is a collection of methods that may be used to determine the genetic data of a target individual, with high accuracy, of one or a small number of cells from that individual, or of a mixture of DNA consisting of DNA from the target individual and DNA from one or a plurality of other individuals, specifically to determine disease-related alleles, other alleles of interest, and/or the ploidy state of one or a plurality of chromosomes in the target individual. PARENTAL SUPPORT™ may refer to any of these methods. PARENTAL SUPPORT™ is an example of an informatics based method. Exemplary embodiments of the PARENTAL SUPPORT™ method are illustrated in FIGS. 29-31G and described in Example 19.

[0591] The PARENTAL SUPPORT™ method makes use of known parental genetic data, i.e. haplotypic and/or diploid genetic data of the mother and/or the father, together with the knowledge of the mechanism of meiosis and the imperfect measurement of the target DNA, and possibly of one or more related individuals, along with population based crossover frequencies, in order to reconstruct, in silico, the genotype at a plurality of alleles, and/or the ploidy state of an embryo or of any target cell(s), and the target DNA at the location of key loci with a high degree of confidence. The PARENTAL SUPPORT™ method can reconstruct not only single nucleotide polymorphisms (SNPs) that were measured poorly, but also insertions and deletions, and SNPs or whole regions of DNA that were not measured at all. Furthermore, the PARENTAL SUPPORT™ method can both measure multiple disease-linked loci as well as screen for aneuploidy, from a single cell. In some embodiments, the PARENTAL SUPPORT™ method may be used to characterize one or more cells from embryos biopsied during an IVF cycle to determine the genetic condition of the one or more cells.

[0592] The PARENTAL SUPPORT™ method allows the cleaning of noisy genetic data. This may be done by inferring the correct genetic alleles in the target genome (embryo) using the genotype of related individuals (parents) as a reference. PARENTAL SUPPORT™ may be particularly relevant where only a small quantity of genetic material is available (e.g. PGD) and where direct measurements of the genotypes are inherently noisy due to the limited amounts of genetic material. PARENTAL SUPPORT™ may be particularly relevant where only a small fraction of the genetic material available is from the target individual (e.g. NPD) and where direct measurements of the genotypes are inherently noisy due to the contaminating DNA signal from another individual. The PARENTAL SUPPORT™ method is able to reconstruct highly accurate ordered diploid allele sequences on the embryo, together with copy number of chromosomes segments, even though the conventional, unordered diploid measurements may be characterized by high rates of allele dropouts, drop-ins, variable amplification biases and other errors. The method may employ both an underlying genetic model and an underlying model of measurement error. The genetic model may determine both allele

probabilities at each SNP and crossover probabilities between SNPs. Allele probabilities may be modeled at each SNP based on data obtained from the parents and model crossover probabilities between SNPs based on data obtained from the HapMap database, as developed by the International HapMap Project. Given the proper underlying genetic model and measurement error model, maximum a posteriori (MAP) estimation may be used, with modifications for computational efficiency, to estimate the correct, ordered allele values at each SNP in the embryo.

[0593] The techniques outlined above, in some cases, are able to determine the genotype of an individual given a very small amount of DNA originating from that individual. This could be the DNA from one or a small number of cells, or it could be from the small amount of fetal DNA found in maternal blood.

Hypotheses

[0594] In the context of this disclosure, a hypothesis refers to a possible genetic state. It may refer to a possible ploidy state. It may refer to a possible allelic state. A set of hypotheses may refer to a set of possible genetic states, a set of possible allelic states, a set of possible ploidy states, or combinations thereof. In some embodiments, a set of hypotheses may be designed such that one hypothesis from the set will correspond to the actual genetic state of any given individual. In some embodiments, a set of hypotheses may be designed such that every possible genetic state may be described by at least one hypothesis from the set. In some embodiments of the present disclosure, one aspect of a method is to determine which hypothesis corresponds to the actual genetic state of the individual in question.

[0595] In another embodiment of the present disclosure, one step involves creating a hypothesis. In some embodiments it may be a copy number hypothesis. In some embodiments it may involve a hypothesis concerning which segments of a chromosome from each of the related individuals correspond genetically to which segments, if any, of the other related individuals. Creating a hypothesis may refer to the act of setting the limits of the variables such that the entire set of possible genetic states that are under consideration are encompassed by those variables.

[0596] A “copy number hypothesis,” also called a “ploidy hypothesis,” or a “ploidy state hypothesis,” may refer to a hypothesis concerning a possible ploidy state for a given chromosome copy, chromosome type, or section of a chromosome, in the target individual. It may also refer to the ploidy state at more than one of the chromosome types in the individual. A set of copy number hypotheses may refer to a set of hypotheses where each hypothesis corresponds to a different possible ploidy state in an individual. A set of hypotheses may concern a set of possible ploidy states, a set of possible parental haplotypes contributions, a set of possible fetal DNA percentages in the mixed sample, or combinations thereof. In some embodiments, the copy number hypotheses include all fetuses in a multiple pregnancy being euploid, all fetuses in a multiple pregnancy being aneuploid (such as any of the aneuploidies disclosed herein), and/or one or more fetuses in a multiple pregnancy being euploid and one or more fetuses in a multiple pregnancy being aneuploidy. In some embodiments, the copy number hypotheses include identical twins (also referred to as monozygotic twins) or fraternal twins (also referred to as dizygotic twins).

In some embodiments, the copy number hypotheses include a molar pregnancy, such as a complete or partial molar pregnancy.

[0597] A normal individual contains one of each chromosome type from each parent. However, due to errors in meiosis and mitosis, it is possible for an individual to have 0, 1, 2, or more of a given chromosome type from each parent. In practice, it is rare to see more than two of a given chromosome from a parent. In this disclosure, some embodiments only consider the possible hypotheses where 0, 1, or 2 copies of a given chromosome come from a parent; it is a trivial extension to consider more or less possible copies originating from a parent. In some embodiments, for a given chromosome, there are nine possible hypotheses: the three possible hypothesis concerning 0, 1, or 2 chromosomes of maternal origin, multiplied by the three possible hypotheses concerning 0, 1, or 2 chromosomes of paternal origin. Let (m,f) refer to the hypothesis where m is the number of a given chromosome inherited from the mother, and f is the number of a given chromosome inherited from the father. Therefore, the nine hypotheses are (0,0), (0,1), (0,2), (1,0), (1,1), (1,2), (2,0), (2,1), and (2,2). These may also be written as H_{00} , H_{01} , H_{02} , H_{10} , H_{11} , H_{12} , H_{20} , H_{21} , and H_{22} . The different hypotheses correspond to different ploidy states. For example, (1,1) refers to a normal disomic chromosome; (2,1) refers to a maternal trisomy, and (0,1) refers to a paternal monosomy. In some embodiments, the case where two chromosomes are inherited from one parent and one chromosome is inherited from the other parent may be further differentiated into two cases: one where the two chromosomes are identical (matched copy error), and one where the two chromosomes are homologous but not identical (unmatched copy error). In these embodiments, there are sixteen possible hypotheses. It should be understood that it is possible to use other sets of hypotheses, and a different number of hypotheses.

[0598] In some embodiments of the present disclosure, the ploidy hypothesis refers to a hypothesis concerning which chromosome from other related individuals correspond to a chromosome found in the target individual's genome. In some embodiments, a key to the method is the fact that related individuals can be expected to share haplotype blocks, and using measured genetic data from related individuals, along with a knowledge of which haplotype blocks match between the target individual and the related individual, it is possible to infer the correct genetic data for a target individual with higher confidence than using the target individual's genetic measurements alone. As such, in some embodiments, the ploidy hypothesis may concern not only the number of chromosomes, but also which chromosomes in related individuals are identical, or nearly identical, with one or more chromosomes in the target individual.

[0599] Once the set of hypotheses have been defined, when the algorithms operate on the input genetic data, they may output a determined statistical probability for each of the hypotheses under consideration. The probabilities of the various hypotheses may be determined by mathematically calculating, for each of the various hypotheses, the value that the probability equals, as stated by one or more of the expert techniques, algorithms, and/or methods described elsewhere in this disclosure, using the relevant genetic data as input.

[0600] Once the probabilities of the different hypotheses are estimated, as determined by a plurality of techniques,

they may be combined. This may entail, for each hypothesis, multiplying the probabilities as determined by each technique. The product of the probabilities of the hypotheses may be normalized. Note that one ploidy hypothesis refers to one possible ploidy state for a chromosome.

[0601] The process of "combining probabilities," also called "combining hypotheses," or combining the results of expert techniques, is a concept that should be familiar to one skilled in the art of linear algebra. One possible way to combine probabilities is as follows: When an expert technique is used to evaluate a set of hypotheses given a set of genetic data, the output of the method is a set of probabilities that are associated, in a one-to-one fashion, with each hypothesis in the set of hypotheses. When a set of probabilities that were determined by a first expert technique, each of which are associated with one of the hypotheses in the set, are combined with a set of probabilities that were determined by a second expert technique, each of which are associated with the same set of hypotheses, then the two sets of probabilities are multiplied. This means that, for each hypothesis in the set, the two probabilities that are associated with that hypothesis, as determined by the two expert methods, are multiplied together, and the corresponding product is the output probability. This process may be expanded to any number of expert techniques. If only one expert technique is used, then the output probabilities are the same as the input probabilities. If more than two expert techniques are used, then the relevant probabilities may be multiplied at the same time. The products may be normalized so that the probabilities of the hypotheses in the set of hypotheses sum to 100%.

[0602] In some embodiments, if the combined probabilities for a given hypothesis are greater than the combined probabilities for any of the other hypotheses, then it may be considered that that hypothesis is determined to be the most likely. In some embodiments, a hypothesis may be determined to be the most likely, and the ploidy state, or other genetic state, may be called if the normalized probability is greater than a threshold. In an embodiment, this may mean that the number and identity of the chromosomes that are associated with that hypothesis may be called as the ploidy state. In an embodiment, this may mean that the identity of the alleles that are associated with that hypothesis may be called as the allelic state. In some embodiments, the threshold may be between about 50% and about 80%. In some embodiments the threshold may be between about 80% and about 90%. In some embodiments the threshold may be between about 90% and about 95%. In some embodiments the threshold may be between about 95% and about 99%. In some embodiments the threshold may be between about 99% and about 99.9%. In some embodiments the threshold may be above about 99.9%.

Parental Contexts

[0603] The parental context refers to the genetic state of a given allele, on each of the two relevant chromosomes for one or both of the two parents of the target. Note that in an embodiment, the parental context does not refer to the allelic state of the target, rather, it refers to the allelic state of the parents. The parental context for a given SNP may consist of four base pairs, two paternal and two maternal; they may be the same or different from one another. It is typically written as " $m_1 m_2 f_1 f_2$," where m_1 and m_2 are the genetic state of the given SNP on the two maternal chromosomes, and f_1 and f_2

are the genetic state of the given SNP on the two paternal chromosomes. In some embodiments, the parental context may be written as " $f_1f_2|m_1m_2$." Note that subscripts "1" and "2" refer to the genotype, at the given allele, of the first and second chromosome; also note that the choice of which chromosome is labeled "1" and which is labeled "2" is arbitrary.

[0604] Note that in this disclosure, A and B are often used to generically represent base pair identities; A or B could equally well represent C (cytosine), G (guanine), A (adenine) or T (thymine). For example, if, at a given SNP based allele, the mother's genotype was T at that SNP on one chromosome, and G at that SNP on the homologous chromosome, and the father's genotype at that allele is G at that SNP on both of the homologous chromosomes, one may say that the target individual's allele has the parental context of ABIBB; it could also be said that the allele has the parental context of ABIAA. Note that, in theory, any of the four possible nucleotides could occur at a given allele, and thus it is possible, for example, for the mother to have a genotype of AT, and the father to have a genotype of GC at a given allele. However, empirical data indicate that in most cases only two of the four possible base pairs are observed at a given allele. It is possible, for example when using single tandem repeats, to have more than two parental, more than four and even more than ten contexts. In this disclosure the discussion assumes that only two possible base pairs will be observed at a given allele, although the embodiments disclosed herein could be modified to take into account the cases where this assumption does not hold.

[0605] A "parental context" may refer to a set or subset of target SNPs that have the same parental context. For example, if one were to measure 1000 alleles on a given chromosome on a target individual, then the context AAIBB could refer to the set of all alleles in the group of 1,000 alleles where the genotype of the mother of the target was homozygous, and the genotype of the father of the target is homozygous, but where the maternal genotype and the paternal genotype are dissimilar at that locus. If the parental data is not phased, and thus AB=BA, then there are nine possible parental contexts: AAIAA, AAIAB, AAIBB, ABIAA, ABIAB, ABIBB, BBIAA, BBIAB, and BBIBB. If the parental data is phased, and thus AB BA, then there are sixteen different possible parental contexts: AAIAA, AAIAB, AAIBA, AAIBB, ABIAA, ABIAB, ABIBA, ABIBB, BAIAA, BAIBA, BAIBB, BBIAA, BBIAB, BBIBA, and BBIBB. Every SNP allele on a chromosome, excluding some SNPs on the sex chromosomes, has one of these parental contexts. The set of SNPs wherein the parental context for one parent is heterozygous may be referred to as the heterozygous context.

Use of Parental Contexts in NPD

[0606] Non-invasive prenatal diagnosis is an important technique that can be used to determine the genetic state of a fetus from genetic material that is obtained in a non-invasive manner, for example from a blood draw on the pregnant mother. The blood could be separated and the plasma isolated, followed by isolation of the plasma DNA. Size selection could be used to isolate the DNA of the appropriate length. The DNA may be preferentially enriched at a set of loci. This DNA can then be measured by a number

of means, such as by hybridizing to a genotyping array and measuring the fluorescence, or by sequencing on a high throughput sequencer.

[0607] When sequencing is used for ploidy calling of a fetus in the context of non-invasive prenatal diagnosis, there are a number of ways to use the sequence data. The most common way one could use the sequence data is to simply count the number of reads that map to a given chromosome. For example, imagine if you are trying to determine the ploidy state of chromosome 21 on the fetus. Further imagine that the DNA in the sample is comprised of 10% DNA of fetal origin, and 90% DNA of maternal origin. In this case, you could look at the average number of reads on a chromosome which can be expected to be disomic, for example chromosome 3, and compare that to the number of read on chromosome 21, where the reads are adjusted for the number of base pairs on that chromosome that are part of a unique sequence. If the fetus were euploid, one would expect the amount of DNA per unit of genome to be about equal at all locations (subject to stochastic variations). On the other hand, if the fetus were trisomic at chromosome 21, then one would expect there to be more slightly more DNA per genetic unit from chromosome 21 than the other locations on the genome. Specifically one would expect there to be about 5% more DNA from chromosome 21 in the mixture. When sequencing is used to measure the DNA, one would expect about 5% more uniquely mappable reads from chromosome 21 per unique segment than from the other chromosomes. One could use the observation of an amount of DNA from a particular chromosome that is higher than a certain threshold, when adjusted for the number of sequences that are uniquely mappable to that chromosome, as the basis for an aneuploidy diagnosis. Another method that may be used to detect aneuploidy is similar to that above, except that parental contexts could be taken into account.

[0608] When considering which alleles to target, one may consider the likelihood that some parental contexts are likely to be more informative than others. For example, AAIBB and the symmetric context BBIAA are the most informative contexts, because the fetus is known to carry an allele that is different from the mother. For reasons of symmetry, both AAIBB and BBIAA contexts may be referred to as AAIBB. Another set of informative parental contexts are AAIAB and BBIAB, because in these cases the fetus has a 50% chance of carrying an allele that the mother does not have. For reasons of symmetry, both AAIAB and BBIAB contexts may be referred to as AAIAB. A third set of informative parental contexts are ABIAA and ABIBB, because in these cases the fetus is carrying a known paternal allele, and that allele is also present in the maternal genome. For reasons of symmetry, both ABIAA and ABIBB contexts may be referred to as ABIAA. A fourth parental context is ABIAB where the fetus has an unknown allelic state, and whatever the allelic state, it is one in which the mother has the same alleles. The fifth parental context is AAIAA, where the mother and father are heterozygous.

Different Implementations of the Presently Disclosed Embodiments

[0609] Methods are disclosed herein for determining the ploidy state of a target individual. The target individual may be a blastomere, an embryo, or a fetus. In some embodiments of the present disclosure, a method for determining the ploidy state of one or more chromosome in a target

individual may include any of the steps described in this document, and combinations thereof:

[0610] In some embodiments the source of the genetic material to be used in determining the genetic state of the fetus may be fetal cells, such as nucleated fetal red blood cells, isolated from the maternal blood. The method may involve obtaining a blood sample from the pregnant mother. The method may involve isolating a fetal red blood cell using visual techniques, based on the idea that a certain combination of colors are uniquely associated with nucleated red blood cell, and a similar combination of colors is not associated with any other present cell in the maternal blood. The combination of colors associated with the nucleated red blood cells may include the red color of the hemoglobin around the nucleus, which color may be made more distinct by staining, and the color of the nuclear material which can be stained, for example, blue. By isolating the cells from maternal blood and spreading them over a slide, and then identifying those points at which one sees both red (from the Hemoglobin) and blue (from the nuclear material) one may be able to identify the location of nucleated red blood cells. One may then extract those nucleated red blood cells using a micromanipulator, use genotyping and/or sequencing techniques to measure aspects of the genotype of the genetic material in those cells.

[0611] In an embodiment, one may stain the nucleated red blood cell with a dye that only fluoresces in the presence of fetal hemoglobin and not maternal hemoglobin, and so remove the ambiguity between whether a nucleated red blood cell is derived from the mother or the fetus. Some embodiments of the present disclosure may involve staining or otherwise marking nuclear material. Some embodiments of the present disclosure may involve specifically marking fetal nuclear material using fetal cell specific antibodies.

[0612] There are many other ways to isolate fetal cells from maternal blood, or fetal DNA from maternal blood, or to enrich samples of fetal genetic material in the presence of maternal genetic material. Some of these methods are listed here, but this is not intended to be an exhaustive list. Some appropriate techniques are listed here for convenience: using fluorescently or otherwise tagged antibodies, size exclusion chromatography, magnetically or otherwise labeled affinity tags, epigenetic differences, such as differential methylation between the maternal and fetal cells at specific alleles, density gradient centrifugation succeeded by CD45/14 depletion and CD71-positive selection from CD45/14 negative-cells, single or double Percoll gradients with different osmolalities, or galactose specific lectin method.

[0613] In an embodiment of the present disclosure, the target individual is a fetus, and the different genotype measurements are made on a plurality of DNA samples from the fetus. In some embodiments of the present disclosure, the fetal DNA samples are from isolated fetal cells where the fetal cells may be mixed with maternal cells. In some embodiments of the present disclosure, the fetal DNA samples are from free floating fetal DNA, where the fetal DNA may be mixed with free floating maternal DNA. In some embodiments, the fetal DNA samples may be derived from maternal plasma or maternal blood that contains a mixture of maternal DNA and fetal DNA. In some embodiments, the fetal DNA may be mixed with maternal DNA in maternal:fetal ratios ranging from 99.9:0.1% to 99:1%;

99:1% to 90:10%; 90:10% to 80:20%; 80:20% to 70:30%; 70:30% to 50:50%; 50:50% to 10:90%; or 10:90% to 1:99%; 1:99% to 0.1:99.9%.

[0614] The genetic data of the target individual and/or of the related individual can be transformed from a molecular state to an electronic state by measuring the appropriate genetic material using tools and or techniques taken from a group including, but not limited to: genotyping microarrays, and high throughput sequencing. Some high throughput sequencing methods include Sanger DNA sequencing, pyrosequencing, the ILLUMINA SOLEXA platform, ILLUMINA's GENOME ANALYZER, or APPLIED BIOSYSTEM's 454 sequencing platform, HELICOS's TRUE SINGLE MOLECULE SEQUENCING platform, HALCYON MOLECULAR's electron microscope sequencing method, or any other sequencing method.

[0615] All of these methods physically transform the genetic data stored in a sample of DNA into a set of genetic data that is typically stored in a memory device in route to being processed.

[0616] A relevant individual's genetic data may be measured by analyzing substances taken from a group including, but not limited to: the individual's bulk diploid tissue, one or more diploid cells from the individual, one or more haploid cells from the individual, one or more blastomeres from the target individual, extra-cellular genetic material found on the individual, extra-cellular genetic material from the individual found in maternal blood, cells from the individual found in maternal blood, one or more embryos created from (a) gamete(s) from the related individual, one or more blastomeres taken from such an embryo, extra-cellular genetic material found on the related individual, genetic material known to have originated from the related individual, and combinations thereof.

[0617] In some embodiments, a set of at least one ploidy state hypothesis may be created for each of the chromosomes types of interest of the target individual. Each of the ploidy state hypotheses may refer to one possible ploidy state of the chromosome or chromosome segment of the target individual. The set of hypotheses may include some or all of the possible ploidy states that the chromosome of the target individual may be expected to have. Some of the possible ploidy states may include nullsomy, monosomy, disomy, uniparental disomy, euploidy, trisomy, matching trisomy, unmatching trisomy, maternal trisomy, paternal trisomy, tetrasomy, balanced (2:2) tetrasomy, unbalanced (3:1) tetrasomy, pentasomy, hexasomy, other aneuploidy, and combinations thereof. Any of these aneuploidy states may be mixed or partial aneuploidy such as unbalanced translocations, balanced translocations, Robertsonian translocations, recombinations, deletions, insertions, crossovers, and combinations thereof.

[0618] In some embodiments, the knowledge of the determined ploidy state may be used to make a clinical decision. This knowledge, typically stored as a physical arrangement of matter in a memory device, may then be transformed into a report. The report may then be acted upon. For example, the clinical decision may be to terminate the pregnancy; alternately, the clinical decision may be to continue the pregnancy. In some embodiments the clinical decision may involve an intervention designed to decrease the severity of the phenotypic presentation of a genetic disorder, or a decision to take relevant steps to prepare for a special needs child.

[0619] In an embodiment of the present disclosure, any of the methods described herein may be modified to allow for multiple targets to come from same target individual, for example, multiple blood draws from the same pregnant mother. This may improve the accuracy of the model, as multiple genetic measurements may provide more data with which the target genotype may be determined. In an embodiment, one set of target genetic data served as the primary data which was reported, and the other served as data to double-check the primary target genetic data. In an embodiment, a plurality of sets of genetic data, each measured from genetic material taken from the target individual, are considered in parallel, and thus both sets of target genetic data serve to help determine which sections of parental genetic data, measured with high accuracy, composes the fetal genome.

[0620] In an embodiment, the method may be used for the purpose of paternity testing. For example, given the SNP-based genotypic information from the mother, and from a man who may or may not be the genetic father, and the measured genotypic information from the mixed sample, it is possible to determine if the genotypic information of the male indeed represents that actual genetic father of the gestating fetus. A simple way to do this is to simply look at the contexts where the mother is AA, and the possible father is AB or BB. In these cases, one may expect to see the father contribution half (AAIAB) or all (AAIBB) of the time, respectively. Taking into account the expected ADO, it is straightforward to determine whether or not the fetal SNPs that are observed are correlated with those of the possible father.

[0621] One embodiment of the present disclosure could be as follows: a pregnant woman wants to know if her fetus is afflicted with Down Syndrome, and/or if it will suffer from Cystic Fibrosis, and she does not wish to bear a child that is afflicted with either of these conditions. A doctor takes her blood, and stains the hemoglobin with one marker so that it appears clearly red, and stains nuclear material with another marker so that it appears clearly blue. Knowing that maternal red blood cells are typically anuclear, while a high proportion of fetal cells contain a nucleus, the doctor is able to visually isolate a number of nucleated red blood cells by identifying those cells that show both a red and blue color. The doctor picks up these cells off the slide with a micro-manipulator and sends them to a lab which amplifies and genotypes ten individual cells. By using the genetic measurements, the PARENTAL SUPPORT™ method is able to determine that six of the ten cells are maternal blood cells, and four of the ten cells are fetal cells. If a child has already been born to a pregnant mother, PARENTAL SUPPORT™ can also be used to determine that the fetal cells are distinct from the cells of the born child by making reliable allele calls on the fetal cells and showing that they are dissimilar to those of the born child. Note that this method is similar in concept to the paternal testing embodiment of the present disclosure. The genetic data measured from the fetal cells may be of very poor quality, comprising many allele drop outs, due to the difficulty of genotyping single cells. The clinician is able to use the measured fetal DNA along with the reliable DNA measurements of the parents to infer aspects of the genome of the fetus with high accuracy using PARENTAL SUPPORT™, thereby transforming the genetic data contained on genetic material from the fetus into the predicted genetic state of the fetus, stored on a computer.

The clinician is able to determine both the ploidy state of the fetus, and the presence or absence of a plurality of disease-linked genes of interest. It turns out that the fetus is euploid, and is not a carrier for cystic fibrosis, and the mother decides to continue the pregnancy.

[0622] In an embodiment of the present disclosure, a pregnant mother would like to determine if her fetus is afflicted with any whole chromosomal abnormalities. She goes to her doctor, and gives a sample of her blood, and she and her husband gives samples of their own DNA from cheek swabs. A laboratory researcher genotypes the parental DNA using the MDA protocol to amplify the parental DNA, and ILLUMINA INFINIUM arrays to measure the genetic data of the parents at a large number of SNPs. The researcher then spins down the blood, takes the plasma, and isolates a sample of free-floating DNA using size exclusion chromatography. Alternately, the researcher uses one or more fluorescent antibodies, such as one that is specific to fetal hemoglobin to isolate a nucleated fetal red blood cell. The researcher then takes the isolated or enriched fetal genetic material and amplifies it using a library of 70-mer oligonucleotides appropriately designed such that two ends of each oligonucleotide corresponded to the flanking sequences on either side of a target allele. Upon addition of a polymerase, ligase, and the appropriate reagents, the oligonucleotides underwent gap-filling circularization, capturing the desired allele. An exonuclease was added, heat-inactivated, and the products were used directly as a template for PCR amplification. The PCR products were sequenced on an ILLUMINA GENOME ANALYZER. The sequence reads were used as input for the PARENTAL SUPPORT™ method, which then predicted the ploidy state of the fetus.

[0623] In another embodiment, a couple—where the mother, who is pregnant, and is of advanced maternal age—wants to know whether the gestating fetus has Down syndrome, Turner Syndrome, Prader Willi syndrome, or some other whole chromosomal abnormality. The obstetrician takes a blood draw from the mother and father. The blood is sent to a laboratory, where a technician centrifuges the maternal sample to isolate the plasma and the buffy coat. The DNA in the buffy coat and the paternal blood sample are transformed through amplification and the genetic data encoded in the amplified genetic material is further transformed from molecularly stored genetic data into electronically stored genetic data by running the genetic material on a high throughput sequencer to measure the parental genotypes. The plasma sample is preferentially enriched at a set of loci using a 5,000-plex hemi-nested targeted PCR method. The mixture of DNA fragments is prepared into a DNA library suitable for sequencing. The DNA is then sequenced using a high throughput sequencing method, for example, the ILLUMINA GAIIX GENOME ANALYZER. The sequencing transforms the information that is encoded molecularly in the DNA into information that is encoded electronically in computer hardware. An informatics based technique that includes the presently disclosed embodiments, such as PARENTAL SUPPORT™, may be used to determine the ploidy state of the fetus. This may involve calculating, on a computer, allele count probabilities at the plurality of polymorphic loci from the DNA measurements made on the prepared sample; creating, on a computer, a plurality of ploidy hypotheses each pertaining to a different possible ploidy state of the chromosome; building, on a computer, a joint distribution model for the expected allele

counts at the plurality of polymorphic loci on the chromosome for each ploidy hypothesis; determining, on a computer, a relative probability of each of the ploidy hypotheses using the joint distribution model and the allele counts measured on the prepared sample; and calling the ploidy state of the fetus by selecting the ploidy state corresponding to the hypothesis with the greatest probability. It is determined that the fetus has Down syndrome. A report is printed out, or sent electronically to the pregnant woman's obstetrician, who transmits the diagnosis to the woman. The woman, her husband, and the doctor sit down and discuss their options. The couple decides to terminate the pregnancy based on the knowledge that the fetus is afflicted with a trisomic condition.

[0624] In an embodiment, a company may decide to offer a diagnostic technology designed to detect aneuploidy in a gestating fetus from a maternal blood draw. Their product may involve a mother presenting to her obstetrician, who may draw her blood. The obstetrician may also collect a genetic sample from the father of the fetus. A clinician may isolate the plasma from the maternal blood, and purify the DNA from the plasma. A clinician may also isolate the buffy coat layer from the maternal blood, and prepare the DNA from the buffy coat. A clinician may also prepare the DNA from the paternal genetic sample. The clinician may use molecular biology techniques described in this disclosure to append universal amplification tags to the DNA in the DNA derived from the plasma sample. The clinician may amplify the universally tagged DNA. The clinician may preferentially enrich the DNA by a number of techniques including capture by hybridization and targeted PCR. The targeted PCR may involve nesting, hemi-nesting or semi-nesting, or any other approach to result in efficient enrichment of the plasma derived DNA. The targeted PCR may be massively multiplexed, for example with 10,000 primers in one reaction volume, where the primers target SNPs on chromosomes 13, 18, 21, X and those loci that are common to both X and Y, and optionally other chromosomes as well. The selective enrichment and/or amplification may involve tagging each individual molecule with different tags, molecular barcodes, tags for amplification, and/or tags for sequencing. The clinician may then sequence the plasma sample, and also possibly also the prepared maternal and/or paternal DNA. The molecular biology steps may be executed either wholly or partly by a diagnostic box. The sequence data may be fed into a single computer, or to another type of computing platform such as may be found in 'the cloud'. The computing platform may calculate allele counts at the targeted polymorphic loci from the measurements made by the sequencer. The computing platform may create a plurality of ploidy hypotheses pertaining to nullsomy, monosomy, disomy, matched trisomy, and unmatched trisomy for each of chromosomes 13, 18, 21, X and Y. The computing platform may build a joint distribution model for the expected allele counts at the targeted loci on the chromosome for each ploidy hypothesis for each of the five chromosomes being interrogated. The computing platform may determine a probability that each of the ploidy hypotheses is true using the joint distribution model and the allele counts measured on the preferentially enriched DNA derived from the plasma sample. The computing platform may call the ploidy state of the fetus, for each of chromosome 13, 18, 21, X and Y by selecting the ploidy state corresponding to the germane hypothesis with the greatest probability. A report

may be generated comprising the called ploidy states, and it may be sent to the obstetrician electronically, displayed on an output device, or a printed hard copy of the report may be delivered to the obstetrician. The obstetrician may inform the patient and optionally the father of the fetus, and they may decide which clinical options are open to them, and which is most desirable.

[0625] In another embodiment, a pregnant woman, hereafter referred to as "the mother" may decide that she wants to know whether or not her fetus(es) are carrying any genetic abnormalities or other conditions. She may want to ensure that there are not any gross abnormalities before she is confident to continue the pregnancy. She may go to her obstetrician, who may take a sample of her blood. He may also take a genetic sample, such as a buccal swab, from her cheek. He may also take a genetic sample from the father of the fetus, such as a buccal swab, a sperm sample, or a blood sample. He may send the samples to a clinician. The clinician may enrich the fraction of free floating fetal DNA in the maternal blood sample. The clinician may enrich the fraction of enucleated fetal blood cells in the maternal blood sample. The clinician may use various aspects of the methods described herein to determine genetic data of the fetus. That genetic data may include the ploidy state of the fetus, and/or the identity of one or a number of disease linked alleles in the fetus. A report may be generated summarizing the results of the prenatal diagnosis. The report may be transmitted or mailed to the doctor, who may tell the mother the genetic state of the fetus. The mother may decide to discontinue the pregnancy based on the fact that the fetus has one or more chromosomal, or genetic abnormalities, or undesirable conditions. She may also decide to continue the pregnancy based on the fact that the fetus does not have any gross chromosomal or genetic abnormalities, or any genetic conditions of interest.

[0626] Another example may involve a pregnant woman who has been artificially inseminated by a sperm donor, and is pregnant. She wants to minimize the risk that the fetus she is carrying has a genetic disease. She has blood drawn at a phlebotomist, and techniques described in this disclosure are used to isolate three nucleated fetal red blood cells, and a tissue sample is also collected from the mother and genetic father. The genetic material from the fetus and from the mother and father are amplified as appropriate and genotyped using the ILLUMINA INFINITIUM BEADARRAY, and the methods described herein clean and phase the parental and fetal genotype with high accuracy, as well as to make ploidy calls for the fetus. The fetus is found to be euploid, and phenotypic susceptibilities are predicted from the reconstructed fetal genotype, and a report is generated and sent to the mother's physician so that they can decide what clinical decisions may be best.

[0627] In an embodiment, the raw genetic material of the mother and the father is transformed by way of amplification to an amount of DNA that is similar in sequence, but larger in quantity. Then, by way of a genotyping method, the genotypic data that is encoded by nucleic acids is transformed into genetic measurements that may be stored physically and/or electronically on a memory device, such as those described above. The relevant algorithms that makeup the PARENTAL SUPPORT™ algorithm, relevant parts of which are discussed in detail herein, are translated into a computer program, using a programming language. Then, through the execution of the computer program on the

computer hardware, instead of being physically encoded bits and bytes, arranged in a pattern that represents raw measurement data, they become transformed into a pattern that represents a high confidence determination of the ploidy state of the fetus. The details of this transformation will rely on the data itself and the computer language and hardware system used to execute the method described herein. Then, the data that is physically configured to represent a high quality ploidy determination of the fetus is transformed into a report which may be sent to a health care practitioner. This transformation may be carried out using a printer or a computer display. The report may be a printed copy, on paper or other suitable medium, or else it may be electronic. In the case of an electronic report, it may be transmitted, it may be physically stored on a memory device at a location on the computer accessible by the health care practitioner; it also may be displayed on a screen so that it may be read. In the case of a screen display, the data may be transformed to a readable format by causing the physical transformation of pixels on the display device. The transformation may be accomplished by way of physically firing electrons at a phosphorescent screen, by way of altering an electric charge that physically changes the transparency of a specific set of pixels on a screen that may lie in front of a substrate that emits or absorbs photons. This transformation may be accomplished by way of changing the nanoscale orientation of the molecules in a liquid crystal, for example, from nematic to cholesteric or smectic phase, at a specific set of pixels. This transformation may be accomplished by way of an electric current causing photons to be emitted from a specific set of pixels made from a plurality of light emitting diodes arranged in a meaningful pattern. This transformation may be accomplished by any other way used to display information, such as a computer screen, or some other output device or way of transmitting information. The health care practitioner may then act on the report, such that the data in the report is transformed into an action. The action may be to continue or discontinue the pregnancy, in which case a gestating fetus with a genetic abnormality is transformed into non-living fetus. The transformations listed herein may be aggregated, such that, for example, one may transform the genetic material of a pregnant mother and the father, through a number of steps outlined in this disclosure, into a medical decision consisting of aborting a fetus with genetic abnormalities, or consisting of continuing the pregnancy. Alternately, one may transform a set of genotypic measurements into a report that helps a physician treat his pregnant patient.

[0628] In an embodiment of the present disclosure, the method described herein can be used to determine the ploidy state of a fetus even when the host mother, i.e. the woman who is pregnant, is not the biological mother of the fetus she is carrying. In an embodiment of the present disclosure, the method described herein can be used to determine the ploidy state of a fetus using only the maternal blood sample, and without the need for a paternal genetic sample.

[0629] Some of the math in the presently disclosed embodiments makes hypotheses concerning a limited number of states of aneuploidy. In some cases, for example, only zero, one or two chromosomes are expected to originate from each parent. In some embodiments of the present disclosure, the mathematical derivations can be expanded to take into account other forms of aneuploidy, such as quadrosomy, where three chromosomes originate from one

parent, pentasomy, hexasomy etc., without changing the fundamental concepts of the present disclosure. At the same time, it is possible to focus on a smaller number of ploidy states, for example, only trisomy and disomy. Note that ploidy determinations that indicate a non-whole number of chromosomes may indicate mosaicism in a sample of genetic material.

[0630] In some embodiments, the genetic abnormality is a type of aneuploidy, such as Down syndrome (or trisomy 21), Edwards syndrome (trisomy 18), Patau syndrome (trisomy 13), Turner Syndrome (45X), Klinefelter's syndrome (a male with 2 X chromosomes), Prader-Willi syndrome, and DiGeorge syndrome (UPD 15). Congenital disorders, such as those listed in the prior sentence, are commonly undesirable, and the knowledge that a fetus is afflicted with one or more phenotypic abnormalities may provide the basis for a decision to terminate the pregnancy, to take necessary precautions to prepare for the birth of a special needs child, or to take some therapeutic approach meant to lessen the severity of a chromosomal abnormality.

[0631] In some embodiments, the methods described herein can be used at a very early gestational age, for example as early as four week, as early as five weeks, as early as six weeks, as early as seven weeks, as early as eight weeks, as early as nine weeks, as early as ten weeks, as early as eleven weeks, and as early as twelve weeks.

[0632] In some embodiments, a method disclosed herein is used in the context of pre-implantation genetic diagnosis (PGD) for embryo selection during in vitro fertilization, where the target individual is an embryo, and the parental genotypic data can be used to make ploidy determinations about the embryo from sequencing data from a single or two cell biopsy from a day 3 embryo or a trophectoderm biopsy from a day 5 or day 6 embryo. In a PGD setting, only the child DNA is measured, and only a small number of cells are tested, generally one to five but as many as ten, twenty or fifty. The total number of starting copies of the A and B alleles (at a SNP) are then trivially determined by the child genotype and the number of cells. In NPD, the number of starting copies is very high and so the allele ratio after PCR is expected to accurately reflect the starting ratio. However, the small number of starting copies in PGD means that contamination and imperfect PCR efficiency have a non-trivial effect on the allele ratio following PCR. This effect may be more important than depth of read in predicting the variance in the allele ratio measured after sequencing. The distribution of measured allele ratio given a known child genotype may be created by Monte Carlo simulation of the PCR process based on the PCR probe efficiency and probability of contamination. Given an allele ratio distribution for each possible child genotype, the likelihoods of various hypotheses can be calculated as described for NIPD.

Maximum Likelihood Estimates

[0633] Most methods known in the art for detecting the presence or absence of biological phenomenon or medical condition involve the use of a single hypothesis rejection test, where a metric that is correlated with the condition is measured, and if the metric is on one side of a given threshold, the condition is present, while if the metric falls on the other side of the threshold, the condition is absent. A single-hypothesis rejection test only looks at the null distribution when deciding between the null and alternate hypotheses. Without taking into account the alternate distribution,

one cannot estimate the likelihood of each hypothesis given the observed data and therefore cannot calculate a confidence on the call. Hence with a single-hypothesis rejection test, one gets a yes or no answer without a feeling for the confidence associated with the specific case.

[0634] In some embodiments, the method disclosed herein is able to detect the presence or absence of biological phenomenon or medical condition using a maximum likelihood method. This is a substantial improvement over a method using a single hypothesis rejection technique as the threshold for calling absence or presence of the condition can be adjusted as appropriate for each case. This is particularly relevant for diagnostic techniques that aim to determine the presence or absence of aneuploidy in a gestating fetus from genetic data available from the mixture of fetal and maternal DNA present in the free floating DNA found in maternal plasma. This is because as the fraction of fetal DNA in the plasma derived fraction changes, the optimal threshold for calling aneuploidy vs. euploidy changes. As the fetal fraction drops, the distribution of data that is associated with an aneuploidy becomes increasingly similar to the distribution of data that is associated with a euploidy.

[0635] The maximum likelihood estimation method uses the distributions associated with each hypothesis to estimate the likelihood of the data conditioned on each hypothesis. These conditional probabilities can then be converted to a hypothesis call and confidence. Similarly, maximum a posteriori estimation method uses the same conditional probabilities as the maximum likelihood estimate, but also incorporates population priors when choosing the best hypothesis and determining confidence.

[0636] Therefore, the use of a maximum likelihood estimate (MLE) technique, or the closely related maximum a posteriori (MAP) technique give two advantages, first it increases the chance of a correct call, and it also allows a confidence to be calculated for each call. In an embodiment, selecting the ploidy state corresponding to the hypothesis with the greatest probability is carried out using maximum likelihood estimates or maximum a posteriori estimates. In an embodiment, a method is disclosed for determining the ploidy state of a gestating fetus that involves taking any method currently known in the art that uses a single hypothesis rejection technique and reformulating it such that it uses a MLE or MAP technique. Some examples of methods that can be significantly improved by applying these techniques can be found in U.S. Pat. Nos. 8,008,018, 7,888,017, or U.S. Pat. No. 7,332,277.

[0637] In an embodiment, a method is described for determining presence or absence of fetal aneuploidy in a maternal plasma sample comprising fetal and maternal genomic DNA, the method comprising: obtaining a maternal plasma sample; measuring the DNA fragments found in the plasma sample with a high throughput sequencer; mapping the sequences to the chromosome and determining the number of sequence reads that map to each chromosome; calculating the fraction of fetal DNA in the plasma sample; calculating an expected distribution of the amount of a target chromosome that would be expected to be present if that if the second target chromosome were euploid and one or a plurality of expected distributions that would be expected if that chromosome were aneuploid, using the fetal fraction and the number of sequence reads that map to one or a plurality of reference chromosomes expected to be euploid;

and using a MLE or MAP determine which of the distributions is most likely to be correct, thereby indicating the presence or absence of a fetal aneuploidy. In an embodiment, the measuring the DNA from the plasma may involve conducting massively parallel shotgun sequencing. In an embodiment, the measuring the DNA from the plasma sample may involve sequencing DNA that has been preferentially enriched, for example through targeted amplification, at a plurality of polymorphic or non-polymorphic loci. The plurality of loci may be designed to target one or a small number of suspected aneuploid chromosomes and one or a small number of reference chromosomes. The purpose of the preferential enrichment is to increase the number of sequence reads that are informative for the ploidy determination.

Ploidy Calling Informatics Methods

[0638] Described herein is a method for determining the ploidy state of a fetus given sequence data. In some embodiments, this sequence data may be measured on a high throughput sequencer. In some embodiments, the sequence data may be measured on DNA that originated from free floating DNA isolated from maternal blood, wherein the free floating DNA comprises some DNA of maternal origin, and some DNA of fetal/placental origin. This section will describe one embodiment of the present disclosure in which the ploidy state of the fetus is determined assuming that fraction of fetal DNA in the mixture that has been analyzed is not known and will be estimated from the data. It will also describe an embodiment in which the fraction of fetal DNA ("fetal fraction") or the percentage of fetal DNA in the mixture can be measured by another method, and is assumed to be known in determining the ploidy state of the fetus. In some embodiments the fetal fraction can be calculated using only the genotyping measurements made on the maternal blood sample itself, which is a mixture of fetal and maternal DNA. In some embodiments the fraction may be calculated also using the measured or otherwise known genotype of the mother and/or the measured or otherwise known genotype of the father. In another embodiment ploidy state of the fetus can be determined solely based on the calculated fraction of fetal DNA for the chromosome in question compared to the calculated fraction of fetal DNA for the reference chromosome assumed disomic.

[0639] In the preferred embodiment, suppose that, for a particular chromosome, we observe and analyze N SNPs, for which we have:

[0640] Set of NR free floating DNA sequence measurements $S=(s_1, \dots, s_{NR})$. Since this method utilizes the SNP measurements, all sequence data that corresponds to non-polymorphic loci can be disregarded. In a simplified version, where we have (A,B) counts on each SNP, where A and B correspond to the two alleles present at a given locus, S can be written as $S=((a_1, b_1), \dots, (a_N, b_N))$, where a_i is the A count on SNP i , b_i is the B count on SNP i , and $\sum_{i=1}^N (a_i + b_i) = NR$

[0641] Parent data consisting of

[0642] Genotypes from a SNP microarray or other intensity based genotyping platform: mother $M=(m_1, \dots, m_N)$, father $F=(f_1, \dots, f_N)$, where $m_i, f_i \in \{AA, AB, BB\}$.

[0643] AND/OR sequence data measurements: NRM mother measurements $SM=(sm_1, \dots, sm_{nrm})$, NRF father measurements $SF=(sf_1, \dots, sf_{nrf})$. Similar to the above

simplification, if we have (A,B) counts on each SNP $SM = ((am_1, bm_1), \dots, (am_N, bm_N))$, $SF = ((af_1, bf_1), \dots, (af_N, bf_N))$

[0644] Collectively, the mother, father child data are denoted as $D = (M, F, SM, SF, S)$. Note that the parent data is desired and increases the accuracy of the algorithm, but is NOT necessary, especially the father data. This means that even in the absence of mother and/or father data, it is possible to get very accurate copy number results.

[0645] It is possible to derive the best copy number estimate (H^*) by maximizing the data log likelihood $LIK(D|H)$ over all hypotheses (H) considered. In particular it is possible to determine the relative probability of each of the ploidy hypotheses using the joint distribution model and the allele counts measured on the prepared sample, and using those relative probabilities to determine the hypothesis most likely to be correct as follows:

$$H^* = \underset{H}{\operatorname{argmax}} LIK(D | H)$$

[0646] Similarly the a posteriori hypothesis likelihood given the data may be written as:

$$H^* = \underset{H}{\operatorname{argmax}} LIK(D | H) * \operatorname{priorprob}(H)$$

[0647] Where $\operatorname{priorprob}(H)$ is the prior probability assigned to each hypothesis H , based on model design and prior knowledge.

[0648] It is also possible to use priors to find the maximum a posteriori estimate:

$$H_{MA} = \underset{H}{\operatorname{argmax}} LIK(D | H)$$

[0649] In an embodiment, the copy number hypotheses that may be considered are:

[0650] Monosomy:

[0651] maternal H10 (one copy from mother)

[0652] paternal H01 (one copy from father)

[0653] Disomy: H11 (one copy each mother and father)

[0654] Simple trisomy, no crossovers considered:

[0655] Maternal: H21_matched (two identical copies from mother, one copy from father),

[0656] H21_unmatched (BOTH copies from mother, one copy from father)

[0657] Paternal: H12_matched (one copy from mother, two identical copies from father),

[0658] H12_unmatched (one copy from mother, both copies from father)

[0659] Composite trisomy, allowing for crossovers (using a joint distribution model):

[0660] maternal H21 (two copies from mother, one from father),

[0661] paternal H12 (one copy from mother, two copies from father)

[0662] In other embodiments, other ploidy states, such as nullsomy (H00), uniparental disomy (H20 and H02), and tetrasomy (H04, H13, H22, H31 and H40), may be considered.

[0663] If there are no crossovers, each trisomy, whether the origin was mitosis, meiosis I, or meiosis II, would be one of the matched or unmatched trisomies. Due to crossovers, true trisomy is usually a combination of the two. First, a method to derive hypothesis likelihoods for simple hypotheses is described. Then a method to derive hypothesis likelihoods for composite hypotheses is described, combining individual SNP likelihood with crossovers.

$LIK(D|H)$ for a Simple Hypothesis

[0664] In an embodiment, $LIK(D|H)$ may be determined for simple hypotheses, as follows. For simple hypotheses H , $LIK(H)$, the log likelihood of hypothesis H on a whole chromosome, may be calculated as the sum of log likelihoods of individual SNPs, assuming known or derived child fraction cf . In an embodiment it is possible to derive cf from the data.

$$LIK(D | H) = \sum_i LIK(D | H, cf, i)$$

This hypothesis does not assume any linkage between SNPs, and therefore does not utilize a joint distribution model.

[0665] In some embodiments, the Log Likelihood may be determined on a per SNP basis. On a particular SNP i , assuming fetal ploidy hypothesis H and percent fetal DNA cf , log likelihood of observed data D is defined as:

$$LIK(D | H, i) = \log P(D | H, cf, i) =$$

$$\log \left(\sum_{m,f,c} P(D | m, f, c, H, cf, i) P(c | m, f, H) P(m | i) P(f | i) \right)$$

where m are possible true mother genotypes, f are possible true father genotypes, where $m, f \in \{AA, AB, BB\}$, and c are possible child genotypes given the hypothesis H . In particular, for monosomy $c \in \{A, B\}$, for disomy $c \in \{AA, AB, BB\}$, for trisomy $c \in \{AAA, AAB, ABB, BBB\}$. Genotype prior frequency: $p(mi)$ is the general prior probability of mother genotype m on SNP i , based on the known population frequency at SNP i , denoted pA_i . In particular

$$\begin{aligned} p(AA|pA_i) &= (pA_i)^2, p(AB|pA_i) = 2(pA_i)(1-pA_i), p \\ (BB|pA_i) &= (1-pA_i)^2 \end{aligned}$$

Father genotype probability, $p(fi)$, may be determined in an analogous fashion.

True child probability: $p(c|m, f, H)$ is the probability of getting true child genotype c , given parents m, f , and assuming hypothesis H , which can be easily calculated. For example, for H11, H21 matched and H21 unmatched, $p(c|m, f, H)$ is given below.

p(c m, f, H)												
m	f	H11			H21 matched				H21 unmatched			
		AA	AB	BB	AAA	AAB	ABB	BBB	AAA	AAB	ABB	BBB
AA	AA	1	0	0	1	0	0	0	1	0	0	0
AB	AA	0.5	0.5	0	0.5	0	0.5	0	0	1	0	0
BB	AA	0	1	0	0	0	1	0	0	0	1	0
AA	AB	0.5	0.5	0	0.5	0.5	0	0	0.5	0.5	0	0
AB	AB	0.25	0.5	0.25	0.25	0.25	0.25	0.25	0	0.5	0.5	0
BB	AB	0	0.5	0.5	0	0	0.5	0.5	0	0	0.5	0.5
AA	BB	0	1	0	0	1	0	0	0	1	0	0
AB	BB	0	0.5	0.5	0	0.5	0	0.5	0	0	1	0
BB	BB	0	0	1	0	0	0	1	0	0	0	1

Data likelihood: $P(D|m, f, c, H, i, cf)$ is the probability of given data D on SNP i , given true mother genotype m , true father genotype f , true child genotype c , hypothesis H and child fraction cf . It can be broken down into the probability of mother, father and child data as follows:

$$P(D|m, f, c, H, cf, i) = P(SM|m, i)P(M|m, i)P(SF|f, i)P(F|f, i)P(S|m, c, H, cf, i)$$

Mother SNP array data likelihood: Probability of mother SNP array genotype data m_i at SNP i compared to true genotype m , assuming SNP array genotypes are correct, is simply

$$P(M | m, i) = \begin{cases} 1 & m_i = m \\ 0 & m_i \neq m \end{cases}$$

Mother sequence data likelihood: the probability of the mother sequence data at SNP i , in the case of counts $S_i = (a_i, b_i)$, with no extra noise or bias involved, is the binomial probability defined as $P(SM|m, i) = P_{X|m}(a_i)$ where $X|m \sim \text{Binom}(p_m(A), a_i + b_i)$ with $p_m(A)$ defined as

m	AA	AB	BB	B	B	nocall
p(A)	1	0.5	0	1	0	0.5

Father data likelihood: a similar equation applies for father data likelihood.

Note that it is possible to determine the child genotype without the parent data, especially father data. For example if no father genotype data F is available, one may just use $P(F|f, i) = 1$. If no father sequence data SF is available, one may just use $P(SF|f, i) = 1$.

[0666] In some embodiments, the method involves building a joint distribution model for the expected allele counts at a plurality of polymorphic loci on the chromosome for each ploidy hypothesis; one method to accomplish such an end is described here. Free fetal DNA data likelihood: $P(S|m, c, H, cf, i)$ is the probability of free fetal DNA sequence data on SNP i , given true mother genotype m , true child genotype c , child copy number hypothesis H , and assuming child fraction cf . It is in fact the probability of sequence data S on SNP i , given the true probability of A content on SNP i $\mu(m, c, cf, H)$

$$P(S|m, c, H, cf, i) = P(S|\mu(m, c, cf, H), i)$$

For counts, where $S_i = (a_i, b_i)$, with no extra noise or bias in data involved,

$$P(S|\mu(m, c, cf, H), i) = P_X(a_i)$$

where $X \sim \text{Binom}(p(A), a_i + b_i)$ with $p(A) = \mu(m, c, cf, H)$. In a more complex case where the exact alignment and (A,B) counts per SNP are not known, $P(S|\mu(m, c, cf, H), i)$ is a combination of integrated binomials.

True A content probability: $p(m, c, cf, H)$, the true probability of A content on SNP i in this mother/child mixture, assuming that true mother genotype= m , true child genotype= c , and overall child fraction= cf , is defined as

$$\mu(m, c, cf, H) = \frac{\#A(m) * (1 - cf) + \#A(c) * cf}{n_m * (1 - cf) + n_c * cf}$$

where $\#A(g)$ =number of A's in genotype g , $n_m=2$ is some of mother and n_c is ploidy of the child under hypothesis H (1 for monosomy, 2 for disomy, 3 for trisomy).

Using a Joint Distribution Model: $\text{LIK}(D|H)$ for a Composite Hypothesis

[0667] In some embodiments, the method involves building a joint distribution model for the expected allele counts at the plurality of polymorphic loci on the chromosome for each ploidy hypothesis; one method to accomplish such an end is described here. In many cases, trisomy is usually not purely matched or unmatched, due to crossovers, so in this section results for composite hypotheses H21 (maternal trisomy) and H12 (paternal trisomy) are derived, which combine matched and unmatched trisomy, accounting for possible crossovers.

[0668] In the case of trisomy, if there were no crossovers, trisomy would be simply matched or unmatched trisomy. Matched trisomy is where child inherits two copies of the identical chromosome segment from one parent. Unmatched trisomy is where child inherits one copy of each homologous chromosome segment from the parent. Due to crossovers, some segments of a chromosome may have matched trisomy, and other parts may have unmatched trisomy. Described in this section is how to build a joint distribution model for the heterozygosity rates for a set of alleles; that is, for the expected allele counts at a number of loci for one or more hypotheses.

[0669] Suppose that on SNP i , $\text{LIK}(D|H_m, i)$ is the fit for matched hypothesis H_m , and $\text{LIK}(D|H_u, i)$ is the fit for

unmatched hypothesis H_u , and $pc(i)$ probability of crossover between SNPs $i-1$ and i . One may then calculate the full likelihood as:

$$LIK(D | H) = \sum_E LIK(D | E, 1:N)$$

where $LIK(D|E, 1:N)$ is the likelihood of ending in hypothesis E , for SNPs $1:N$. E =hypothesis of the last SNP, $E \in (H_m, H_u)$. Recursively, one may calculate:

$$LIK(D | E, 1:i) = LIK(D | E, i) + \log(\exp(LIK(D | E, 1:i-1)) * (1 - pc(i)) + \exp(LIK(D | \sim E, 1:i-1)) * pc(i))$$

where $\sim E$ is the hypothesis other than E (not E), where hypotheses considered are H_m and H_u . In particular, one may calculate the likelihood of $1:i$ SNPs, based on likelihood of 1 to $(i-1)$ SNPs with either the same hypothesis and no crossover, or the opposite hypothesis and a crossover, multiplied by the likelihood of the SNP i .

For SNP 1, $i=1$, $LIK(D|E, 1:1)=LIK(D|E, 1)$.

For SNP 2, $i=2$, $LIK(D|E, 1:2)=LIK(D|E, 2)+\log(\exp(LIK(D|E, 1))*(1-pc(2))+\exp(LIK(D|\sim E, 1))*pc(2))$, and so on for $i=3:N$.

[0670] In some embodiments, the child fraction may be determined. The child fraction may refer to the proportion of sequences in a mixture of DNA that originate from the child. In the context of non-invasive prenatal diagnosis, the child fraction may refer to the proportion of sequences in the maternal plasma that originate from the fetus or the portion of the placenta with fetal genotype. It may refer to the child fraction in a sample of DNA that has been prepared from the maternal plasma, and may be enriched in fetal DNA. One purpose of determining the child fraction in a sample of DNA is for use in an algorithm that can make ploidy calls on the fetus, therefore, the child fraction could refer to whatever sample of DNA was analyzed by sequencing for the purpose of non-invasive prenatal diagnosis.

[0671] Some of the algorithms presented in this disclosure that are part of a method of non-invasive prenatal aneuploidy diagnosis assume a known child fraction, which may not always be the case. In an embodiment, it is possible to find the most likely child fraction by maximizing the likelihood for disomy on selected chromosomes, with or without the presence of the parental data.

[0672] In particular, suppose that $LIK(D|H11, cf, chr)=\log$ likelihood as described above, for the disomy hypothesis, and for child fraction cf on chromosome chr . For selected chromosomes in $Cset$ (usually $1:16$), assumed to be euploid, the full likelihood is:

$$LIK(cf) = \sum_{chr \in Cset} LIK(D|H11, cf, chr)$$

The most likely child fraction (cf^*) is derived as

$$cf^* = \underset{cf}{\operatorname{argmax}} LIK(cf).$$

[0673] It is possible to use any set of chromosomes. It is also possible to derive child fraction without assuming

euploidy on the reference chromosomes. Using this method it is possible to determine the child fraction for any of the following situations: (1) one has array data on the parents and shotgun sequencing data on the maternal plasma; (2) one has array data on the parents and targeted sequencing data on the maternal plasma; (3) one has targeted sequencing data on both the parents and maternal plasma; (4) one has targeted sequencing data on both the mother and the maternal plasma fraction; (5) one has targeted sequencing data on the maternal plasma fraction; (6) other combinations of parental and child fraction measurements.

[0674] In some embodiments the informatics method may incorporate data dropouts; this may result in ploidy determinations of higher accuracy. Elsewhere in this disclosure it has been assumed that the probability of getting an A is a direct function of the true mother genotype, the true child genotype, the fraction of the child in the mixture, and the child copy number. It is also possible that mother or child alleles can drop out, for example instead of measuring true child AB in the mixture, it may be the case that only sequences mapping to allele A are measured. One may denote the parent dropout rate for genomic Illumina data d_{pg} , parent dropout rate for sequence data d_{ps} and child dropout rate for sequence data d_{cs} . In some embodiments, the mother dropout rate may be assumed to be zero, and child dropout rates are relatively low; in this case, the results are not severely affected by dropouts. In some embodiments the possibility of allele dropouts may be sufficiently large that they result in a significant effect of the predicted ploidy call. For such a case, allele dropouts have been incorporated into the algorithm here:

[0675] Parent SNP array data dropouts: For mother genomic data M , suppose that the genotype after the dropout is m_d , then

$$P(M | m, i) = \sum_{m_d} P(M | m_d, i) P(m_d | m)$$

where

$$P(m_d | m, i) = \begin{cases} 1 & m_i = m_d \\ 0 & m_i \neq m_d \end{cases}$$

as before, and $P(m_d|m)$ is the likelihood of genotype m_d after the possible dropout given the true genotype m , defined as below, for dropout rate d

m	md					nocall
	AA	AB	BB	A	B	
AA	$(1-d)^2$	0	0	$2d(1-d)$	0	d^2
AB	0	$(1-d)^2$	0	$d(1-d)$	$d(1-d)$	d^2

-continued

m	md					nocall
	AA	AB	BB	A	B	
BB	0	0	$(1 - d)^2$	0	$2d(1 - d)$	d^2

A similar equation applies for father SNP array data.
Parent sequence data dropouts: For mother sequence data SM

$$P(SM | m, i) = \sum_{m_d} P_{X|m_d}(am_i) P(m_d | m)$$

where $P(m_d | m)$ is defined as in previous section and $P_{X|m_d}(am_i)$ probability from a binomial distribution is defined as before in the parent data likelihood section. A similar equation applies to the paternal sequence data.

[0676] Free floating DNA sequence data dropout:

$$P(S | m, c, H, cf, i) = \sum_{m_d, c_d} P(S | \mu(m_d, c_d, cf, H), i) P(m_d | m) P(c_d | c)$$

where $P(S | m_d, c_d, cf, H, i)$ is as defined in the section on free floating data likelihood.

[0677] In an embodiment, $p(m_d | m)$ is the probability of observed mother genotype m , given true mother genotype m , assuming dropout rate d_{ps} , and $p(c_d | c)$ is the probability of observed child genotype c_d , given true child genotype c , assuming dropout rate d_{cs} . If nA_T =number of A alleles in true genotype c , nA_D =number of A alleles in observed genotype c_d , where $nA_T \leq nA_D$, and similarly nB_T =number of B alleles in true genotype c , nB_D =number of B alleles in observed genotype c_d , where $nB_T \geq nB_D$ and d =dropout rate, then

$$p(c_d | c) = \binom{nA_T}{nA_D} * d^{nA_T - nA_D} * (1 - d)^{nA_D} * \binom{nB_T}{nB_D} * d^{nB_T - nB_D} * (1 - d)^{nB_D}$$

[0678] In an embodiment, the informatics method may incorporate random and consistent bias.

[0679] In an ideal world there is no per SNP consistent sampling bias or random noise (in addition to the binomial distribution variation) in the number of sequence counts. In particular, on SNP i , for mother genotype m , true child genotype c and child fraction cf , and X =the number of A's in the set of (A+B) reads on SNP i , X acts like a X-Binomial ($p, A+B$), where $p=\mu(m, c, cf, H)$ =true probability of A content.

[0680] In an embodiment, the informatics method may incorporate random bias. As is often the case, suppose that there is a bias in the measurements, so that the probability of getting an A on this SNP is equal to q , which is a bit different than p as defined above. How much different p is from q depends on the accuracy of the measurement process and number of other factors and can be quantified by standard deviations of q away from p . In an embodiment, it is possible to model q as having a beta distribution, with

parameters α, β depending on the mean of that distribution being centered at p , and some specified standard deviation s . In particular, this gives $X | q \sim \text{Bin}(q, D_i)$, where $q \sim \text{Beta}(\alpha, \beta)$. If we let $E(q)=p, V(q)=s^2$, and parameters α, β can be derived as $\alpha=pN, \beta=(1-p)N$, where

$$N = \frac{p(1-p)}{s^2} - 1.$$

[0681] This is the definition of a beta-binomial distribution, where one is sampling from a binomial distribution with variable parameter q , where q follows a beta distribution with mean p . So, in a setup with no bias, on SNP i , the parent sequence data (SM) probability assuming true mother genotype (m), given mother sequence A count on SNP i (am_i) and mother sequence B count on SNP i (bm_i) may be calculated as:

$$P(SM | m, i) = P_{X|m}(am_i) \text{ where } X | m \sim \text{Binom}(p_m(A), am_i + bm_i)$$

Now, including random bias with standard deviation s , this becomes:

$$X | m \sim \text{BetaBinom}(p_m(A), am_i + bm_i, s)$$

[0682] In the case with no bias, the maternal plasma DNA sequence data (S) probability assuming true mother genotype (m), true child genotype (c), child fraction (cf), assuming child hypothesis H , given free floating DNA sequence A count on SNP i (a_i) and free floating sequence B count on SNP i (b_i) may be calculated as

$$P(S | m, c, cf, H, i) = P_X(a_i)$$

where $X \sim \text{Binom}(p(A), a_i + b_i)$ with $p(A) = \mu(m, c, cf, H)$.

[0683] In an embodiment, including random bias with standard deviation s , this becomes $X \sim \text{BetaBinom}(p(A), a_i + b_i, s)$, where the amount of extra variation is specified by the deviation parameter s , or equivalently N . The smaller the value of s (or the larger the value of N) the closer this distribution is to the regular binomial distribution. It is possible to estimate the amount of bias, i.e. estimate N above, from unambiguous contexts AA|AA, BB|BB, AA|BB, BB|AA and use estimated $\hat{N}$ in the above probability. Depending on the behavior of the data, N may be made to be a constant irrespective of the depth of read $a_i + b_i$, or a function of $a_i + b_i$, making bias smaller for larger depths of read.

[0684] In an embodiment, the informatics method may incorporate consistent per-SNP bias. Due to artifacts of the sequencing process, some SNPs may have consistently lower or higher counts irrespective of the true amount of A content. Suppose that SNP i consistently adds a bias of w_i percent to the number of A counts. In some embodiments, this bias can be estimated from the set of training data derived under same conditions, and added back in to the parent sequence data estimate as:

$$P(SM | m, i) = P_{X|m}(am_i) \text{ where } X | m \sim \text{BetaBinom}(p_m(A) + w_i, am_i + bm_i, s)$$

[0685] and with the free floating DNA sequence data probability estimate as:

$$P(S | m, c, cf, H, i) = P_X(a_i) \text{ where } X \sim \text{BetaBinom}(p(A) + w_i, a_i + b_i, s),$$

[0686] In some embodiments, the method may be written to specifically take into account additional noise, differential sample quality, differential SNP quality, and random sampling bias. An example of this is given here. This method has been shown to be particularly useful in the context of data generated using the massively multiplexed mini-PCR protocol, and was used in Examples 7 through 13. The method involves several steps that each introduce different kind of noise and/or bias to the final model:

[0687] Suppose the first sample that comprises a mixture of maternal and fetal DNA contains an original amount of DNA of size N_0 molecules, usually in the range 1,000-40,000, where p =true % refs

[0688] In the amplification using the universal ligation adaptors, assume that N_1 molecules are sampled; usually $N_1 \sim N_0/2$ molecules and random sampling bias is introduced due to sampling. The amplified sample may contain a number of molecules N_2 where $N_2 \gg N_1$. Let X_1 represent the amount of reference loci (on per SNP basis) out of N_1 sampled molecules, with a variation in $p_1 = X_1/N_1$ that introduces random sampling bias throughout the rest of protocol. This sampling bias is included in the model by using a Beta-Binomial (BB) distribution instead of using a simple Binomial distribution model. Parameter N of the Beta-Binomial distribution may be estimated later on per sample basis from training data after adjusting for leakage and amplification bias, on SNPs with $0 < p < 1$. Leakage is the tendency for a SNP to be read incorrectly.

[0689] The amplification step will amplify any allelic bias, thus amplification bias introduced due to possible uneven amplification. Suppose that one allele at a locus is amplified f times another allele at that locus is amplified g times, where $f = g e^b$, where $b=0$ indicates no bias. The bias parameter, b , is centered at 0, and indicates how much more or less the A allele get amplified as opposed to the B allele on a particular SNP. The parameter b may differ from SNP to SNP. Bias parameter b may be estimated on per SNP basis, for example from training data.

[0690] The sequencing step involves sequencing a sample of amplified molecules. In this step there may be leakage, where leakage is the situation where a SNP is read incorrectly. Leakage may result from any number of problems, and may result in a SNP being read not as the correct allele A, but as another allele B found at that locus or as an allele C or D not typically found at that locus. Suppose the sequencing measures the sequence data of a number of DNA molecules from an amplified sample of size N_3 , where $N_3 < N_2$. In some embodiments, N_3 may be in the range of 20,000 to 100,000; 100,000 to 500,000; 500,000 to 4,000,000; 4,000,000 to 20,000,000; or 20,000,000 to 100,000,000. Each molecule sampled has a probability p_g of being read correctly, in which case it will show up correctly as allele A. The sample will be incorrectly read as an allele unrelated to the original molecule with probability $1 - p_g$, and will look like allele A with probability p_r , allele B with probability p_m or allele C or allele D with probability p_o , where $p_r + p_m + p_o = 1$. Parameters p_g , p_r , p_m , p_o are estimated on per SNP basis from the training data.

[0691] Different protocols may involve similar steps with variations in the molecular biology steps resulting in different amounts of random sampling, different levels of amplification and different leakage bias. The following model may be equally well applied to each of these cases. The model for the amount of DNA sampled, on per SNP basis, is given by:

$$X_3 \sim \text{BetaBinomial}(L(F(p,b), p_r, p_g), N * H(p,b))$$

where p =the true amount of reference DNA, b =per SNP bias, and as described above, p_g is the probability of a correct read, p_r is the probability of read being read incorrectly but

serendipitously looking like the correct allele, in case of a bad read, as described above, and:

$$F(p,b) = p e^b / (p e^b + (1-p)), H(p,b) = (e^b p + (1-p))^2 / e^b, L(p, p_r, p_g) = p * p_g + p_r * (1 - p_g).$$

[0692] In some embodiments, the method uses a Beta-Binomial distribution instead of a simple binomial distribution; this takes care of the random sampling bias. Parameter N of the Beta-Binomial distribution is estimated on per sample basis on an as needed basis. Using bias correction $F(p,b)$, $H(p,b)$, instead of just p , takes care of the amplification bias. Parameter b of the bias is estimated on per SNP basis from training data ahead of time.

[0693] In some embodiments the method uses leakage correction $L(p, p_r, p_g)$, instead of just p ; this takes care of the leakage bias, i.e. varying SNP and sample quality. In some embodiments, parameters p_g , p_r , p_o are estimated on per SNP basis from the training data ahead of time. In some embodiments, the parameters p_g , p_r , p_o may be updated with the current sample on the go, to account for varying sample quality.

[0694] The model described herein is quite general and can account for both differential sample quality and differential SNP quality. Different samples and SNPs are treated differently, as exemplified by the fact that some embodiments use Beta-Binomial distributions whose mean and variance are a function of the original amount of DNA, as well as sample and SNP quality.

Platform Modeling

[0695] Consider a single SNP where the expected allele ratio present in the plasma is r (based on the maternal and fetal genotypes). The expected allele ratio is defined as the expected fraction of A alleles in the combined maternal and fetal DNA. For maternal genotype g_m and child genotype g_c , the expected allele ratio is given by equation 1, assuming that the genotypes are represented as allele ratios as well.

$$r = f g_c + (1-f) g_m \quad (1)$$

[0696] The observation at the SNP consists of the number of mapped reads with each allele present, n_a and n_b , which sum to the depth of read d . Assume that thresholds have already been applied to the mapping probabilities and phred scores such that the mappings and allele observations can be considered correct. A phred score is a numerical measure that relates to the probability that a particular measurement at a particular base is wrong. In an embodiment, where the base has been measured by sequencing, the phred score may be calculated from the ratio of the dye intensity corresponding to the called base to the dye intensity of the other bases. The simplest model for the observation likelihood is a binomial distribution which assumes that each of the d reads is drawn independently from a large pool that has allele ratio r . Equation 2 describes this model.

$$P(n_a, n_b | r) = p_{\text{bino}}(n_a; n_a + n_b, r) = \binom{n_a + n_b}{n_a} r^{n_a} (1-r)^{n_b} \quad (2)$$

[0697] The binomial model can be extended in a number of ways. When the maternal and fetal genotypes are either all A or all B, the expected allele ratio in plasma will be 0 or 1, and the binomial probability will not be well-defined. In practice, unexpected alleles are sometimes observed in prac-

tice. In an embodiment, it is possible to use a corrected allele ratio $\hat{r}=1/(n_a+n_b)$ to allow a small number of the unexpected allele. In an embodiment, it is possible to use training data to model the rate of the unexpected allele appearing on each SNP, and use this model to correct the expected allele ratio. When the expected allele ratio is not 0 or 1, the observed allele ratio may not converge with a sufficiently high depth of read to the expected allele ratio due to amplification bias or other phenomena. The allele ratio can then be modeled as a beta distribution centered at the expected allele ratio, leading to a beta-binomial distribution for $P(n_a, n_b|r)$ which has higher variance than the binomial.

[0698] The platform model for the response at a single SNP will be defined as $F(a, b, g_c, g_m, f)$ (3), or the probability of observing $n_a=a$ and $n_b=b$ given the maternal and fetal genotypes, which also depends on the fetal fraction through equation 1. The functional form of F may be a binomial distribution, beta-binomial distribution, or similar functions as discussed above.

$$F(a, b, g_c, g_m, f) = P(n_a = a, n_b = b | g_c, g_m, f) = P(n_a = a, n_b = b | r(g_c, g_m, f)) \quad (3)$$

[0699] In an embodiment, the child fraction may be determined as follows. A maximum likelihood estimate of the fetal fraction f for a prenatal test may be derived without the use of paternal information. This may be relevant where the paternal genetic data is not available, for example where the father of record is not actually the genetic father of the fetus. The fetal fraction is estimated from the set of SNPs where the maternal genotype is 0 or 1, resulting in a set of only two possible fetal genotypes. Define S_0 as the set of SNPs with maternal genotype 0 and S_1 as the set of SNPs with maternal genotype 1. The possible fetal genotypes on S_0 are 0 and 0.5, resulting in a set of possible allele ratios $R_0(f) = \{0, f/2\}$. Similarly, $R_1(f) = \{1-f/2, 1\}$. This method can be trivially extended to include SNPs where maternal genotype is 0.5, but these SNPs will be less informative due to the larger set of possible allele ratios.

[0700] Define N_{a0} and N_{b0} as the vectors formed by n_{as} and n_{bs} for SNPs s in S_0 , and N_{a1} and N_{b1} similarly for S_1 . The maximum likelihood estimate $\hat{f}$ of f is defined by equation 4.

$$\hat{f} = \arg \max_f P(N_{a0}, N_{b0} | f) P(N_{a1}, N_{b1} | f) \quad (4)$$

[0701] Assuming that the allele counts at each SNP are independent conditioned on the SNP's plasma allele ratio, the probabilities can be expressed as products over the SNPs in each set (5).

$$P(N_{a0}, N_{b0} | f) = \prod_{s \in S_0} P(n_{as}, n_{bs} | f) \\ P(N_{a1}, N_{b1} | f) = \prod_{s \in S_1} P(n_{as}, n_{bs} | f) \quad (5)$$

[0702] The dependence on f is through the sets of possible allele ratios $R_0(f)$ and $R_1(f)$. The SNP probability $P(n_{as}, n_{bs} | f)$ can be approximated by assuming the maximum likelihood genotype conditioned on f . At reasonably high fetal fraction and depth of read, the selection of the maximum likelihood genotype will be high confidence. For example, at fetal fraction of 10 percent and depth of read of 1000, consider a SNP where the mother has genotype zero. The expected allele ratios are 0 and 5 percent, which will be easily distinguishable at sufficiently high depth of read. Substitution of the estimated child genotype into equation 5 results in the complete equation (6) for the fetal fraction estimate.

$$\hat{f} = \arg \max_f \left[\prod_{s \in S_0} \left(\max_{r_s \in R_0(f)} P(n_{as}, n_{bs} | r_s) \right) \prod_{s \in S_1} \left(\max_{r_s \in R_1(f)} P(n_{as}, n_{bs} | r_s) \right) \right] \quad (6)$$

[0703] The fetal fraction must be in the range $[0, 1]$ and so the optimization can be easily implemented by a constrained one-dimensional search.

[0704] In the presence of low depth of read or high noise level, it may be preferable not to assume the maximum likelihood genotype, which may result in artificially high confidences.

[0705] Another method would be to sum over the possible genotypes at each SNP, resulting in the following expression (7) for $P(n_a, n_b | f)$ for a SNP in S_0 . The prior probability $P(r)$ could be assumed uniform over $R_0(f)$, or could be based on population frequencies. The extension to group S_1 is trivial.

$$P(n_a, n_b | f) = \sum_{r \in R_0(f)} P(n_a, n_b | r) P(r) \quad (7)$$

[0706] In some embodiments the probabilities may be derived as follows. A confidence can be calculated from the data likelihoods of the two hypotheses H_t and H_f . The likelihood of each hypothesis is derived based on the response model, the estimated fetal fraction, the mother genotypes, allele population frequencies, and the plasma allele counts.

Define the following notation:

[0707] G_m, G_c true maternal and child genotypes

[0708] G_{af}, G_{tf} true genotypes of alleged father and of true father

[0709] $G(g_c, g_m, g_{tf}) = P(G_c = g_c | G_m = g_m, G_{tf} = g_{tf})$ inheritance probabilities

[0710] $P(g) = P(G_{tf} = g)$ population frequency of genotype g at particular SNP

[0711] Assuming that the observation at each SNP is independent conditioned on the plasma allele ratio, the likelihood of a paternity hypothesis is the product of the likelihoods on the SNPs. The following equations derive the likelihood for a single SNP. Equation 8 is a general expression for the likelihood of any hypothesis h , which will then be broken down into the specific cases of H_t and H_f .

$$P(n_a, n_b | h, G_m, G_{tf}, f) = \sum_{g_c \in (0, 0.5, 1)} P(n_a, n_b | G_c = g_c, G_m, G_{tf}, h, f) P(G_c = g_c | G_m, G_{tf}, h, f) = \sum_{g_c \in (0, 0.5, 1)} P(n_a, n_b | G_c = g_c, G_m, f) P(G_c = g_c | G_m, G_{tf}, h) \\ = \sum_{g_c \in (0, 0.5, 1)} F(n_a, n_b | g_c, g_m, f) P(G_c = g_c | G_m, G_{tf}, h) \quad (8)$$

[0712] In the case of H_t , the alleged father is the true father and the fetal genotypes are inherited from the maternal genotypes and alleged father genotypes according to equation 9.

$$P(n_a, n_b | H_t, G_m, G_{tf}, f) = \sum_{g_c \in (0, 0.5, 1)} F(n_a, n_b | g_c, g_m, f) P(G_c = g_c | G_m, G_{tf}, H_t, f) = \sum_{g_c \in (0, 0.5, 1)} F(n_a, n_b | g_c, g_m, f) P(g_c, G_m, G_{tf}) \quad (9)$$

[0713] In the case of H_f , the alleged father is not the true father. The best estimate of the true father genotypes are given by the population frequencies at each SNP. Thus, the probabilities of child genotypes are determined by the known mother genotypes and the population frequencies, as in equation 10.

$$\begin{aligned}
P(n_a, n_b \mid H_t, G_m, G_{tf}, f) &= \sum_{g_c \in (0,0.5,1)} F(n_a, n_b, g_c, g_m, f) \\
&\quad P(G_c = g_c \mid G_m, G_{tf}, H_f) \\
&= \sum_{g_c \in (0,0.5,1)} F(n_a, n_b, g_c, g_m, f) \\
&\quad P(G_c = g_c \mid G_m) \\
&= \sum_{g_c \in (0,0.5,1)} \sum_{g_{tf} \in (0,0.5,1)} F(n_a, \\
&\quad n_b, g_c, g_m, f) P(G_c = g_c \mid G_m, \\
&\quad G_{tf} = g_{tf}) P(G_{tf} = g_{tf}) \\
&= \sum_{g_c \in (0,0.5,1)} \sum_{g_{tf} \in (0,0.5,1)} F(n_a, \\
&\quad n_b, g_c, g_m, f) G(g_c, G_m, g_{tf}) P(g_{tf})
\end{aligned}$$

[0714] The confidence C_p on correct paternity is calculated from the product over SNPs of the two likelihoods using Bayes rule (11).

$$C_p = \frac{\Pi_s P(n_{as}, n_{bs} \mid H_t, G_{ms}, G_{tf}, f)}{\Pi_s P(n_{as}, n_{bs} \mid H_t, G_{ms}, G_{tf}, f) + \Pi_s P(n_{as}, n_{bs} \mid H_f, G_{ms}, G_{tf}, f)} \quad (11)$$

Exemplary Methods for Identifying and Analyzing Multiple Pregnancies

[0715] In some embodiments, any of the methods of the present invention are used to detect the presence of a multiple pregnancy, such as a twin pregnancy, where at least one of the fetuses is genetically different from at least one other fetus. In some embodiments, fraternal twins are identified based on the presence of two fetus with different allele, different allele ratios, or different allele distributions at some (or all) of the tested loci. In some embodiments, fraternal twins are identified by determining the expected allele ratio at each locus (such as SNP loci) for two fetuses that may have the same or different fetal fractions in the sample (such as a plasma sample). In some embodiments, the likelihood of a particular pair of fetal fractions (where f1 is the fetal fraction for fetus 1, and f2 is the fetal fraction for fetus 2) is calculated by considering some or all of the possible genotypes of the two fetuses, conditioned on the mother's genotype and genotype population frequencies. The mixture of two fetal and one maternal genotype, combined with the fetal fractions, determine the expected allele ratio at a SNP. For example, if the mother is AA, fetus 1 is AA, and fetus 2 is AB, the overall fraction of B allele at the SNP is one-half of f2. The likelihood calculation asks how well all of the SNPs together match the expected allele ratios based on all of the possible combinations of fetal genotypes. The fetal fraction pair (f1, f2) that best matches the data is selected. It is not necessary to calculated specific genotypes of the fetuses; instead, one can, for example, considered all of the possible genotypes in a statistical combination. In some embodiments, if the method does not distinguish between singleton and identical twins, an ultrasound can be performed to determine whether there is a singleton or identical twin pregnancy. If the ultrasound detects a twin pregnancy it can be assumed that the pregnancy is an identical twin pregnancy because a fraternal twin pregnancy would have been detected based on the SNP analysis discussed above.

[0716] In some embodiments, a pregnant mother is known to have a multiple pregnancy (such as a twin pregnancy) based on prior testing, such as an ultrasound. Any of the methods of the present invention can be used to determine whether the multiple pregnancy includes identical or fraternal twins. For example, the measured allele ratios can be compared to what would be expected for identical twins (the same allele ratios as a singleton pregnancy) or for fraternal twins (such as the calculation of allele ratios as described above). Some identical twins are monochorionic twins, which have a risk of twin-to-twin transfusion syndrome. Thus, twins determined to be identical twins using a method of the invention are desirably tested (such as by ultrasound) to determine if they are monochorionic twins, and if so, these twins can be monitored (such as bi-weekly ultrasounds from 16 weeks) for signs of win-to-twin transfusion syndrome.

[0717] In some embodiments, any of the methods of the present invention are used to determine whether any of the fetuses in a multiple pregnancy, such as a twin pregnancy, are aneuploid. Aneuploidy testing for twins begins with the fetal fraction estimate. In some embodiments, the fetal fraction pair (f1, f2) that best matches the data is selected as described above. In some embodiments, a maximum likelihood estimate is performed for the parameter pair (f1, f2) over the range of possible fetal fractions. In some embodiments, the range of f2 is from 0 to f1 because f2 is defined as the smaller fetal fraction. Given a pair (f1, f2), data likelihood is calculated from the allele ratios observed at a set of loci such as SNP loci. In some embodiments, the data likelihood reflects the genotypes of the mother, the father if available, population frequencies, and the resulting probabilities of fetal genotypes. In some embodiments, SNPs are assumed independent. The estimated fetal fraction pair is the one that produces the highest data likelihood. If f2 is 0 then the data is best explained by only one set of fetal genotypes, indicating identical twins, where f1 is the combined fetal fraction. Otherwise f1 and f2 are the estimates of the individual twin fetal fractions. Having established the best estimate of (f1, f2), one can predict the overall fraction of B allele in the plasma for any combination of maternal and fetal genotypes, if desired. It is not necessary to assign individual sequence reads to the individual fetuses. Ploidy testing is performed using another maximum likelihood estimate which compares the data likelihood of two hypotheses. In some embodiments for identical twins, one consider the hypotheses (i) both twins are euploid, and (ii) both twins are trisomic. In some embodiments for fraternal twins, one considers the hypotheses (i) both twins are euploid and (ii) at least one twin is trisomic. The trisomy hypotheses for fraternal twins are based on the lower fetal fraction, since a trisomy in the twin with a higher fetal fraction would also be detected. Ploidy likelihoods are calculated using a method which predicts the expected number of reads at each targeted genome locus conditioned on either the disomy or trisomy hypothesis. There is no requirement for a disomy reference chromosome. The variance model for the expected number of reads takes into account the performance of individual target loci as well as the correlation between loci (see, for example, U.S. Ser. No. 62/008,235, filed Jun. 5, 2014, and U.S. Ser. No. 62/032,785, filed Aug. 4, 2014, which are each hereby incorporated by reference in its entirety). If the smaller twin has fetal fraction f1, our ability to detect a trisomy in that twin is equivalent to our ability to detect a

trisomy in a singleton pregnancy at the same fetal fraction. This is because the part of the method that detects the trisomy in some embodiments does not depend on genotypes and does not distinguish between multiple or singleton pregnancy. It simply looks for an increased number of reads in accordance with the determined fetal fraction.

[0718] In some embodiments, the method includes detecting the presence of twins based on SNP loci (such as described above). If twins are detected, SPNs are used to determine the fetal fraction of each fetus (f1, f2) such as described above. In some embodiments, samples that have high confidence disomy calls are used to determine the amplification bias on a per-SNP basis. In some embodiments, these samples with high confidence disomy calls are analyzed in the same run as one or more samples of interest. In some embodiments, the amplification bias on a per-SNP basis is used to model the distribution of reads for one or more chromosomes or chromosome segments of interest such as chromosome 21 that are expected or the disomy hypothesis and the trisomy hypothesis given the lower of the two twin fetal fraction. The likelihood or probability of disomy or trisomy is calculated given the two models and the measured quantity of the chromosome or chromosome segment of interest.

[0719] In some embodiments, the threshold for a positive aneuploidy call (such as a trisomy call) is set based on the twin with the lower fetal fraction. This way, if the other twin is positive, or if both are positive, the total chromosome representation is definitely above the threshold.

Maximum Likelihood Model using Percent Fetal Fraction

[0720] Determining the ploidy status of a fetus by measuring the free floating DNA contained in maternal serum, or by measuring the genotypic material in any mixed sample, is a non-trivial exercise. There are a number of methods, for example, performing a read count analysis where the presumption is that if the fetus is trisomic at a particular chromosome, then the overall amount of DNA from that chromosome found in the maternal blood will be elevated with respect to a reference chromosome. One way to detect trisomy in such fetuses is to normalize the amount of DNA expected for each chromosome, for example, according to the number of SNPs in the analysis set that correspond to a given chromosome, or according to the number of uniquely mappable portions of the chromosome. Once the measurements have been normalized, any chromosomes for which the amount of DNA measured exceeds a certain threshold are determined to be trisomic. This approach is described in Fan, et al. PNAS, 2008; 105(42); pp. 16266-16271, and also in Chiu et al. BMJ 2011; 342:c7401. In the Chiu et al. paper, the normalization was accomplished by calculating a Z score as follows:

$$\text{Z score for percentage chromosome 21 in test case} = \frac{(\text{percentage chromosome 21 in test case}) - (\text{mean percentage chromosome 21 in reference controls})}{(\text{standard deviation of percentage chromosome 21 in reference controls})}$$

[0721] These methods determine the ploidy status of the fetus using a single hypothesis rejection method. However, they suffer from some significant shortcomings. Since these methods for determining ploidy in the fetus are invariant according to the percentage of fetal DNA in the sample, they use one cut off value; the result of this is that the accuracies of the determinations are not optimal, and those cases where

the percentage of fetal DNA in the mixture are relatively low will suffer the worst accuracies.

[0722] In an embodiment, a method of the present disclosure is used to determine the ploidy state of the fetus involves taking into account the fraction of fetal DNA in the sample. In another embodiment of the present disclosure, the method involves the use of maximum likelihood estimations. In an embodiment, a method of the present disclosure involves calculating the percent of DNA in a sample that is fetal or placental in origin. In an embodiment, the threshold for calling aneuploidy is adaptively adjusted based on the calculated percent fetal DNA. In some embodiments, the method for estimating the percentage of DNA that is of fetal origin in a mixture of DNA, comprises obtaining a mixed sample that comprises genetic material from the mother, and genetic material from the fetus, obtaining a genetic sample from the father of the fetus, measuring the DNA in the mixed sample, measuring the DNA in the father sample, and calculating the percentage of DNA that is of fetal origin in the mixed sample using the DNA measurements of the mixed sample, and of the father sample.

[0723] In an embodiment of the present disclosure, the fraction of fetal DNA, or the percentage of fetal DNA in the mixture can be measured. In some embodiments the fraction can be calculated using only the genotyping measurements made on the maternal plasma sample itself, which is a mixture of fetal and maternal DNA. In some embodiments the fraction may be calculated also using the measured or otherwise known genotype of the mother and/or the measured or otherwise known genotype of the father. In some embodiments the percent fetal DNA may be calculated using the measurements made on the mixture of maternal and fetal DNA along with the knowledge of the parental contexts. In an embodiment, the fraction of fetal DNA may be calculated using population frequencies to adjust the model on the probability on particular allele measurements.

[0724] In an embodiment of the present disclosure, a confidence may be calculated on the accuracy of the determination of the ploidy state of the fetus. In an embodiment, the confidence of the hypothesis of greatest likelihood (H_{major}) may be calculated as $(1 - H_{major}) / \sum(\text{all } H)$. It is possible to determine the confidence of a hypothesis if the distributions of all of the hypotheses are known. It is possible to determine the distribution of all of the hypotheses if the parental genotype information is known. It is possible to calculate a confidence of the ploidy determination if the knowledge of the expected distribution of data for the euploid fetus and the expected distribution of data for the aneuploid fetus are known. It is possible to calculate these expected distributions if the parental genotype data are known. In an embodiment one may use the knowledge of the distribution of a test statistic around a normal hypothesis and around an abnormal hypothesis to determine both the reliability of the call as well as refine the threshold to make a more reliable call. This is particularly useful when the amount and/or percent of fetal DNA in the mixture is low. It will help to avoid the situation where a fetus that is actually aneuploid is found to be euploid because a test statistic, such as the Z statistic does not exceed a threshold that is made based on a threshold that is optimized for the case where there is a higher percent fetal DNA.

[0725] In an embodiment, a method disclosed herein can be used to determine a fetal aneuploidy by determining the number of copies of maternal and fetal target chromosomes

in a mixture of maternal and fetal genetic material. This method may entail obtaining maternal tissue comprising both maternal and fetal genetic material; in some embodiments this maternal tissue may be maternal plasma or a tissue isolated from maternal blood. This method may also entail obtaining a mixture of maternal and fetal genetic material from said maternal tissue by processing the aforementioned maternal tissue. This method may entail distributing the genetic material obtained into a plurality of reaction samples, to randomly provide individual reaction samples that comprise a target sequence from a target chromosome and individual reaction samples that do not comprise a target sequence from a target chromosome, for example, performing high throughput sequencing on the sample. This method may entail analyzing the target sequences of genetic material present or absent in said individual reaction samples to provide a first number of binary results representing presence or absence of a presumably euploid fetal chromosome in the reaction samples and a second number of binary results representing presence or absence of a possibly aneuploid fetal chromosome in the reaction samples. Either of the number of binary results may be calculated, for example, by way of an informatics technique that counts sequence reads that map to a particular chromosome, to a particular region of a chromosome, to a particular locus or set of loci. This method may involve normalizing the number of binary events based on the chromosome length, the length of the region of the chromosome, or the number of loci in the set. This method may entail calculating an expected distribution of the number of binary results for a presumably euploid fetal chromosome in the reaction samples using the first number. This method may entail calculating an expected distribution of the number of binary results for a presumably aneuploid fetal chromosome in the reaction samples using the first number and an estimated fraction of fetal DNA found in the mixture, for example, by multiplying the expected read count distribution of the number of binary results for a presumably euploid fetal chromosome by $(1+n/2)$ where n is the estimated fetal fraction. In some embodiments, the sequence reads may be treated at probabilistic mappings rather than binary results; this method would yield higher accuracies, but require more computing power. The fetal fraction may be estimated by a plurality of methods, some of which are described elsewhere in this disclosure. This method may involve using a maximum likelihood approach to determine whether the second number corresponds to the possibly aneuploid fetal chromosome being euploid or being aneuploid. This method may involve calling the ploidy status of the fetus to be the ploidy state that corresponds to the hypothesis with the maximum likelihood of being correct given the measured data.

[0726] Note that the use of a maximum likelihood model may be used to increase the accuracy of any method that determines the ploidy state of a fetus. Similarly, a confidence maybe calculated for any method that determines the ploidy state of the fetus. The use of a maximum likelihood model would result in an improvement of the accuracy of any method where the ploidy determination is made using a single hypothesis rejection technique. A maximum likelihood model may be used for any method where a likelihood distribution can be calculated for both the normal and abnormal cases. The use of a maximum likelihood model implies the ability to calculate a confidence for a ploidy call.

[0727] Further Discussion of the Method

[0728] In an embodiment, a method disclosed herein utilizes a quantitative measure of the number of independent observations of each allele at a polymorphic locus, where this does not involve calculating the ratio of the alleles. This is different from methods, such as some microarray based methods, which provide information about the ratio of two alleles at a locus but do not quantify the number of independent observations of either allele. Some methods known in the art can provide quantitative information regarding the number of independent observations, but the calculations leading to the ploidy determination utilize only the allele ratios, and do not utilize the quantitative information. To illustrate the importance of retaining information about the number of independent observations consider the sample locus with two alleles, A and B. In a first experiment twenty A alleles and twenty B alleles are observed, in a second experiment 200 A alleles and 200 B alleles are observed. In both experiments the ratio $(A/(A+B))$ is equal to 0.5, however the second experiment conveys more information than the first about the certainty of the frequency of the A or B allele. The instant method, rather than utilizing the allele ratios, uses the quantitative data to more accurately model the most likely allele frequencies at each polymorphic locus.

[0729] In an embodiment, the instant methods build a genetic model for aggregating the measurements from multiple polymorphic loci to better distinguish trisomy from disomy and also to determine the type of trisomy. Additionally, the instant method incorporates genetic linkage information to enhance the accuracy of the method. This is in contrast to some methods known in the art where allele ratios are averaged across all polymorphic loci on a chromosome. The method disclosed herein explicitly models the allele frequency distributions expected in disomy as well as and trisomy resulting from nondisjunction during meiosis I, nondisjunction during meiosis II, and nondisjunction during mitosis early in fetal development. To illustrate why this is important, if there were no crossovers nondisjunction during meiosis I would result a trisomy in which two different homologs were inherited from one parent; nondisjunction during meiosis II or during mitosis early in fetal development would result in two copies of the same homolog from one parent. Each scenario results in different expected allele frequencies at each polymorphic locus and also at all physically linked loci (i.e. loci on the same chromosome) considered jointly. Crossovers, which result in the exchange of genetic material between homologs, make the inheritance pattern more complex, but the instant method accommodates for this by using genetic linkage information, i.e. recombination rate information and the physical distance between loci. To better distinguish between meiosis I nondisjunction and meiosis II or mitotic nondisjunction the instant method incorporates into the model an increasing probability of crossover as the distance from the centromere increases. Meiosis II and mitotic nondisjunction can distinguished by the fact that mitotic nondisjunction typically results in identical or nearly identical copies of one homolog while the two homologs present following a meiosis II nondisjunction event often differ due to one or more crossovers during gametogenesis.

[0730] In an embodiment, a method of the present disclosure may not determine the haplotypes of the parents if disomy is assumed. In an embodiment, in case of trisomy, the instant method can make a determination about the

haplotypes of one or both parents by using the fact that plasma takes two copies from one parent, and parent phase information can be determined by noting which two copies have been inherited from the parent in question. In particular, a child can inherit either two of the same copies of the parent (matched trisomy) or both copies of the parent (unmatched trisomy). At each SNP one can calculate the likelihood of the matched trisomy and of the unmatched trisomy. A ploidy calling method that does not use the linkage model accounting for crossovers would calculate the overall likelihood of the trisomy as a simple weighted average of the matched and unmatched trisomies over all chromosomes. However, due to the biological mechanisms that result in disjunction error and crossing over, trisomy can change from matched to unmatched (and vice versa) on a chromosome only if a crossover occurs. The instant method probabilistically takes into account the likelihood of crossover, resulting in ploidy calls that are of greater accuracy than those methods that do not.

[0731] In an embodiment, a reference chromosome is used to determine the child fraction and noise level amount or probability distribution. In an embodiment, the child fraction, noise level, and/or probability distribution is determined using only the genetic information available from the chromosome whose ploidy state is being determined. The instant method works without the reference chromosome, as well as without fixing the particular child fraction or noise level. This is a significant improvement and point of differentiation from methods known in the art where genetic data from a reference chromosome is necessary to calibrate the child fraction and chromosome behavior.

[0732] In an embodiment where a reference chromosome is not needed to determine the fetal fraction, determining the hypothesis is done as follows:

$$H^* = \underset{H}{\operatorname{argmax}} \operatorname{LIK}(D | H) * \operatorname{priorprob}(H)$$

[0733] With the algorithm with reference chromosome, one typically assumes that the reference chromosome is a disomy, and then one may either (a) fix the most likely child fraction and random noise level N based on this assumption and reference chromosome data:

$$[cfr^*, N^*] = \underset{cfr, N}{\operatorname{argmax}} \operatorname{LIK}(D(\operatorname{ref.chrom}) | H11, cfr, N)$$

And then reduce

$$\operatorname{LIK}(D|H) = \operatorname{LIK}(D|H, cfr^*, N^*)$$

or (b) estimate the child fraction and noise level distribution based on this assumption and reference chromosome data. In particular, one would not fix just one value for cfr and N, but assign probability $p(cfr, N)$ for the wider range of possible cfr, N values:

$$p(cfr, N) \sim \operatorname{LIK}(D(\operatorname{ref.chrom}) | H11, cfr, N) * \operatorname{priorprob}(cfr, N)$$

where $\operatorname{priorprob}(cfr, N)$ is the prior probability of particular child fraction and noise level, determined by prior knowledge and experiments. If desired, just uniform over the range of cfr, N. One may then write:

$$\operatorname{LIK}(D | H) = \sum_{cfr, N} \operatorname{LIK}(D | H, cfr, N) * p(cfr, N)$$

[0734] Both methods above give good results.

[0735] Note that in some instances using a reference chromosome is not desirable, possible or feasible. In such a case, it is possible to derive the best ploidy call for each chromosome separately. In particular:

$$\operatorname{LIK}(D | H) = \sum_{cfr, N} \operatorname{LIK}(D | H, cfr, N) * p(cfr, N | H)$$

[0736] $p(cfr, N | H)$ may be determined as above, for each chromosome separately, assuming hypothesis H, not just for the reference chromosome assuming disomy. It is possible, using this method, to keep both noise and child fraction parameters fixed, fix either of the parameters, or keep both parameters in probabilistic form for each chromosome and each hypothesis.

[0737] Measurements of DNA are noisy and/or error prone, especially measurements where the amount of DNA is small, or where the DNA is mixed with contaminating DNA. This noise results in less accurate genotypic data, and less accurate ploidy calls. In some embodiments, platform modeling or some other method of noise modeling may be used to counter the deleterious effects of noise on the ploidy determination. The instant method uses a joint model of both channels, which accounts for the random noise due to the amount of input DNA, DNA quality, and/or protocol quality.

[0738] This is in contrast to some methods known in the art where the ploidy determinations are made using the ratio of allele intensities at a locus. This method precludes accurate SNP noise modeling. In particular, errors in the measurements typically do not specifically depend on the measured channel intensity ratio, which reduces the model to using one-dimensional information. Accurate modeling of noise, channel quality and channel interaction requires a two-dimensional joint model, which can not be modeled using allele ratios.

[0739] In particular, projecting two channel information to the ratio r where $f(x, y)$ is $r = x/y$, does not lend itself to accurate channel noise and bias modeling. Noise on a particular SNP is not a function of the ratio, i.e. $\operatorname{noise}(x, y) \neq f(x, y)$ but is in fact a joint function of both channels. For example, in the binomial model, noise of the measured ratio has a variance of $r(1-r)/(x+y)$ which is not a function purely of r . In such a model, where any channel bias or noise is included, suppose that on SNP i , the observed channel X value is $x = a_i X + b_i$, where X is the true channel value, b_i is the extra channel bias and random noise. Similarly, suppose that $y = c_i Y + d_i$. The observed ratio $r = x/y$ can not accurately predict the true ratio X/Y or model the leftover noise, since $(a_i X + b_i)/(c_i Y + d_i)$ is not a function of X/Y .

[0740] The method disclosed herein describes an effective way to model noise and bias using joint binomial distributions of all of the measurement channels individually. Relevant equations may be found elsewhere in the document in sections which speaks of per SNP consistent bias, $P(\text{good})$ and $P(\text{reflbad})$, $P(\text{mutlbad})$ which effectively adjust SNP behavior. In an embodiment, a method of the present disclosure uses a BetaBinomial distribution, which avoids the

limiting practice of relying on the allele ratios only, but instead models the behavior based on both channel counts.

[0741] In an embodiment, a method disclosed herein can call the ploidy of a gestating fetus from genetic data found in maternal plasma by using all available measurements. In an embodiment, a method disclosed herein can call the ploidy of a gestating fetus from genetic data found in maternal plasma by using the measurements from only a subset of parental contexts. Some methods known in the art only use measured genetic data where the parental context is from the AAIBB context, that is, where the parents are both homozygous at a given locus, but for a different allele. One problem with this method is that a small proportion of polymorphic loci are from the AAIBB context, typically less than 10%. In an embodiment of a method disclosed herein, the method does not use genetic measurements of the maternal plasma made at loci where the parental context is AAIBB. In an embodiment, the instant method uses plasma measurements for only those polymorphic loci with the AAIBB, ABIAA, and ABIAB parental context.

[0742] Some methods known in the art involve averaging allele ratios from SNPs in the AAIBB context, where both parent genotypes are present, and claim to determine the ploidy calls from the average allele ratio on these SNPs. This method suffers from significant inaccuracy due differential SNP behavior. Note that this method assumes that have both parent genotypes are known. In contrast, in some embodiments, the instant method uses a joint channel distribution model that does not assume the presence of either of the parents, and does not assume the uniform SNP behavior. In some embodiments, the instant method accounts for the different SNP behavior/weighting. In some embodiments, the instant method does not require the knowledge of one or both parental genotypes. An example of how the instant method may accomplish this follows:

[0743] In some embodiments, the log likelihood of a hypothesis may be determined on a per SNP basis. On a particular SNP i , assuming fetal ploidy hypothesis H and percent fetal DNA cf , the log likelihood of observed data D is defined as:

$$LIK(D | H, i) = \log P(D | H, cf, i) = \log \left(\sum_{m,f,c} P(D | m, f, c, H, cf, i) P(c | m, f, H) P(m | i) P(f | i) \right)$$

[0744] where m are possible true mother genotypes, f are possible true father genotypes, where $m, f \in \{AA, AB, BB\}$, and where c are possible child genotypes given the hypothesis H . In particular, for monosomy $c \in \{A, B\}$, for disomy $c \in \{AA, AB, BB\}$, for trisomy $c \in \{AAA, AAB, ABB, BBB\}$. Note that including parental genotypic data typically results in more accurate ploidy determinations, however, parental genotypic data is not necessary for the instant method to work well.

[0745] Some methods known in the art involve averaging allele ratios from SNPs where the mother is homozygous but a different allele is measured in the plasma (either AAIBB or AAIBB contexts), and claim to determine the ploidy calls from the average allele ratio on these SNPs. This method is intended for cases where the paternal genotype is not available. Note that it is questionable how accurately one can claim that plasma is heterozygous on a particular SNP

without the presence of homozygous and opposite father BB: for cases with low child fraction, what looks like presence of B allele could be just presence of noise; additionally, what looks like no B present could be simple allele drop out of the fetal measurements. Even in a case where one can actually determine heterozygosity of the plasma, this method will not be able to distinguish paternal trisomies. In particular, for SNPs where mother is AA, and where some B is measured in the plasma, if the father is GG, the resulting child genotype is ΔGG , resulting in an average ratio of 33% A (for child fraction=100%). But in the case where the father is ΔG , the resulting child genotype could be ΔGG for matched trisomy, contributing to the 33% A ratio, or AAG for unmatched trisomy, drawing the average ratio more toward 66% A. Given that many trisomies are on chromosomes with crossovers, the overall chromosome can have anywhere between no unmatched trisomy and all unmatched trisomy, this ratio can vary anywhere between 33-66%. For a plain disomy, the ratio should be around 50%. Without the use of a linkage model or an accurate error model of the average, this method would miss many cases of paternal trisomy. In contrast, the method disclosed herein assigns parental genotype probabilities for each parental genotypic candidate, based on available genotypic information and population frequency, and does not explicitly require parental genotypes. Additionally, the method disclosed herein is able to detect trisomy even in the absence or presence of parent genotypic data, and can compensate by identifying the points of possible crossovers from matched to unmatched trisomy using a linkage model.

[0746] Some methods known in the art claim a method for averaging allele ratios from SNPs where neither the maternal or paternal genotype is known, and for determining the ploidy calls from average ratio on these SNPs. However, a method to accomplish these ends is not disclosed. The method disclosed herein is able to make accurate ploidy calls in such a situation, and the reduction to practice is disclosed elsewhere in this document, using a joint probability maximum likelihood method and optionally utilizes SNP noise and bias models, as well as a linkage model.

[0747] Some methods known in the art involve averaging allele ratios and claim to determine the ploidy calls from the average allele ratio at one or a few SNPs. However, such methods do not utilize the concept of linkage. The methods disclosed herein do not suffer from these drawbacks.

Using Sequence Length as a Prior to Determine the Origin of DNA

[0748] It has been reported that the distribution of length of sequences differ for maternal and fetal DNA, with fetal generally being shorter. In an embodiment of the present disclosure, it is possible to use previous knowledge in the form of empirical data, and construct prior distribution for expected length of both mother ($P(X|maternal)$) and fetal DNA ($P(X|fetal)$). Given new unidentified DNA sequence of length x , it is possible to assign a probability that a given sequence of DNA is either maternal or fetal DNA, based on prior likelihood of x given either maternal or fetal. In particular if $P(x|maternal) > P(x|fetal)$, then the DNA sequence can be classified as maternal, with $P(x|maternal) = P(x|maternal) / [P(x|maternal) + P(x|fetal)]$, and if $P(x|maternal) < P(x|fetal)$, then the DNA sequence can be classified as fetal, $P(x|fetal) = P(x|fetal) / [P(x|maternal) + P(x|fetal)]$. In an embodiment of the present disclosure, a distributions of

maternal and fetal sequence lengths can be determined that is specific for that sample by considering the sequences that can be assigned as maternal or fetal with high probability, and then that sample specific distribution can be used as the expected size distribution for that sample.

Variable Read Depth to Minimize Sequencing Cost

[0749] In many clinical trials concerning a diagnostic, for example, in Chiu et al. BMJ 2011; 342:c7401, a protocol with a number of parameters is set, and then the same protocol is executed with the same parameters for each of the patients in the trial. In the case of determining the ploidy status of a fetus gestating in a mother using sequencing as a method to measure genetic material one pertinent parameter is the number of reads. The number of reads may refer to the number of actual reads, the number of intended reads, fractional lanes, full lanes, or full flow cells on a sequencer. In these studies, the number of reads is typically set at a level that will ensure that all or nearly all of the samples achieve the desired level of accuracy. Sequencing is currently an expensive technology, a cost of roughly \$200 per 5 mappable million reads, and while the price is dropping, any method which allows a sequencing based diagnostic to operate at a similar level of accuracy but with fewer reads will necessarily save a considerable amount of money.

[0750] The accuracy of a ploidy determination is typically dependent on a number of factors, including the number of reads and the fraction of fetal DNA in the mixture. The accuracy is typically higher when the fraction of fetal DNA in the mixture is higher. At the same time, the accuracy is typically higher if the number of reads is greater. It is possible to have a situation with two cases where the ploidy state is determined with comparable accuracies wherein the first case has a lower fraction of fetal DNA in the mixture than the second, and more reads were sequenced in the first case than the second. It is possible to use the estimated fraction of fetal DNA in the mixture as a guide in determining the number of reads necessary to achieve a given level of accuracy.

[0751] In an embodiment of the present disclosure, a set of samples can be run where different samples in the set are sequenced to different reads depths, wherein the number of reads run on each of the samples is chosen to achieve a given level of accuracy given the calculated fraction of fetal DNA in each mixture. In an embodiment of the present disclosure, this may entail making a measurement of the mixed sample to determine the fraction of fetal DNA in the mixture; this estimation of the fetal fraction may be done with sequencing, it may be done with TAQMAN, it may be done with qPCR, it may be done with SNP arrays, it may be done with any method that can distinguish different alleles at a given loci. The need for a fetal fraction estimate may be eliminated by including hypotheses that cover all or a selected set of fetal fractions in the set of hypotheses that are considered when comparing to the actual measured data. After the fraction fetal DNA in the mixture has been determined, the number of sequences to be read for each sample may be determined.

[0752] In an embodiment of the present disclosure, 100 pregnant women visit their respective OB's, and their blood is drawn into blood tubes with an anti-lysan and/or something to inactivate DNAase. They each take home a kit for the father of their gestating fetus who gives a saliva sample. Both sets of genetic materials for all 100 couples are sent

back to the laboratory, where the mother blood is spun down and the buffy coat is isolated, as well as the plasma. The plasma comprises a mixture of maternal DNA as well as placentally derived DNA. The maternal buffy coat and the paternal blood is genotyped using a SNP array, and the DNA in the maternal plasma samples are targeted with SURESELECT hybridization probes. The DNA that was pulled down with the probes is used to generate 100 tagged libraries, one for each of the maternal samples, where each sample is tagged with a different tag. A fraction from each library is withdrawn, each of those fractions are mixed together and added to two lanes of a ILLUMINA HISEQ DNA sequencer in a multiplexed fashion, wherein each lane resulted in approximately 50 million mappable reads, resulting in approximately 100 million mappable reads on the 100 multiplexed mixtures, or approximately 1 million reads per sample. The sequence reads were used to determine the fraction of fetal DNA in each mixture. 50 of the samples had more than 15% fetal DNA in the mixture, and the 1 million reads were sufficient to determine the ploidy status of the fetuses with a 99.9% confidence.

[0753] Of the remaining mixtures, 25 had between 10 and 15% fetal DNA; a fraction of each of the relevant libraries prepped from these mixtures were multiplexed and run down one lane of the HISEQ generating an additional 2 million reads for each sample. The two sets of sequence data for each of the mixture with between 10 and 15% fetal DNA were added together, and the resulting 3 million reads per sample which were sufficient to determine the ploidy state of those fetuses with 99.9% confidence.

[0754] Of the remaining mixtures, 13 had between 6 and 10% fetal DNA; a fraction of each of the relevant libraries prepped from these mixtures were multiplexed and run down one lane of the HISEQ generating an additional 4 million reads for each sample. The two sets of sequence data for each of the mixture with between 6 and 10% fetal DNA were added together, and the resulting 5 million total reads per mixture which were sufficient to determine the ploidy state of those fetuses with 99.9% confidence.

[0755] Of the remaining mixtures, 8 had between 4 and 6% fetal DNA; a fraction of each of the relevant libraries prepped from these mixtures were multiplexed and run down one lane of the HISEQ generating an additional 6 million reads for each sample. The two sets of sequence data for each of the mixture with between 4 and 6% fetal DNA were added together, and the resulting 7 million total reads per mixture which were sufficient to determine the ploidy state of those fetuses with 99.9% confidence.

[0756] Of the remaining four mixtures, all of them had between 2 and 4% fetal DNA; a fraction of each of the relevant libraries prepped from these mixtures were multiplexed and run down one lane of the HISEQ generating an additional 12 million reads for each sample. The two sets of sequence data for each of the mixture with between 2 and 4% fetal DNA were added together, and the resulting 13 million total reads per mixture which were sufficient to determine the ploidy state of those fetuses with 99.9% confidence.

[0757] This method required six lanes of sequencing on a HISEQ machine to achieve 99.9% accuracy over 100 samples. If the same number of runs had been required for every sample, to ensure that every ploidy determination was made with a 99.9% accuracy, it would have taken 25 lanes

of sequencing, and if a no-call rate or error rate of 4% was tolerated, it could have been achieved with 14 lanes of sequencing.

Using Raw Genotyping Data

[0758] There are a number of methods that can accomplish NPD using fetal genetic information measured on fetal DNA found in maternal blood. Some of these methods involve making measurements of the fetal DNA using SNP arrays, some methods involve untargeted sequencing, and some methods involve targeted sequencing. The targeted sequencing may target SNPs, it may target STRs, it may target other polymorphic loci, it may target non-polymorphic loci, or some combination thereof. Some of these methods may involve using a commercial or proprietary allele caller that calls the identity of the alleles from the intensity data that comes from the sensors in the machine doing the measuring. For example, the ILLUMINA INFINITIUM system or the AFFYMETRIX GENECHIP microarray system involves beads or microchips with attached DNA sequences that can hybridize to complementary segments of DNA; upon hybridization, there is a change in the fluorescent properties of the sensor molecule that can be detected. There are also sequencing methods, for example the ILLUMINA SOLEXA GENOME SEQUENCER or the ABI SOLID GENOME SEQUENCER, wherein the genetic sequence of fragments of DNA are sequenced; upon extension of the strand of DNA complementary to the strand being sequenced, the identity of the extended nucleotide is typically detected via a fluorescent or radio tag appended to the complementary nucleotide. In all of these methods the genotypic or sequencing data is typically determined on the basis of fluorescent or other signals, or the lack thereof. These systems are typically combined with low level software packages that make specific allele calls (secondary genetic data) from the analog output of the fluorescent or other detection device (primary genetic data). For example, in the case of a given allele on a SNP array, the software will make a call, for example, that a certain SNP is present or not present if the fluorescent intensity is measure above or below a certain threshold. Similarly, the output of a sequencer is a chromatogram that indicates the level of fluorescence detected for each of the dyes, and the software will make a call that a certain base pair is A or T or C or G. High throughput sequencers typically make a series of such measurements, called a read, that represents the most likely structure of the DNA sequence that was sequenced. The direct analog output of the chromatogram is defined here to be the primary genetic data, and the base pair/SNP calls made by the software are considered here to be the secondary genetic data. In an embodiment, primary data refers to the raw intensity data that is the unprocessed output of a genotyping platform, where the genotyping platform may refer to a SNP array, or to a sequencing platform. The secondary genetic data refers to the processed genetic data, where an allele call has been made, or the sequence data has been assigned base pairs, and/or the sequence reads have been mapped to the genome.

[0759] Many higher level applications take advantage of these allele calls, SNP calls and sequence reads, that is, the secondary genetic data, that the genotyping software produces. For example, DNA NEXUS, ELAND or MAQ will take the sequencing reads and map them to the genome. For example, in the context of non-invasive prenatal diagnosis, complex informatics, such as PARENTAL SUPPORT™,

may leverage a large number of SNP calls to determine the genotype of an individual. Also, in the context of preimplantation genetic diagnosis, it is possible to take a set of sequence reads that are mapped to the genome, and by taking a normalized count of the reads that are mapped to each chromosome, or section of a chromosome, it may be possible to determine the ploidy state of an individual. In the context of non-invasive prenatal diagnosis it may be possible to take a set of sequence reads that have been measured on DNA present in maternal plasma, and map them to the genome. One may then take a normalized count of the reads that are mapped to each chromosome, or section of a chromosome, and use that data to determine the ploidy state of an individual. For example, it may be possible to conclude that those chromosomes that have a disproportionately large number of reads are trisomic in the fetus that is gestating in the mother from which the blood was drawn.

[0760] However, in reality, the initial output of the measuring instruments is an analog signal. When a certain base pair is called by the software that is associated with the sequencing software, for example the software may call the base pair a T, in reality the call is the call that the software believes to be most likely. In some cases, however, the call may be of low confidence, for example, the analog signal may indicate that the particular base pair is only 90% likely to be a T, and 10% likely to be an A. In another example, the genotype calling software that is associated with a SNP array reader may call a certain allele to be G. However, in reality, the underlying analog signal may indicate that it is only 70% likely that the allele is G, and 30% likely that the allele is T. In these cases, when the higher level applications use the genotype calls and sequence calls made by the lower level software, they are losing some information. That is, the primary genetic data, as measured directly by the genotyping platform, may be messier than the secondary genetic data that is determined by the attached software packages, but it contains more information. In mapping the secondary genetic data sequences to the genome, many reads are thrown out because some bases are not read with enough clarity and or mapping is not clear. When the primary genetic data sequence reads are used, all or many of those reads that may have been thrown out when first converted to secondary genetic data sequence read can be used by treating the reads in a probabilistic manner.

[0761] In an embodiment of the present disclosure, the higher level software does not rely on the allele calls, SNP calls, or sequence reads that are determined by the lower level software. Instead, the higher level software bases its calculations on the analog signals directly measured from the genotyping platform. In an embodiment of the present disclosure, an informatics based method such as PARENTAL SUPPORT™ is modified so that its ability to reconstruct the genetic data of the embryo/fetus/child is engineered to directly use the primary genetic data as measured by the genotyping platform. In an embodiment of the present disclosure, an informatics based method such as PARENTAL SUPPORT™ is able to make allele calls, and/or chromosome copy number calls using primary genetic data, and not using the secondary genetic data. In an embodiment of the present disclosure, all genetic calls, SNPs calls, sequence reads, sequence mapping is treated in a probabilistic manner by using the raw intensity data as measured directly by the genotyping platform, rather than converting the primary genetic data to secondary genetic calls. In an

embodiment, the DNA measurements from the prepared sample used in calculating allele count probabilities and determining the relative probability of each hypothesis comprise primary genetic data.

[0762] In some embodiments, the method can increase the accuracy of genetic data of a target individual which incorporates genetic data of at least one related individual, the method comprising obtaining primary genetic data specific to a target individual's genome and genetic data specific to the genome(s) of the related individual(s), creating a set of one or more hypotheses concerning possibly which segments of which chromosomes from the related individual(s) correspond to those segments in the target individual's genome, determining the probability of each of the hypotheses given the target individual's primary genetic data and the related individual(s)'s genetic data, and using the probabilities associated with each hypothesis to determine the most likely state of the actual genetic material of the target individual. In some embodiments, the method can determining the number of copies of a segment of a chromosome in the genome of a target individual, the method comprising creating a set of copy number hypotheses about how many copies of the chromosome segment are present in the genome of a target individual, incorporating primary genetic data from the target individual and genetic information from one or more related individuals into a data set, estimating the characteristics of the platform response associated with the data set, where the platform response may vary from one experiment to another, computing the conditional probabilities of each copy number hypothesis, given the data set and the platform response characteristics, and determining the copy number of the chromosome segment based on the most probable copy number hypothesis. In an embodiment, a method of the present disclosure can determine a ploidy state of at least one chromosome in a target individual, the method comprising obtaining primary genetic data from the target individual and from one or more related individuals, creating a set of at least one ploidy state hypothesis for each of the chromosomes of the target individual, using one or more expert techniques to determine a statistical probability for each ploidy state hypothesis in the set, for each expert technique used, given the obtained genetic data, combining, for each ploidy state hypothesis, the statistical probabilities as determined by the one or more expert techniques, and determining the ploidy state for each of the chromosomes in the target individual based on the combined statistical probabilities of each of the ploidy state hypotheses. In an embodiment, a method of the present disclosure can determine an allelic state in a set of alleles, in a target individual, and from one or both parents of the target individual, and optionally from one or more related individuals, the method comprising obtaining primary genetic data from the target individual, and from the one or both parents, and from any related individuals, creating a set of at least one allelic hypothesis for the target individual, and for the one or both parents, and optionally for the one or more related individuals, where the hypotheses describe possible allelic states in the set of alleles, determining a statistical probability for each allelic hypothesis in the set of hypotheses given the obtained genetic data, and determining the allelic state for each of the alleles in the set of alleles for the target individual, and for the one or both parents, and optionally for the one or more related individuals, based on the statistical probabilities of each of the allelic hypotheses.

[0763] In some embodiments, the genetic data of the mixed sample may comprise sequence data wherein the sequence data may not uniquely map to the human genome. In some embodiments, the genetic data of the mixed sample may comprise sequence data wherein the sequence data maps to a plurality of locations in the genome, wherein each possible mapping is associated with a probability that the given mapping is correct. In some embodiments, the sequence reads are not assumed to be associated with a particular position in the genome. In some embodiments, the sequence reads are associated with a plurality of positions in the genome, and an associated probability belonging to that position.

Counting Method to Determine Chromosome Copy Number

[0764] In one aspect, the invention features methods of testing for an abnormal distribution of a fetal chromosome by comparing the number of sequence tags that align to different chromosomes (see, e.g., U.S. Pat. No. 8,296,076, filed Apr. 20, 2012, which is hereby incorporated by reference in its entirety). As is known in the art, the term "sequence tag" refers to a relatively short (e.g., 15-100) nucleic acid sequence that can be used to identify a certain larger sequence, e.g., be mapped to a chromosome or genomic region or gene. In some embodiments, the method involves (i) contacting a sample that includes a mixture of maternal and fetal DNA with a library of primers that simultaneously hybridize to at least 25; 50; 75; 100; 300; 500; 750; 1,000; 2,000; 5,000; 7,500; 10,000; 15,000; 19,000; 20,000; 25,000; 27,000; 28,000; 30,000; 40,000; 50,000; 75,000; or 100,000 different target loci to produce a reaction mixture; wherein the target loci are from a plurality of different chromosomes; and wherein the plurality of different chromosomes comprise at least one first chromosome suspected of having an abnormal distribution in the sample and at least one second chromosome presumed to be normally distributed in the sample; (ii) subjecting the reaction mixture to primer extension reaction conditions to produce amplified products; (iii) sequencing the amplified products to obtain a plurality of sequence tags aligning to the target loci; wherein the sequence tags are of sufficient length to be assigned to a specific target locus; (iv) assigning on a computer the plurality of sequence tags to their corresponding target loci; (v) determining on a computer a number of sequence tags aligning to the target loci of the first chromosome and a number of sequence tags aligning to the target loci of the second chromosome; and (vi) comparing the numbers from step (v) to determine the presence or absence of an abnormal distribution of the first chromosome.

[0765] In one aspect, the invention provides methods for detecting the presence or absence of a fetal aneuploidy by comparing the relative frequency of target amplicons between chromosomes (see, e.g., PCT Publ. No. WO 2012/103031, filed Jan. 23, 2012, which is hereby incorporated by reference in its entirety). In some embodiments, the method involves (i) contacting a sample that includes a mixture of maternal and fetal DNA with a library of primers that simultaneously hybridize to at least 25; 50; 75; 100; 300; 500; 750; 1,000; 2,000; 5,000; 7,500; 10,000; 15,000; 19,000; 20,000; 25,000; 27,000; 28,000; 30,000; 40,000; 50,000; 75,000; or 100,000 different non-polymorphic target loci to produce a reaction mixture; wherein the target loci are from a plurality of different chromosomes; (ii) subjecting the reaction mixture to primer extension reaction conditions to

produce amplified products that includes target amplicons; (iii) quantifying on a computer a relative frequency of the target amplicons from the first and second chromosomes of interest; (iv) comparing on a computer the relative frequency of the target amplicons from the first and second chromosomes of interest; and (v) identifying the presence or absence of an aneuploidy based on the compared relative frequencies of the first and second chromosome of interest. In some embodiments, the first chromosome is a chromosome suspected of being euploid. In some embodiments, the second chromosome is a chromosome suspected of being aneuploidy.

Combining Methods of Prenatal Diagnosis

[0766] There are many methods that may be used for prenatal diagnosis or prenatal screening of aneuploidy or other genetic defects. Described elsewhere in this document, and in U.S. Utility application Ser. No. 11/603,406, filed Nov. 28, 2006; U.S. Utility application Ser. No. 12/076,348, filed Mar. 17, 2008, and PCT Application Serial No. PCT/S09/52730 is one such method that uses the genetic data of related individuals to increase the accuracy with which genetic data of a target individual, such as a fetus, is known, or estimated. Other methods used for prenatal diagnosis involve measuring the levels of certain hormones in maternal blood, where those hormones are correlated with various genetic abnormalities. An example of this is called the triple test, a test wherein the levels of several (commonly two, three, four or five) different hormones are measured in maternal blood. In a case where multiple methods are used to determine the likelihood of a given outcome, where none of the methods are definitive in and of themselves, it is possible to combine the information given by those methods to make a prediction that is more accurate than any of the individual methods. In the triple test, combining the information given by the three different hormones can result in a prediction of genetic abnormalities that is more accurate than the individual hormone levels may predict.

[0767] Disclosed herein is a method for making more accurate predictions about the genetic state of a fetus, specifically the possibility of genetic abnormalities in a fetus that comprises combining predictions of genetic abnormalities in a fetus where those predictions were made using a variety of methods. A “more accurate” method may refer to a method for diagnosing an abnormality that has a lower false negative rate at a given false positive rate. In a favored embodiment of the present disclosure, one or more of the predictions are made based on the genetic data known about the fetus, where the genetic knowledge was determined using the PARENTAL SUPPORT™ method, that is, using genetic data of individual related to the fetus to determine the genetic data of the fetus with greater accuracy. In some embodiments the genetic data may include ploidy states of the fetus. In some embodiments, the genetic data may refer to a set of allele calls on the genome of the fetus. In some embodiments some of the predictions may have been made using the triple test. In some embodiments, some of the predictions may have been made using measurements of other hormone levels in maternal blood. In some embodiments, predictions made by methods considered diagnoses may be combined with predictions made by methods considered screening. In some embodiments, the method involves measuring maternal blood levels of alpha-fetoprotein (AFP). In some embodiments, the method involves

measuring maternal blood levels of unconjugated estriol (UE_3). In some embodiments, the method involves measuring maternal blood levels of beta human chorionic gonadotropin (beta-hCG). In some embodiments, the method involves measuring maternal blood levels of invasive trophoblast antigen (ITGA). In some embodiments, the method involves measuring maternal blood levels of inhibin. In some embodiments, the method involves measuring maternal blood levels of pregnancy-associated plasma protein A (PAPP-A). In some embodiments, the method involves measuring maternal blood levels of other hormones or maternal serum markers. In some embodiments, some of the predictions may have been made using other methods. In some embodiments, some of the predictions may have been made using a fully integrated test such as one that combines ultrasound and blood test at around 12 weeks of pregnancy and a second blood test at around 16 weeks. In some embodiments, the method involves measuring the fetal nuchal translucency (NT). In some embodiments, the method involves using the measured levels of the aforementioned hormones for making predictions. In some embodiments the method involves a combination of the aforementioned methods.

[0768] There are many ways to combine the predictions, for example, one could convert the hormone measurements into a multiple of the median (MoM) and then into likelihood ratios (LR). Similarly, other measurements could be transformed into LRs using the mixture model of NT distributions. The LRs for NT and the biochemical markers could be multiplied by the age and gestation-related risk to derive the risk for various conditions, such as trisomy 21. Detection rates (DRs) and false-positive rates (FPRs) could be calculated by taking the proportions with risks above a given risk threshold.

[0769] In an embodiment, a method to call the ploidy state involves combining the relative probabilities of each of the ploidy hypotheses determined using the joint distribution model and the allele count probabilities with relative probabilities of each of the ploidy hypotheses that are calculated using statistical techniques taken from other methods that determine a risk score for a fetus being trisomic, including but not limited to: a read count analysis, comparing heterozygosity rates, a statistic that is only available when parental genetic information is used, the probability of normalized genotype signals for certain parent contexts, a statistic that is calculated using an estimated fetal fraction of the first sample or the prepared sample, and combinations thereof.

[0770] Another method could involve a situation with four measured hormone levels, where the probability distribution around those hormones is known: $p(x_1, x_2, x_3, x_4|e)$ for the euploid case and $p(x_1, x_2, x_3, x_4|a)$ for the aneuploid case. Then one could measure the probability distribution for the DNA measurements, $g(y|e)$ and $g(y|a)$ for the euploid and aneuploid cases respectively. Assuming they are independent given the assumption of euploid/aneuploid, one could combine as $p(x_1, x_2, x_3, x_4|a)g(y|a)$ and $p(x_1, x_2, x_3, x_4|e)g(y|e)$ and then multiply each by the prior $p(a)$ and $p(e)$ given the maternal age. One could then choose the one that is highest.

[0771] In an embodiment, it is possible to evoke central limit theorem to assume distribution on $g(y|a)$ or e is Gaussian, and measure mean and standard deviation by looking at multiple samples. In another embodiment, one

could assume they are not independent given the outcome and collect enough samples to estimate the joint distribution $p(x_1, x_2, x_3, x_4|a \text{ or } e)$.

[0772] In an embodiment, the ploidy state for the target individual is determined to be the ploidy state that is associated with the hypothesis whose probability is the greatest. In some cases, one hypothesis will have a normalized, combined probability greater than 90%. Each hypothesis is associated with one, or a set of, ploidy states, and the ploidy state associated with the hypothesis whose normalized, combined probability is greater than 90%, or some other threshold value, such as 50%, 80%, 95%, 98%, 99%, or 99.9%, may be chosen as the threshold required for a hypothesis to be called as the determined ploidy state.

DNA from Children from Previous Pregnancies in Maternal Blood

[0773] One difficulty to non-invasive prenatal diagnosis is differentiating fetal cells from the current pregnancy from fetal cells from previous pregnancies. Some believe that genetic matter from prior pregnancies will go away after some time, but conclusive evidence has not been shown. In an embodiment of the present disclosure, it is possible to determine fetal DNA present in the maternal blood of paternal origin (that is, DNA that the fetus inherited from the father) using the PARENTAL SUPPORT™ (PS) method, and the knowledge of the paternal genome. This method may utilize phased parental genetic information. It is possible to phase the parental genotype from unphased genotypic information using grandparental genetic data (such as measured genetic data from a sperm from the grandfather), or genetic data from other born children, or a sample of a miscarriage. One could also phase unphased genetic information by way of a HapMap-based phasing, or a haplotyping of paternal cells. Successful haplotyping has been demonstrated by arresting cells at phase of mitosis when chromosomes are tight bundles and using microfluidics to put separate chromosomes in separate wells. In another embodiment it is possible to use the phased parental haplotypic data to detect the presence of more than one homolog from the father, implying that the genetic material from more than one child is present in the blood. By focusing on chromosomes that are expected to be euploid in a fetus, one could rule out the possibility that the fetus was afflicted with a trisomy. Also, it is possible to determine if the fetal DNA is not from the current father, in which case one could use other methods such as the triple test to predict genetic abnormalities.

[0774] There may be other sources of fetal genetic material available via methods other than a blood draw. In the case of the fetal genetic material available in maternal blood, there are two main categories: (1) whole fetal cells, for example, nucleated fetal red blood cells or erythroblasts, and (2) free floating fetal DNA. In the case of whole fetal cells, there is some evidence that fetal cells can persist in maternal blood for an extended period of time such that it is possible to isolate a cell from a pregnant woman that contains the DNA from a child or fetus from a prior pregnancy. There is also evidence that the free floating fetal DNA is cleared from the system in a matter of weeks. One challenge is how to determine the identity of the individual whose genetic material is contained in the cell, namely to ensure that the measured genetic material is not from a fetus from a prior pregnancy. In an embodiment of the present disclosure, the knowledge of the maternal genetic material can be used to

ensure that the genetic material in question is not maternal genetic material. There are a number of methods to accomplish this end, including informatics based methods such as PARENTAL SUPPORT™, as described in this document or any of the patents referenced in this document.

[0775] In an embodiment of the present disclosure, the blood drawn from the pregnant mother may be separated into a fraction comprising free floating fetal DNA, and a fraction comprising nucleated red blood cells. The free floating DNA may optionally be enriched, and the genotypic information of the DNA may be measured. From the measured genotypic information from the free floating DNA, the knowledge of the maternal genotype may be used to determine aspects of the fetal genotype. These aspects may refer to ploidy state, and/or a set of allele identities. Then, individual nucleated red blood cells may be genotyped using methods described elsewhere in this document, and other referent patents, especially those mentioned in the first section of this document. The knowledge of the maternal genome would allow one to determine whether or not any given single blood cell is genetically maternal. And the aspects of the fetal genotype that were determined as described above would allow one to determine if the single blood cell is genetically derived from the fetus that is currently gestating. In essence, this aspect of the present disclosure allows one to use the genetic knowledge of the mother, and possibly the genetic information from other related individuals, such as the father, along with the measured genetic information from the free floating DNA found in maternal blood to determine whether an isolated nucleated cell found in maternal blood is either (a) genetically maternal, (b) genetically from the fetus currently gestating, or (c) genetically from a fetus from a prior pregnancy.

Prenatal Sex Chromosome Aneuploidy Determination

[0776] In methods known in the art, people attempting to determine the sex of a gestating fetus from the blood of the mother have used the fact that fetal free floating DNA (ffDNA) is present in the plasma of the mother. If one is able to detect Y-specific loci in the maternal plasma, this implies that the gestating fetus is a male. However, the lack of detection of Y-specific loci in the plasma does not always guarantee that the gestating fetus is a female when using methods known in the art, as in some cases the amount of ffDNA is too low to ensure that the Y-specific loci would be detected in the case of a male fetus.

[0777] Presented here is a novel method that does not require the measurement of Y-specific nucleic acids, that is, DNA that is from loci that are exclusively paternally derived. The Parental Support method, disclosed previously, uses crossover frequency data, parental genotypic data, and informatics techniques, to determine the ploidy state of a gestating fetus. The sex of a fetus is simply the ploidy state of the fetus at the sex chromosomes. A child that is XX is female, and XY is male. The method described herein is also able to determine the ploidy state of the fetus. Note that sexing is effectively synonymous with ploidy determination of the sex chromosomes; in the case of sexing, an assumption is often made that the child is euploid, therefore there are fewer possible hypotheses.

[0778] The method disclosed herein involves looking at loci that are common to both the X and Y chromosome to create a baseline in terms of expected amount of fetal DNA present for a fetus. Then, those regions that are specific only

to the X chromosome can be interrogated to determine if the fetus is female or male. In the case of a male, we expect to see less fetal DNA from loci that are specific to the X chromosome than from loci that are specific to both the X and the Y. In contrast, in female fetuses, we expect the amount of DNA for each of these groups to be the same. The DNA in question can be measured by any technique that can quantitate the amount of DNA present on a sample, for example, qPCR, SNP arrays, genotyping arrays, or sequencing. For DNA that is exclusively from an individual we would expect to see the following:

	DNA specific to X	DNA specific to X and Y	DNA specific to Y
Male (XY)	A	2A	A
Female (XX)	2A	2A	0

[0779] In the case of DNA from a fetus that is mixed with DNA from the mother, and where the fraction of fetal DNA in the mixture is F, and where the fraction of maternal DNA in the mixture is M, such that F+M=100%, we would expect to see the following:

	DNA specific to X	DNA specific to X and Y	DNA specific to Y
Male fetus (XY)	M + $\frac{1}{2}$ F	M + F	$\frac{1}{2}$ F
Female fetus (XX)	M + F	M + F	0

[0780] In the case where F and M are known, the expected ratios can be computed, and the observed data can be compared to the expected data. In the case where M and F are not known, a threshold can be selected based on historical data. In both cases, the measured amount of DNA at loci specific to both X and Y can be used as a baseline, and the test for the sex of the fetus can be based on the amount of DNA observed on loci specific to only the X chromosome. If that amount is lower than the baseline by an amount roughly equal to $\frac{1}{2}$ F, or by an amount that causes it to fall below a predefined threshold, the fetus is determined to be male, and if that amount is about equal to the baseline, or if is not lower by an amount that causes it to fall below a predefined threshold, the fetus is determined to be female.

[0781] In another embodiment, one can look only at those loci that are common to both the X and the Y chromosomes, often termed the Z chromosome. A subset of the loci on the Z chromosome are typically always A on the X chromosome, and B on the Y chromosome. If SNPs from the Z chromosome are found to have the B genotype, then the fetus is called a male; if the SNPs from the Z chromosome are found to only have A genotype, then the fetus is called a female. In another embodiment, one can look at the loci that are found only on the X chromosome. Contexts such as AA|B are particularly informative as the presence of a B indicates that the fetus has an X chromosome from the father. Contexts such as AB|B are also informative, as we expect to see B present only half as often in the case of a female fetus as compared to a male fetus. In another embodiment, one can look at the SNPs on the Z chromo-

some where both A and B alleles are present on both the X and the Y chromosome, and where it is known which SNPs are from the paternal Y chromosome, and which are from the paternal X chromosome.

[0782] In an embodiment, it is possible to amplify single nucleotide positions known to varying between the homologous non-recombining (HNR) region shared by chromosome Y and chromosome X. The sequence within this HNR region is largely identical between the X and Y chromosomes. Within this identical region are single nucleotide positions that, while invariant among X chromosomes and among Y chromosomes in the population, are different between the X and Y chromosomes. Each PCR assay could amplify a sequence from loci that are present on both the X and Y chromosomes. Within each amplified sequence would be a single base that can be detected using sequencing or some other method (see, for example, U.S. Publication No. 2011/0178719, filed Feb. 3, 2011, which is hereby incorporated by reference in its entirety).

[0783] In an embodiment, the sex of the fetus could be determined from the fetal free floating DNA found in maternal plasma, the method comprising some or all of the following steps: 1) Design PCR (either regular or mini-PCR, plus multiplexing if desired) primers amplify X/Y variant single nucleotide positions within HNR region, 2) obtain maternal plasma, 3) PCR Amplify targets from maternal plasma using HNR X/Y PCR assays, 4) sequence the amplicons, 5) Examine sequence data for presence of Y-allele within one or more of the amplified sequences. The presence of one or more would indicate a male fetus. Absence of all Y-alleles from all amplicons indicates a female fetus.

[0784] In an embodiment, one could use targeted sequencing to measure the DNA in the maternal plasma and/or the parental genotypes. In an embodiment, one could ignore all sequences that clearly originate from paternally sourced DNA. For example, in the context AA|AB, one could count the number of A sequences and ignore all the B sequences. In order to determine a heterozygosity rate for the above algorithm, one could compare the number of observed A sequences to the expected number of total sequences for the given probe. There are many ways one could calculate an expected number of sequences for each probe on a per sample basis. In an embodiment, it is possible to use historical data to determine what fraction of all sequence reads belongs to each specific probe and then use this empirical fraction, combined with the total number of sequence reads, to estimate the number of sequences at each probe. Another approach could be to target some known homozygous alleles and then use historical data to relate the number of reads at each probe with the number of reads at the known homozygous alleles. For each sample, one could then measure the number of reads at the homozygous alleles and then use this measurement, along with the empirically derived relationships, to estimate the number of sequence reads at each probe.

[0785] In some embodiments, it is possible to determine the sex of the fetus by combining the predictions made by a plurality of methods. In some embodiments the plurality of methods are taken from methods described in this disclosure. In some embodiments, at least one of the plurality of methods are taken from methods described in this disclosure.

[0786] In some embodiments the method described herein can be used to determine the ploidy state of the gestating

fetus. In an embodiment, the ploidy calling method uses loci that are specific to the X chromosome, or common to both the X and Y chromosome, but does not make use of any Y-specific loci. In an embodiment, the ploidy calling method uses one or more of the following: loci that are specific to the X chromosome, loci that are common to both the X and Y chromosome, and loci that are specific to the Y chromosome. In an embodiment, where the ratios of sex chromosomes are similar, for example 45,X (Turner Syndrome), 46,XX (normal female) and 47,XXX (trisomy X), the differentiation can be accomplished by comparing the allele distributions to expected allele distributions according to the various hypotheses. In another embodiment, this can be accomplished by comparing the relative number of sequence reads for the sex chromosomes to one or a plurality of reference chromosomes that are assumed to be euploid. Also note that these methods can be expanded to include aneuploid cases.

Single Gene Disease Screening

[0787] In an embodiment, a method for determining the ploidy state of the fetus may be extended to enable simultaneous testing for single gene disorders. Single-gene disease diagnosis leverages the same targeted approach used for aneuploidy testing, and requires additional specific targets. In an embodiment, the single gene NPD diagnosis is through linkage analysis. In many cases, direct testing of the cfDNA sample is not reliable, as the presence of maternal DNA makes it virtually impossible to determine if the fetus has inherited the mother's mutation. Detection of a unique paternally-derived allele is less challenging, but is only fully informative if the disease is dominant and carried by the father, limiting the utility of the approach. In an embodiment, the method involves PCR or related amplification approaches.

[0788] In some embodiments, the method involves phasing the abnormal allele with surrounding very tightly linked SNPs in the parents using information from first-degree relatives. Then Parental Support may be run on the targeted sequencing data obtained from these SNPs to determine which homologs, normal or abnormal, were inherited by the fetus from both parents. As long as the SNPs are sufficiently linked, the inheritance of the genotype of the fetus can be determined very reliably. In some embodiments, the method comprises (a) adding a set of SNP loci to densely flank a specified set of common diseases to our multiplex pool for aneuploidy testing; (b) reliably phasing the alleles from these added SNPs with the normal and abnormal alleles based on genetic data from various relatives; and (c) reconstructing the fetal haplotype, or set of phased SNP alleles on the inherited maternal and paternal homologs in the region surrounding the disease locus to determine fetal genotype. In some embodiments additional probes that are closely linked to a disease linked locus are added to the set of polymorphic locus being used for aneuploidy testing.

[0789] Reconstructing fetal diploidy is challenging because the sample is a mixture of maternal and fetal DNA. In some embodiments, the method incorporates relative information to phase the SNPs and disease alleles, then take into account physical distance of the SNPs and recombination data from location specific recombination likelihoods and the data observed from the genetic measurements of the maternal plasma to obtain the most likely genotype of the fetus.

[0790] In an embodiment, a number of additional probes per disease linked locus are included in the set of targeted polymorphic loci; the number of additional probes per disease linked locus may be between 4 and 10, between 11 and 20, between 21 and 40, between 41 and 60, between 61 and 80, or combinations thereof.

[0791] Phasing the diploid data from the parents can be challenging, and there are a number of ways this can be accomplished. Some are discussed in this disclosure, others are described in greater detail in other disclosures (see, e.g., PCT Publ. No. WO2009105531, filed Feb. 9, 2009, and PCT Publ. No. WO2010017214, filed Aug. 4, 2009, which are each hereby incorporated by reference in its entirety). In one embodiment, a parent can be phased by inference by measuring tissue from the parent that is haploid, for example by measuring one or more sperm or eggs. In one embodiment the parent can be phased by inference using the measured genotypic data of a first degree relative such as the parent's parent(s) or siblings. In one embodiment, the parent can be phased by dilution where the DNA is diluted, in one or a plurality of wells, to the point where there is expected to be no more than approximately one copy of each haplotype in each well, and then measuring the DNA in the one or more wells. In one embodiment, the parent genotype can be phased by using computer programs that use population based haplotype frequencies to infer the most likely phase. In one embodiment, the parent can be phased if the phased haplotypic data is known for the other parent, along with the unphased genetic data of one or more genetic offspring of the parents. In some embodiments, the genetic offspring of the parents may be one or more embryos, fetuses, and/or born children. Some of these methods and other methods for phasing one or both parents are disclosed in greater detail in, e.g., U.S. Publ. No. 2011/0033862, filed Aug. 19, 2010; U.S. Publ. No. 2011/0178719, filed Feb. 3, 2011; U.S. Publ. No. 2007/0184467, filed Nov. 22, 2006; U.S. Publ. No. 2008/0243398, filed Mar. 17, 2008, which are each hereby incorporated by reference in its entirety.

Fetal Genome Reconstruction

[0792] In one aspect, the invention features methods for determining a haplotype of a fetus. In various embodiments, this method allows one to determine which polymorphic loci (such as SNPs) were inherited by the fetus and to reconstruct which homologs (including recombination events) are present in the fetus (and thereby interpolate the sequence between the polymorphic loci). If desired, essentially the entire genome of the fetus can be reconstructed. If there is some remaining ambiguity in the genome of the fetus (such as in intervals with a crossover), this ambiguity can be minimized if desired by analyzing additional polymorphic loci. In various embodiments, the polymorphic loci are chosen to cover one or more of the chromosomes at a density to reduce any ambiguity to a desired level. This method has important applications for the detection of polymorphisms or other mutations of interest in a fetus since it enables their detection based on linkage (such as the presence of linked polymorphic loci in the fetal genome) rather than by directly detecting the polymorphism or other mutation of interest in the fetal genome. For example, if a parent is a carrier for a mutation associated with cystic fibrosis (CF), a nucleic acid sample that includes maternal DNA from the mother of the fetus and fetal DNA from the fetus can be analyzed to determine whether the fetal DNA include the haplotype

containing the CF mutation. In particular, polymorphic loci can be analyzed to determine whether the fetal DNA includes the haplotype containing the CF mutation without having to detect the CF mutation itself in the fetal DNA. This is useful in screening for one or more mutations, such as disease-linked mutations, without having to directly detect the mutations.

[0793] In some embodiments, the method involves determining a parental haplotype (e.g., a haplotype of the mother or father of the fetus). In some embodiments, this determination is made without using data from a relative of the mother or father. In some embodiments, a parental haplotype is determined using a dilution approach followed by SNP genotyping or sequencing as described herein and elsewhere (see, e.g., U.S. Publ. No. 2011/0033862, filed Aug. 19, 2010, which is hereby incorporated by reference in its entirety). Because the DNA is diluted, it is unlikely that more than one haplotype is in the same fraction (or tube). Thus, there may be effectively a single molecule of DNA in the tube, which allows the haplotype on a single DNA molecule to be determined. In some embodiments, the method includes dividing a DNA sample into a plurality of fractions such that at least one of the fractions includes one chromosome or one chromosome segment from a pair of chromosomes, and genotyping (e.g., determining the presence of two or more polymorphic loci) the DNA sample in at least one of the fractions, thereby determining a parental haplotype. In some embodiments, the genotyping involves sequencing (such as shotgun sequencing). In some embodiments, the genotyping involves use of a SNP array to detect polymorphic loci, such as at least 25; 50; 75; 100; 300; 500; 750; 1,000; 2,000; 5,000; 7,500; 10,000; 15,000; 19,000; 20,000; 25,000; 27,000; 28,000; 30,000; 40,000; 50,000; 75,000; or 100,000 different polymorphic loci. In some embodiments, the genotyping involves the use of multiplex PCR. In some embodiments, the method involves contacting the sample in a fraction with a library of primers that simultaneously hybridize to at least 25; 50; 75; 100; 300; 500; 750; 1,000; 2,000; 5,000; 7,500; 10,000; 15,000; 19,000; 20,000; 25,000; 27,000; 28,000; 30,000; 40,000; 50,000; 75,000; or 100,000 different polymorphic loci (such as SNPs) to produce a reaction mixture; and subjecting the reaction mixture to primer extension reaction conditions to produce amplified products that are measured with a high throughput sequencer to produce sequencing data.

[0794] In some embodiments, a haplotype of the mother is determined by any of the methods described herein using data from a relative of the mother. In some embodiments, a haplotype of the father is determined by any of the methods described herein using data from a relative of the father. In some embodiments, a haplotype is determined for both the father and the mother. In some embodiments, a SNP array is used to determine the presence of at least 25; 50; 75; 100; 300; 500; 750; 1,000; 2,000; 5,000; 7,500; 10,000; 15,000; 19,000; 20,000; 25,000; 27,000; 28,000; 30,000; 40,000; 50,000; 75,000; or 100,000 different polymorphic loci in a DNA sample from the mother (or father) and a relative of the mother (or father). In some embodiments, the method involves contacting a DNA sample from the mother (or father) and/or a relative of the mother (or father) with a library of primers that simultaneously hybridize to at least 25; 50; 75; 100; 300; 500; 750; 1,000; 2,000; 5,000; 7,500; 10,000; 15,000; 19,000; 20,000; 25,000; 27,000; 28,000; 30,000; 40,000; 50,000; 75,000; or 100,000 different poly-

morphic loci (such as SNPs) to produce a reaction mixture; and subjecting the reaction mixture to primer extension reaction conditions to produce amplified products that are measured with a high throughput sequencer to produce sequencing data. The parental haplotype may be determined based on the SNP array or sequencing data. In some embodiments, parental data may be phased by methods described or referred to elsewhere in this document.

[0795] This parental haplotype data can be used to determine if the fetus inherited the parental haplotype. In some embodiments, a nucleic acid sample that includes maternal DNA from the mother of the fetus and fetal DNA from the fetus is analyzed using a SNP array to detect at least 25; 50; 75; 100; 300; 500; 750; 1,000; 2,000; 5,000; 7,500; 10,000; 15,000; 19,000; 20,000; 25,000; 27,000; 28,000; 30,000; 40,000; 50,000; 75,000; or 100,000 different polymorphic loci.

[0796] In some embodiments, a nucleic acid sample that includes maternal DNA from the mother of the fetus and fetal DNA from the fetus is analyzed by contacting the sample with a library of primers that simultaneously hybridize to at least 25; 50; 75; 100; 300; 500; 750; 1,000; 2,000; 5,000; 7,500; 10,000; 15,000; 19,000; 20,000; 25,000; 27,000; 28,000; 30,000; 40,000; 50,000; 75,000; or 100,000 different polymorphic loci (such as SNPs) to produce a reaction mixture. In some embodiments, the reaction mixture is subjected to primer extension reaction conditions to produce amplified products. In some embodiments, the amplified products are measured with a high throughput sequencer to produce sequencing data. In various embodiments, the SNP array or sequencing data is used to determine a parental haplotype by using data about the probability of chromosomes crossing over at different locations in a chromosome (such as by using recombination data such as may be found in the HapMap database to create a recombination risk score for any interval) to model dependence between polymorphic alleles on the chromosome. In some embodiments, allele counts at the polymorphic loci are calculated on a computer based on the sequencing data. In some embodiments, a plurality of ploidy hypotheses each pertaining to a different possible ploidy state of the chromosome are created on a computer; a model (such as a joint distribution model) for the expected allele counts at the polymorphic loci on the chromosome is built on a computer for each ploidy hypothesis; a relative probability of each of the ploidy hypotheses is determined on a computer using the joint distribution model and the allele counts; and the ploidy state of the fetus is called by selecting the ploidy state corresponding to the hypothesis with the greatest probability. In some embodiments, building a joint distribution model for allele counts and the step of determining the relative probability of each hypothesis are done using a method that does not require the use of a reference chromosome.

[0797] In some embodiments, a fetal haplotype is determined for one or more chromosomes taken from the group consisting of chromosomes 13, 18, 21, X, and Y. In some embodiments, a fetal haplotype is determined for all of the fetal chromosomes. In various embodiments, the method determines essentially the entire genome of the fetus. In some embodiments, the haplotype is determined for at least 30, 40, 50, 60, 70, 80, 90, or 95% of the genome of the fetus.

[0798] In some embodiments, the haplotype determination of the fetus includes information about which allele is present for at least 25; 50; 75; 100; 300; 500; 750; 1,000;

2,000; 5,000; 7,500; 10,000; 15,000; 19,000; 20,000; 25,000; 27,000; 28,000; 30,000; 40,000; 50,000; 75,000; or 100,000 different polymorphic loci.

Compositions of DNA

[0799] When performing an informatics analysis on sequencing data measured on a mixture of fetal and maternal blood to determine genomic information pertaining to the fetus, for example the ploidy state of the fetus, it may be advantageous to measure the allele distributions at a set of alleles. Unfortunately, in many cases, such as when attempting to determine the ploidy state of a fetus from the DNA mixture found in the plasma of a maternal blood sample, the amount of DNA available is not sufficient to directly measure the allele distributions with good fidelity in the mixture. In these cases, amplification of the DNA mixture will provide sufficient numbers of DNA molecules that the desired allele distributions may be measured with good fidelity. However, current methods of amplification typically used in the amplification of DNA for sequencing are often very biased, meaning that they do not amplify both alleles at a polymorphic locus by the same amount. A biased amplification can result in allele distributions that are quite different from the allele distributions in the original mixture. For most purposes, highly accurate measurements of the relative amounts of alleles present at polymorphic loci are not needed. In contrast, in an embodiment of the present disclosure, amplification or enrichment methods that specifically enrich polymorphic alleles and preserve allelic ratios is advantageous.

[0800] A number of methods are described herein that may be used to preferentially enrich a sample of DNA at a plurality of loci in a way that minimizes allelic bias. Some examples are using circularizing probes to target a plurality of loci where the 3' ends and 5' ends of the pre-circularized probe are designed to hybridize to bases that are one or a few positions away from the polymorphic sites of the targeted allele. Another is to use PCR probes where the 3' end PCR probe is designed to hybridize to bases that are one or a few positions away from the polymorphic sites of the targeted allele. Another is to use a split and pool approach to create mixtures of DNA where the preferentially enriched loci are enriched with low allelic bias without the drawbacks of direct multiplexing. Another is to use a hybrid capture approach where the capture probes are designed such that the region of the capture probe that is designed to hybridize to the DNA flanking the polymorphic site of the target is separated from the polymorphic site by one or a small number of bases.

[0801] In the case where measured allele distributions at a set of polymorphic loci are used to determine the ploidy state of an individual, it is desirable to preserve the relative amounts of alleles in a sample of DNA as it is prepared for genetic measurements. This preparation may involve WGA amplification, targeted amplification, selective enrichment techniques, hybrid capture techniques, circularizing probes or other methods meant to amplify the amount of DNA and/or selectively enhance the presence of molecules of DNA that correspond to certain alleles.

[0802] In some embodiments of the present disclosure, there is a set of DNA probes designed to target loci where the loci have maximal minor allele frequencies. In some embodiments of the present disclosure, there is a set of probes that are designed to target where the loci have the

maximum likelihood of the fetus having a highly informative SNP at those loci. In some embodiments of the present disclosure, there is a set of probes that are designed to target loci where the probes are optimized for a given population subgroup. In some embodiments of the present disclosure, there is a set of probes that are designed to target loci where the probes are optimized for a given mix of population subgroups. In some embodiments of the present disclosure, there is a set of probes that are designed to target loci where the probes are optimized for a given pair of parents which are from different population subgroups that have different minor allele frequency profiles. In some embodiments of the present disclosure, there is a circularized strand of DNA that comprises at least one base pair that annealed to a piece of DNA that is of fetal origin. In some embodiments of the present disclosure, there is a circularized strand of DNA that comprises at least one base pair that annealed to a piece of DNA that is of placental origin. In some embodiments of the present disclosure, there is a circularized strand of DNA that circularized while at least some of the nucleotides were annealed to DNA that was of fetal origin. In some embodiments of the present disclosure, there is a circularized strand of DNA that circularized while at least some of the nucleotides were annealed to DNA that was of placental origin. In some embodiments of the present disclosure, there is a set of probes wherein some of the probes target single tandem repeats, and some of the probes target single nucleotide polymorphisms. In some embodiments, the loci are selected for the purpose of non-invasive prenatal diagnosis. In some embodiments, the probes are used for the purpose of non-invasive prenatal diagnosis. In some embodiments, the loci are targeted using a method that could include circularizing probes, MIPs, capture by hybridization probes, probes on a SNP array, or combinations thereof. In some embodiments, the probes are used as circularizing probes, MIPs, capture by hybridization probes, probes on a SNP array, or combinations thereof. In some embodiments, the loci are sequenced for the purpose of non-invasive prenatal diagnosis.

[0803] In the case where the relative informativeness of a sequence is greater when combined with relevant parent contexts, it follows that maximizing the number of sequence reads that contain a SNP for which the parental context is known may maximize the informativeness of the set of sequencing reads on the mixed sample. In an embodiment, the number of sequence reads that contain a SNP for which the parent contexts are known may be enhanced by using qPCR to preferentially amplify specific sequences. In an embodiment, the number of sequence reads that contain a SNP for which the parent contexts are known may be enhanced by using circularizing probes (for example, MIPs) to preferentially amplify specific sequences. In an embodiment, the number of sequence reads that contain a SNP for which the parent contexts are known may be enhanced by using a capture by hybridization method (for example SURESELECT) to preferentially amplify specific sequences. Different methods may be used to enhance the number of sequence reads that contain a SNP for which the parent contexts are known. In an embodiment, the targeting may be accomplished by extension ligation, ligation without extension, capture by hybridization, or PCR.

[0804] In a sample of fragmented genomic DNA, a fraction of the DNA sequences map uniquely to individual chromosomes; other DNA sequences may be found on different chromosomes. Note that DNA found in plasma,

whether maternal or fetal in origin is typically fragmented, often at lengths under 500 bp. In a typical genomic sample, roughly 3.3% of the mappable sequences will map to chromosome 13; 2.2% of the mappable sequences will map to chromosome 18; 1.35% of the mappable sequences will map to chromosome 21; 4.5% of the mappable sequences will map to chromosome X in a female; 2.25% of the mappable sequences will map to chromosome X (in a male); and 0.73% of the mappable sequences will map to chromosome Y (in a male). These are the chromosomes that are most likely to be aneuploid in a fetus. Also, among short sequences, approximately 1 in 20 sequences will contain a SNP, using the SNPs contained on dbSNP. The proportion may well be higher given that there may be many SNPs that have not been discovered.

[0805] In an embodiment of the present disclosure, targeting methods may be used to enhance the fraction of DNA in a sample of DNA that map to a given chromosome such that the fraction significantly exceeds the percentages listed above that are typical for genomic samples. In an embodiment of the present disclosure, targeting methods may be used to enhance the fraction of DNA in a sample of DNA such that the percentage of sequences that contain a SNP are significantly greater than what may be found in typical for genomic samples. In an embodiment of the present disclosure, targeting methods may be used to target DNA from a chromosome or from a set of SNPs in a mixture of maternal and fetal DNA for the purposes of prenatal diagnosis.

[0806] Note that a method has been reported (U.S. Pat. No. 7,888,017) for determining fetal aneuploidy by counting the number of reads that map to a suspect chromosome and comparing it to the number of reads that map to a reference chromosome, and using the assumption that an overabundance of reads on the suspect chromosome corresponds to a triploidy in the fetus at that chromosome. Those methods for prenatal diagnosis would not make use of targeting of any sort, nor do they describe the use of targeting for prenatal diagnosis.

[0807] By making use of targeting approaches in sequencing the mixed sample, it may be possible to achieve a certain level of accuracy with fewer sequence reads. The accuracy may refer to sensitivity, it may refer to specificity, or it may refer to some combination thereof. The desired level of accuracy may be between 90% and 95%; it may be between 95% and 98%; it may be between 98% and 99%; it may be between 99% and 99.5%; it may be between 99.5% and 99.9%; it may be between 99.9% and 99.99%; it may be between 99.99% and 99.999%, it may be between 99.999% and 100%. Levels of accuracy above 95% may be referred to as high accuracy.

[0808] There are a number of published methods in the prior art that demonstrate how one may determine the ploidy state of a fetus from a mixed sample of maternal and fetal DNA, for example: G. J. W. Liao et al. *Clinical Chemistry* 2011; 57(1) pp. 92-101. These methods focus on thousands of locations along each chromosome. The number of locations along a chromosome that may be targeted while still resulting in a high accuracy ploidy determination on a fetus, for a given number of sequence reads, from a mixed sample of DNA is unexpectedly low. In an embodiment of the present disclosure, an accurate ploidy determination may be made by using targeted sequencing, using any method of targeting, for example qPCR, ligand mediated PCR, other PCR methods, capture by hybridization, or circularizing

probes, wherein the number of loci along a chromosome that need to be targeted may be between 5,000 and 2,000 loci; it may be between 2,000 and 1,000 loci; it may be between 1,000 and 500 loci; it may be between 500 and 300 loci; it may be between 300 and 200 loci; it may be between 200 and 150 loci; it may be between 150 and 100 loci; it may be between 100 and 50 loci; it may be between 50 and 20 loci; it may be between 20 and 10 loci. Optimally, it may be between 100 and 500 loci. The high level of accuracy may be achieved by targeting a small number of loci and executing an unexpectedly small number of sequence reads. The number of reads may be between 100 million and 50 million reads; the number of reads may be between 50 million and 20 million reads; the number of reads may be between 20 million and 10 million reads; the number of reads may be between 10 million and 5 million reads; the number of reads may be between 5 million and 2 million reads; the number of reads may be between 2 million and 1 million; the number of reads may be between 1 million and 500,000; the number of reads may be between 500,000 and 200,000; the number of reads may be between 200,000 and 100,000; the number of reads may be between 100,000 and 50,000; the number of reads may be between 50,000 and 20,000; the number of reads may be between 20,000 and 10,000; the number of reads may be below 10,000. Fewer number of read are necessary for larger amounts of input DNA.

[0809] In some embodiments, there is a composition comprising a mixture of DNA of fetal origin, and DNA of maternal origin, wherein the percent of sequences that uniquely map to chromosome 13 is greater than 4%, greater than 5%, greater than 6%, greater than 7%, greater than 8%, greater than 9%, greater than 10%, greater than 12%, greater than 15%, greater than 20%, greater than 25%, or greater than 30%. In some embodiments of the present disclosure, there is a composition comprising a mixture of DNA of fetal origin, and DNA of maternal origin, wherein the percent of sequences that uniquely map to chromosome 18 is greater than 3%, greater than 4%, greater than 5%, greater than 6%, greater than 7%, greater than 8%, greater than 9%, greater than 10%, greater than 12%, greater than 15%, greater than 20%, greater than 25%, or greater than 30%. In some embodiments of the present disclosure, there is a composition comprising a mixture of DNA of fetal origin, and DNA of maternal origin, wherein the percent of sequences that uniquely map to chromosome 21 is greater than 2%, greater than 3%, greater than 4%, greater than 5%, greater than 6%, greater than 7%, greater than 8%, greater than 9%, greater than 10%, greater than 12%, greater than 15%, greater than 20%, greater than 25%, or greater than 30%. In some embodiments of the present disclosure, there is a composition comprising a mixture of DNA of fetal origin, and DNA of maternal origin, wherein the percent of sequences that uniquely map to chromosome X is greater than 6%, greater than 7%, greater than 8%, greater than 9%, greater than 10%, greater than 12%, greater than 15%, greater than 20%, greater than 25%, or greater than 30%. In some embodiments of the present disclosure, there is a composition comprising a mixture of DNA of fetal origin, and DNA of maternal origin, wherein the percent of sequences that uniquely map to chromosome Y is greater than 1%, greater than 2%, greater than 3%, greater than 4%, greater than 5%, greater than 6%, greater than 7%, greater than 8%, greater

than 9%, greater than 10%, greater than 12%, greater than 15%, greater than 20%, greater than 25%, or greater than 30%.

[0810] In some embodiments, a composition is described comprising a mixture of DNA of fetal origin, and DNA of maternal origin, wherein the percent of sequences that uniquely map to a chromosome, and that contains at least one single nucleotide polymorphism is greater than 0.2%, greater than 0.3%, greater than 0.4%, greater than 0.5%, greater than 0.6%, greater than 0.7%, greater than 0.8%, greater than 0.9%, greater than 1%, greater than 1.2%, greater than 1.4%, greater than 1.6%, greater than 1.8%, greater than 2%, greater than 2.5%, greater than 3%, greater than 4%, greater than 5%, greater than 6%, greater than 7%, greater than 8%, greater than 9%, greater than 10%, greater than 12%, greater than 15%, or greater than 20%, and where the chromosome is taken from the group 13, 18, 21, X, or Y. In some embodiments of the present disclosure, there is a composition comprising a mixture of DNA of fetal origin, and DNA of maternal origin, wherein the percent of sequences that uniquely map to a chromosome and that contain at least one single nucleotide polymorphism from a set of single nucleotide polymorphisms is greater than 0.15%, greater than 0.2%, greater than 0.3%, greater than 0.4%, greater than 0.5%, greater than 0.6%, greater than 0.7%, greater than 0.8%, greater than 0.9%, greater than 1%, greater than 1.2%, greater than 1.4%, greater than 1.6%, greater than 1.8%, greater than 2%, greater than 2.5%, greater than 3%, greater than 4%, greater than 5%, greater than 6%, greater than 7%, greater than 8%, greater than 9%, greater than 10%, greater than 12%, greater than 15%, or greater than 20%, where the chromosome is taken from the set of chromosome 13, 18, 21, X and Y, and where the number of single nucleotide polymorphisms in the set of single nucleotide polymorphisms is between 1 and 10, between 10 and 20, between 20 and 50, between 50 and 100, between 100 and 200, between 200 and 500, between 500 and 1,000, between 1,000 and 2,000, between 2,000 and 5,000, between 5,000 and 10,000, between 10,000 and 20,000, between 20,000 and 50,000, and between 50,000 and 100,000.

[0811] In theory, each cycle in the amplification doubles the amount of DNA present; however, in reality, the degree of amplification is slightly lower than two. In theory, amplification, including targeted amplification, will result in bias free amplification of a DNA mixture; in reality, however, different alleles tend to be amplified to a different extent than other alleles.

[0812] When DNA is amplified, the degree of allelic bias typically increases with the number of amplification steps. In some embodiments, the methods described herein involve amplifying DNA with a low level of allelic bias. Since the allelic bias compounds with each additional cycle, one can determine the per cycle allelic bias by calculating the n th root of the overall bias where n is the base 2 logarithm of degree of enrichment. In some embodiments, there is a composition comprising a second mixture of DNA, where the second mixture of DNA has been preferentially enriched at a plurality of polymorphic loci from a first mixture of DNA where the degree of enrichment is at least 10, at least 100, at least 1,000, at least 10,000, at least 100,000 or at least 1,000,000, and where the ratio of the alleles in the second mixture of DNA at each locus differs from the ratio of the alleles at that locus in the first mixture of DNA by a factor

that is, on average, less than 1,000%, 500%, 200%, 100%, 50%, 20%, 10%, 5%, 2%, 1%, 0.5%, 0.2%, 0.1%, 0.05%, 0.02%, or 0.01%. In some embodiments, there is a composition comprising a second mixture of DNA, where the second mixture of DNA has been preferentially enriched at a plurality of polymorphic loci from a first mixture of DNA where the per cycle allelic bias for the plurality of polymorphic loci is, on average, less than 10%, 5%, 2%, 1%, 0.5%, 0.2%, 0.1%, 0.05%, or 0.02%. In some embodiments, the plurality of polymorphic loci comprises at least 10 loci, at least 20 loci, at least 50 loci, at least 100 loci, at least 200 loci, at least 500 loci, at least 1,000 loci, at least 2,000 loci, at least 5,000 loci, at least 10,000 loci, at least 20,000 loci, or at least 50,000 loci.

Some Embodiments

[0813] In some embodiments, a method is disclosed herein for generating a report disclosing the determined ploidy status of a chromosome in a gestating fetus, the method comprising: obtaining a first sample that contains DNA from the mother of the fetus and DNA from the fetus; obtaining genotypic data from one or both parents of the fetus; preparing the first sample by isolating the DNA so as to obtain a prepared sample; measuring the DNA in the prepared sample at a plurality of polymorphic loci; calculating, on a computer, allele counts or allele count probabilities at the plurality of polymorphic loci from the DNA measurements made on the prepared sample; creating, on a computer, a plurality of ploidy hypotheses concerning expected allele count probabilities at the plurality of polymorphic loci on the chromosome for different possible ploidy states of the chromosome; building, on a computer, a joint distribution model for allele count probability of each polymorphic locus on the chromosome for each ploidy hypothesis using genotypic data from the one or both parents of the fetus; determining, on a computer, a relative probability of each of the ploidy hypotheses using the joint distribution model and the allele count probabilities calculated for the prepared sample; calling the ploidy state of the fetus by selecting the ploidy state corresponding to the hypothesis with the greatest probability; and generating a report disclosing the determined ploidy status.

[0814] In some embodiments, the method is used to determine the ploidy state of a plurality of gestating fetuses in a plurality of respective mothers, the method further comprising: determining the percent of DNA that is of fetal origin in each of the prepared samples; and wherein the step of measuring the DNA in the prepared sample is done by sequencing a number of DNA molecules in each of the prepared samples, where more molecules of DNA are sequenced from those prepared samples that have a smaller fraction of fetal DNA than those prepared samples that have a larger fraction of fetal DNA.

[0815] In some embodiments, the method is used to determine the ploidy state of a plurality of gestating fetuses in a plurality of respective mothers, and where the measuring the DNA in the prepared sample is done, for each of the fetuses, by sequencing a first fraction of the prepared sample of DNA to give a first set of measurements, the method further comprising: making a first relative probability determination for each of the ploidy hypotheses for each of the fetuses, given the first set of DNA measurements; resequencing a second fraction of the prepared sample from those fetuses where the first relative probability determination for each of

the ploidy hypotheses indicates that a ploidy hypothesis corresponding to an aneuploid fetus has a significant but not conclusive probability, to give a second set of measurements; making a second relative probability determination for ploidy hypotheses for the fetuses using the second set of measurements and optionally also the first set of measurements; and calling the ploidy states of the fetuses whose second sample was resequenced by selecting the ploidy state corresponding to the hypothesis with the greatest probability as determined by the second relative probability determination.

[0816] In some embodiments, a composition of matter is disclosed, the composition of matter comprising: a sample of preferentially enriched DNA, wherein the sample of preferentially enriched DNA has been preferentially enriched at a plurality of polymorphic loci from a first sample of DNA, wherein the first sample of DNA consisted of a mixture of maternal DNA and fetal DNA derived from maternal plasma, where the degree of enrichment is at least a factor of 2, and wherein the allelic bias between the first sample and the preferentially enriched sample is, on average, selected from the group consisting of less than 2%, less than 1%, less than 0.5%, less than 0.2%, less than 0.1%, less than 0.05%, less than 0.02%, and less than 0.01%. In some embodiments, a method is disclosed to create a sample of such preferentially enriched DNA.

[0817] In some embodiment, a method is disclosed for determining the presence or absence of a fetal aneuploidy in a maternal tissue sample comprising fetal and maternal genomic DNA, wherein the method comprises: (a) obtaining a mixture of fetal and maternal genomic DNA from said maternal tissue sample; (b) selectively enriching the mixture of fetal and maternal DNA at a plurality of polymorphic alleles; (c) distributing selectively enriched fragments from the mixture of fetal and maternal genomic DNA of step a to provide reaction samples comprising a single genomic DNA molecule or amplification products of a single genomic DNA molecule; (d) conducting massively parallel DNA sequencing of the selectively enriched fragments of genomic DNA in the reaction samples of step c) to determine the sequence of said selectively enriched fragments; (e) identifying the chromosomes to which the sequences obtained in step d) belong; (f) analyzing the data of step d) to determine i) the number of fragments of genomic DNA from step d) that belong to at least one first target chromosome that is presumed to be diploid in both the mother and the fetus, and ii) the number of fragments of genomic DNA from step d) that belong to a second target chromosome, wherein said second chromosome is suspected to be aneuploid in the fetus; (g) calculating an expected distribution of the number of fragments of genomic DNA from step d) for the second target chromosome if the second target chromosome is euploid, using the number determined in step f) part i); (h) calculating an expected distribution of the number of fragments of genomic DNA from step d) for the second target chromosome if the second target chromosome is aneuploid, using the first number is step f) part i) and an estimated fraction of fetal DNA found in the mixture of step b); and (i) using a maximum likelihood or maximum a posteriori approach to determine whether the number of fragments of genomic DNA determined in step f) part ii) is more likely to be part of the distribution calculated in step g) or the distribution calculated in step h); thereby indicating the presence or absence of a fetal aneuploidy.

Exemplary Cancer Diagnostic Methods

[0818] Note that it has been demonstrated that DNA that originated from cancer that is living in a host can be found in the blood of the host. In the same way that genetic diagnoses can be made from the measurement of mixed DNA found in maternal blood, genetic diagnoses can equally well be made from the measurement of mixed DNA found in host blood. The genetic diagnoses may include aneuploidy states, or gene mutations. Any claim in the instant disclosure that reads on determining the ploidy state or genetic state of a fetus from the measurements made on maternal blood can equally well read on determining the ploidy state or genetic state of a cancer from the measurements on host blood.

[0819] In some embodiments, a method of the present disclosure allows one to determine the ploidy status of a cancer, the method including obtaining a mixed sample that contains genetic material from the host, and genetic material from the cancer; measuring the DNA in the mixed sample; calculating the fraction of DNA that is of cancer origin in the mixed sample; and determining the ploidy status of the cancer using the measurements made on the mixed sample and the calculated fraction. In some embodiments, the method may further include administering a cancer therapeutic based on the determination of the ploidy state of the cancer. In some embodiments, the method may further include administering a cancer therapeutic based on the determination of the ploidy state of the cancer, wherein the cancer therapeutic is taken from the group comprising a pharmaceutical, a biologic therapeutic, and antibody based therapy and combination thereof.

Exemplary Clinical Actions

[0820] In some embodiments, any of the methods include taking a clinical action based on a result of a method of the invention (such as the determination of the presence or absence of a polymorphism or mutation, ploidy state, or paternity). In some embodiments in which an embryo or fetus has one or more one or more polymorphisms or mutations of interest (such as a CNV) based on a result of a method of the invention, the clinical action includes performing additional testing (such as testing to confirm the presence of the polymorphism or mutation), not implanting the embryo for IVF, implanting a different embryo for IVF, terminating a pregnancy, preparing for a special needs child, or undergoing an intervention designed to decrease the severity of the phenotypic presentation of a genetic disorder. In some embodiments, the clinical action is selected from the group consisting of performing an ultrasound, amniocentesis on the fetus, amniocentesis on a subsequent fetus that inherits genetic material from the mother and/or father, chorion villus biopsy on the fetus, chorion villus biopsy on a subsequent fetus that inherits genetic material from the mother and/or father, in vitro fertilization, preimplantation genetic diagnosis on one or more embryos that inherited genetic material from the mother and/or father, karyotyping on the mother, karyotyping on the father, fetal echocardiogram (such as an echocardiogram of a fetus with trisomy 21, 18, or 13, monosomy X, or a microdeletion), and combinations thereof. In some embodiments, the clinical action is selected from the group consisting of administering growth hormone to a born child with monosomy X (such as administration starting at ~9 months), administering calcium to a

born child with a 22q deletion (such as DiGeorge syndrome), administering an androgen such as testosterone to a born child with 47,XXY (such as one injection per month for 3 months of 25 mg testosterone enanthate to an infant or toddler), performing a test for cancer on a woman with a complete or partial molar pregnancy (such as a triploid fetus), administering a therapy for cancer such as a chemotherapeutic agent to a woman with a complete or partial molar pregnancy (such as a triploid fetus), screening a fetus determined to be male (such as a fetus determined to be male using a method of the invention) for one or more X-linked genetic disorders such as Duchenne muscular dystrophy (DMD), adrenoleukodystrophy, or hemophilia, performing amniocentesis on a male fetus at risk for an X-linked disorder, administering dexamethasone to a woman with a female fetus at risk male (such as a fetus determined to be female using a method of the invention) for congenital adrenal hyperplasia, performing amniocentesis on a female fetus at risk for congenital adrenal hyperplasia, administering killed vaccines (instead of live vaccines) or not administering certain vaccines to a born child who is (or is suspected of being) immune deficient from a 22q11.2 deletion, performing occupational and/or physical therapy, performing early intervention in education, delivering the baby at a tertiary care center with a NICU and/or having pediatric specialists available at delivery, behavioral intervention for born child (such as a child with XXX, XXY, or XYY), and combinations thereof.

[0821] In some embodiments, ultrasound or another screening test is performed on a women determined to have multiple pregnancies (such as twins) to determine whether or not two or more of the fetus are monochorionic. Monozygotic twins result from ovulation and fertilization of a single oocyte, with subsequent division of the zygote; placental may be dichorionic or monochorionic. Dizygotic twins occur from ovulation and fertilization of two oocytes, which usually results in dichorionic placental. Monochorionic twins have a risk of twin-to-twin transfusion syndrome, which may cause unequal distribution of blood between fetuses that results in differences in their growth and development, sometimes resulting in stillbirth. Thus, twins determined to be monozygotic twins using a method of the invention are desirably tested (such as by ultrasound) to determine if they are monochorionic twins, and if so, these twins can be monitored (such as bi-weekly ultrasounds from 16 weeks) for signs of win-to-twin transfusion syndrome.

[0822] In some embodiments in which an embryo or fetus does not have one or more one or more polymorphisms or mutations of interest (such as a CNV) based on a result of a method of the invention, the clinical action includes implanting the embryo for IVF or continuing a pregnancy. In some embodiments, the clinical action is additional testing to confirm the absence of the polymorphism or mutation selected from the group consisting of performing an ultrasound, amniocentesis, chorion villus biopsy, and combinations thereof.

[0823] In some embodiments in which an individual has one or more polymorphisms or mutations (such as a polymorphism or mutation associated with a disease or disorder such as cancer or an increased risk for a disease or disorder such as cancer) based on a result of a method of the invention, the clinical action includes performing additional testing or administering one or more therapies for a disease or disorder (such as a therapy for cancer, a therapy for the

specific type of cancer or type of mutation the individual is diagnosed with, or any of the therapies disclosed herein). In some embodiments, the clinical action is additional testing to confirm the presence or absence of a polymorphism or mutation selected from the group consisting of biopsy, surgery, medical imaging (such as a mammogram or an ultrasound), and combinations thereof.

[0824] In some embodiments, the additional testing includes performing the same or a different method (such as any of the methods described herein) to confirm the presence or absence of the polymorphism or mutation (such as a CNV), such as testing either a second fraction of the same sample that was tested or a different sample from the same individual (such as the same pregnant mother, fetus, embryo, or individual at increased risk for cancer). In some embodiments, the additional testing is performed for an individual for whom the probability of a polymorphism or mutation (such as a CNV) is above a threshold value. In some embodiments, the additional testing is performed for an individual for whom the confidence or z-score for the determination of a polymorphism or mutation (such as a CNV) is above a threshold value (such as additional testing to confirm the presence of a likely polymorphism or mutation). In some embodiments, the additional testing is performed for an individual for whom the confidence or z-score for the determination of a polymorphism or mutation (such as a CNV) is between minimum and maximum threshold values (such as additional testing to increase the confidence that the initial result is correct). In some embodiments, the additional testing is performed for an individual for whom the confidence for the determination of the presence or absence of a polymorphism or mutation (such as a CNV) is below a threshold value (such as a "no call" result due to not being able to determine the presence or absence of the CNV with sufficient confidence). An exemplary Z core is calculated in Chiu et al. BMJ 2011; 342:c7401 (which is hereby incorporated by reference in its entirety) in which chromosome 21 is used as an example and can be replaced with any other chromosome or chromosome segment in the test sample.

$$\begin{aligned} \text{Z score for percentage chromosome 21 in test case} = & \\ & ((\text{percentage chromosome 21 in test case}) - \\ & (\text{mean percentage chromosome 21 in reference} \\ & \text{controls})) / (\text{standard deviation of percentage} \\ & \text{chromosome 21 in reference controls}). \end{aligned}$$

[0825] In some embodiments, the additional testing is performed for an individual for whom the initial sample did not meet quality control guidelines or had a fetal fraction or a tumor fraction below a threshold value. In some embodiments, the method includes selecting an individual for additional testing based on the result of a method of the invention, the probability of the result, the confidence of the result, or the z-score; and performing the additional testing on the individual (such as on the same or a different sample). In some embodiments, a subject diagnosed with a disease or disorder (such as cancer) undergoes repeat testing using a method of the invention or known testing for the disease or disorder at multiple time points to monitor the progression of the disease or disorder or the remission or reoccurrence of the disease or disorder.

Exemplary Implementation Methods

[0826] Any of the embodiments disclosed herein may be implemented in digital electronic circuitry, integrated cir-

cuitry, specially designed ASICs (application-specific integrated circuits), computer hardware, firmware, software, or in combinations thereof. Apparatus of the presently disclosed embodiments can be implemented in a computer program product tangibly embodied in a machine-readable storage device for execution by a programmable processor; and method steps of the presently disclosed embodiments can be performed by a programmable processor executing a program of instructions to perform functions of the presently disclosed embodiments by operating on input data and generating output. The presently disclosed embodiments can be implemented advantageously in one or more computer programs that are executable and/or interpretable on a programmable system including at least one programmable processor, which may be special or general purpose, coupled to receive data and instructions from, and to transmit data and instructions to, a storage system, at least one input device, and at least one output device. Each computer program can be implemented in a high-level procedural or object-oriented programming language or in assembly or machine language if desired; and in any case, the language can be a compiled or interpreted language. A computer program may be deployed in any form, including as a stand-alone program, or as a module, component, subroutine, or other unit suitable for use in a computing environment. A computer program may be deployed to be executed or interpreted on one computer or on multiple computers at one site, or distributed across multiple sites and interconnected by a communication network.

[0827] Computer readable storage media, as used herein, refers to physical or tangible storage (as opposed to signals) and includes without limitation volatile and non-volatile, removable and non-removable media implemented in any method or technology for the tangible storage of information such as computer-readable instructions, data structures, program modules or other data. Computer readable storage media includes, but is not limited to, RAM, ROM, EPROM, EEPROM, flash memory or other solid state memory technology, CD-ROM, DVD, or other optical storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other physical or material medium which can be used to tangibly store the desired information or data or instructions and which can be accessed by a computer or processor.

[0828] In some embodiments, the invention features a computer configured to accomplish one or more of the in vitro methods described herein. In some embodiments, the data is analyzed by the computer system as described herein. In some embodiments, genetic data (such as sequencing or microarray data) from at least 100; 200; 500; 750; 1,000; 2,000; 5,000; 7,500; 10,000; 20,000; 25,000; 30,000; 40,000; 50,000; 75,000; or 100,000 different loci is analyzed by the computer is less than 200, 100, 60, 30, 20, 10, 5, or 1 minute, or in less than 30 or 10 seconds to detect the present or absence of a mutation (such as a CNV or SNV) at the loci.

[0829] Any of the methods described herein may include the output of data in a physical format, such as on a computer screen, or on a paper printout. In explanations of any embodiments elsewhere in this document, it should be understood that the described methods may be combined with the output of the actionable data in a format that can be acted upon by a physician. In addition, the described methods may be combined with the actual execution of a clinical

decision that results in a clinical treatment, or the execution of a clinical decision to make no action. Some of the embodiments described in the document for determining genetic data pertaining to a target individual may be combined with the decision to select one or more embryos for transfer in the context of IVF, optionally combined with the process of transferring the embryo to the womb of the prospective mother. Some of the embodiments described in the document for determining genetic data pertaining to a target individual may be combined with the notification of a potential chromosomal abnormality, or lack thereof, with a medical professional, optionally combined with the decision to abort, or to not abort, a fetus in the context of prenatal diagnosis. Some of the embodiments described herein may be combined with the output of the actionable data, and the execution of a clinical decision that results in a clinical treatment, or the execution of a clinical decision to make no action.

Exemplary Diagnostic Boxes

[0830] In an embodiment, the present disclosure comprises a diagnostic box that is capable of partly or completely carrying out any of the methods described in this disclosure. In an embodiment, the diagnostic box may be located at a physician's office, a hospital laboratory, or any suitable location reasonably proximal to the point of patient care. The box may be able to run the entire method in a wholly automated fashion, or the box may require one or a number of steps to be completed manually by a technician. In an embodiment, the box may be able to analyze at least the genotypic data measured on the maternal plasma. In an embodiment, the box may be linked to means to transmit the genotypic data measured on the diagnostic box to an external computation facility which may then analyze the genotypic data, and possibly also generate a report. The diagnostic box may include a robotic unit that is capable of transferring aqueous or liquid samples from one container to another. It may comprise a number of reagents, both solid and liquid. It may comprise a high throughput sequencer. It may comprise a computer.

Experimental Section

[0831] The presently disclosed embodiments are described in the following Examples, which are set forth to aid in the understanding of the disclosure, and should not be construed to limit in any way the scope of the disclosure as defined in the claims which follow thereafter. The following examples are put forth so as to provide those of ordinary skill in the art with a complete disclosure and description of how to use the described embodiments, and are not intended to limit the scope of the disclosure nor are they intended to represent that the Examples below are all or the only experiments performed. Efforts have been made to ensure accuracy with respect to numbers used (e.g. amounts, temperature, etc.) but some experimental errors and deviations should be accounted for. Unless indicated otherwise, parts are parts by volume, and temperature is in degrees Centigrade. It should be understood that variations in the methods as described may be made without changing the fundamental aspects that the Examples are meant to illustrate.

Example 1

[0832] The objective was to show that a Bayesian maximum likelihood estimation (MLE) algorithm that uses par-

ent genotypes to calculate fetal fraction improves accuracy of non-invasive prenatal trisomy diagnosis compared to published methods.

[0833] Simulated sequencing data for maternal cfDNA was created by sampling reads obtained on trisomy-21 and respective mother cell lines. The rate of correct disomy and trisomy calls were determined from 500 simulations at various fetal fractions for a published method (Chiu et al. BMJ 2011; 342:c7401) and our MLE-based algorithm. We validated the simulations by obtaining 5 million shotgun reads from four pregnant mothers and respective fathers collected under an IRB-approved protocol. Parental genotypes were obtained on a 290K SNP array. (See FIG. 14)

[0834] In simulations, the MLE-based approach achieved 99.0% accuracy for fetal fractions as low as 9% and reported confidences that corresponded well to overall accuracy. We validated these results using four real samples wherein we obtained all correct calls with a computed confidence exceeding 99%. In contrast, our implementation of the published algorithm for Chiu et al. required 18% fetal fraction to achieve 99.0% accuracy, and achieved only 87.8% accuracy at 9% fetal DNA.

[0835] Fetal fraction determination from parental genotypes in conjunction with a MLE-based approach achieves greater accuracy than published algorithms at the fetal fractions expected during the 1st and early 2nd trimester. Furthermore, the method disclosed herein produces a confidence metric that is crucial in determining the reliability of the result, especially at low fetal fractions where ploidy detection is more difficult. Published methods use a less accurate threshold method for calling ploidy based on large sets of disomy training data, an approach that predefines a false positive rate. In addition, without a confidence metric, published methods are at risk of reporting false negative results when there is insufficient fetal cfDNA to make a call.

[0836] In some embodiments, a confidence estimate is calculated for the called ploidy state.

Example 2

[0837] The objective was to improve non-invasive detection of fetal trisomy 18, 21, and X particularly in samples consisting of low fetal fraction by using a targeted sequencing approach combined with parent genotypes and Hapmap data in a Bayesian Maximum Likelihood Estimation (MLE) algorithm.

[0838] Maternal samples from four euploid and two trisomy-positive pregnancies and respective paternal samples were obtained under an IRB-approved protocol from patients where fetal karyotype was known. Maternal cfDNA was extracted from plasma and roughly 10 million sequence reads were obtained following preferential enrichment that targeted specific SNPs. Parent samples were similarly sequenced to obtain genotypes.

[0839] The described algorithm correctly called chromosome 18 and 21 disomy for all euploid samples and normal chromosomes of aneuploid samples. Trisomy 18 and 21 calls were correct, as were chromosome X copy numbers in male and female fetuses. The confidence produced by the algorithm was in excess of 98% in all cases.

[0840] The method described accurately reported the ploidy of all tested chromosomes from six samples, including samples comprised of less than 12% fetal DNA, which account for roughly 30% of 1st and early 2nd-trimester samples. The crucial difference between the instant MLE

algorithm and published methods is that it leverages parent genotypes and Hapmap data to improve accuracy and generate a confidence metric. At low fetal fractions, all methods become less accurate; it is important to correctly identify samples without sufficient fetal cfDNA to make a reliable call. Others have used chromosome Y specific probes to estimate fetal fraction of male fetuses, but concurrent parental genotyping enables estimation of fetal fraction for both sexes. Another inherent limitation of published methods using untargeted shotgun sequencing is that accuracy of ploidy calling varies among chromosomes due to differences in factors such as GC richness. The instant targeted sequencing approach is largely independent of such chromosome-scale variations and yields more consistent performance between chromosomes.

Example 3

[0841] The objective was to determine if trisomy is detectable with high confidence on a triploid fetus, using novel informatics to analyze SNP loci of free floating fetal DNA in maternal plasma.

[0842] 20 mL of blood was drawn from a pregnant patient following abnormal ultrasound. After centrifugation, maternal DNA was extracted from the buffy coat (DNEASY, QIAGEN); cell-free DNA was extracted from plasma (QIAAMP QIAGEN). Targeted sequencing was applied to SNP loci on chromosomes 2, 21, and X in both DNA samples. Maximum-Likelihood Bayesian estimation selected the most likely hypothesis from the set of all possible ploidy states. The method determines fetal DNA fraction, ploidy state and explicit confidences in the ploidy determination. No assumptions are made about the ploidy of a reference chromosome. The diagnostic uses a test statistic that is independent of sequence read counts, which is the recent state of the art.

[0843] The instant method accurately diagnosed trisomy of chromosomes 2 and 21. Child fraction was estimated at 11.9% [CI 11.7-12.1]. The fetus was found to have one maternal and two paternal copies of chromosomes 2 and 21 with confidence of effectively 1 (error probability < 10⁻³⁰). This was achieved with 92,600 and 258,100 reads on chromosomes 2 and 21 respectively.

[0844] This is the first demonstration of non-invasive prenatal diagnosis of trisomic chromosomes from maternal blood where the fetus was triploid, as confirmed by metaphase karyotype. Extant methods of non-invasive diagnosis would not detect aneuploidy in this sample. Current methods rely on a surplus of sequence reads on a trisomic chromosome relative to disomic reference chromosomes; but a triploid fetus has no disomic reference. Furthermore, extant methods would not achieve similarly high-confidence ploidy determination with this fraction of fetal DNA and number of sequence reads. It is straightforward to extend the approach to all 24 chromosomes.

Example 4

[0845] The following protocol was used for 800-plex amplification of DNA isolated from maternal plasma from a euploid pregnancy and also genomic DNA from a triploidy 21 cell line using standard PCR (meaning no nesting was used). Library preparation and amplification involved single tube blunt ending followed by A-tailing. Adaptor ligation was run using the ligation kit found in the AGILENT

SURESELECT kit, and PCR was run for 7 cycles. Then, 15 cycles of STA (95° C. for 30 s; 72° C. for 1 min; 60° C. for 4 min; 65° C. for 1 min; 72° C. for 30 s) using 800 different primer pairs targeting SNPs on chromosomes 2, 21 and X. The reaction was run with 12.5 nM primer concentration. The DNA was then sequenced with an ILLUMINA IIGAX sequencer. The sequencer output 1.9 million reads, of which 92% mapped to the genome; of those reads that mapped to the genome, more than 99% mapped to one of the regions targeted by the targeted primers. The numbers were essentially the same for both the plasma DNA and the genomic DNA. FIG. 15 shows the ratio of the two alleles for the ~780 SNPs that were detected by the sequencer in the genomic DNA that was taken from a cell line with known trisomy at chromosome 21. Note that the allele ratios are plotted here for ease of visualization, because the allele distributions are not straightforward to read visually. The circles represent SNPs on disomic chromosomes, while the stars represent SNPs on a trisomic chromosome. FIG. 16 is another representation of the same data as in FIG. 15, where the Y-axis is the relative number of A and B measured for each SNP, and where the X-axis is the SNP number where the SNPs are separated by chromosome. In FIG. 16, SNP 1 to 312 are found on chromosome 2, from SNP 313 to 605 are found on chromosome 21 which is trisomic, and from SNP 606 to 800 are on chromosome X. The data from chromosomes 2 and X show a disomic chromosome, as the relative sequence counts lie in three clusters: AA at the top of the graph, BB at the bottom of the graph, and AB in the middle of the graph. The data from chromosome 21, which is trisomic, shows four clusters: AAA at the top of the graph, AAB around the 0.65 line (2/3), ABB around the 0.35 line (1/3), and BBB at the bottom of the graph.

[0846] FIGS. 17A-D show data for the same 800-plex protocol, but measured on DNA that was amplified from four plasma samples from pregnant women. For these four samples, we expect to see seven clusters of dots: (1) along the top of the graph are those loci where both the mother and the fetus are AA, (2) slightly below the top of the graph are those loci where the mother is AA and the fetus is AB, (3) slightly above the 0.5 line are those loci where the mother is AB and the fetus is AA, (4) along the 0.5 line are those loci where the mother and the fetus are both AB, (5) slightly below the 0.5 line are those loci where the mother is AB and the fetus is BB, (6) slightly above the bottom of the graph are those loci where the mother is BB and the fetus is AB, (7) along the bottom of the graph are those loci where both the mother and the fetus are BB. The smaller the fetal fraction, the less the separation between clusters (1) and (2), between clusters (3), (4) and (5), and between clusters (6) and (7). The separation is expected to be half of the fraction of DNA that is of fetal origin. For example if the DNA is 20% fetal, and 80% maternal, we expect (1) through (7) to be centered at 1.0, 0.9, 0.6, 0.5, 0.4, 0.1 and 0.0 respectively; see for example FIG. 17D, POOL1_BC5_ref_rate. If, instead the DNA is 8% fetal, and 92% maternal, we expect (1) through (7) to be centered at 1.00, 0.96, 0.54, 0.50, 0.46, 0.04 and 0.00 respectively; see for example FIG. 17B, POOL1_BC2_ref_rate. If there is not fetal DNA detected, we do not expect to see (2), (3), (5), or (6); alternately we could say that the separation is zero, and therefore (1) and (2) are on top of each other, as are (3), (4) and (5), and also

(6) and (7); see e.g. FIG. 17C, POOL1_BC7_ref_rate. Note that the fetal fraction for FIG. 17A, POOL1_BC1_ref_rate is about 25%.

Example 5

[0847] Most methods of DNA amplification and measurement will produce some allele bias, wherein the two alleles that are typically found at a locus are detected with intensities or counts that are not representative of the actual amounts of alleles in the sample of DNA. For example, for a single individual, at a heterozygous locus we expect to see a 1:1 ratio of the two alleles, which is the theoretical ratio expected for a heterozygous locus; however due to allele bias, we may see 55:45, or even 60:40. Also note that in the context of sequencing, if the depth of read is low, then simple stochastic noise could result in significant allele bias. In an embodiment, it is possible to model the behavior of each SNP such that if a consistent bias is observed for particular alleles, this bias can be corrected for. FIG. 18 shows the fraction of data that can be explained by binomial variance, before and after bias correction. In FIG. 18, the stars represent the observed allele bias on raw sequence data for the 800-plex experiment; the circles represent the allele bias after correction. Note that if there were no allele bias at all, we would expect the data to fall along the $x=y$ line. A similar set of data that was produced by amplifying DNA using a 150-plex targeted amplification produced data that fell very closely on the 1:1 line after bias correction.

Example 6

[0848] Universal amplification of DNA using ligated adaptors with primers specific to the adaptor tags, where the primer annealing and extension times are limited to a few minutes has the effect of enriching the proportion of shorter DNA strands. Most library protocols designed for creating DNA libraries suitable for sequencing contain such a step, and example protocols are published and well known to those in the art. In some embodiments of the invention, adaptors with a universal tag are ligated to the plasma DNA, and amplified using primers specific to the adaptor tag. In some embodiments, the universal tag can be the same tag as used for sequencing, it can be a universal tag only for PCR amplification, or it can be a set of tags. Since the fetal DNA is typically short in nature, while the maternal DNA can be both short and long in nature, this method has the effect of enriching the proportion of fetal DNA in the mixture. The free floating DNA, thought to be DNA from apoptotic cells, and which contains both fetal and maternal DNA, is short—mostly under 200 bp. Cellular DNA released by cell lysis, a common phenomenon after phlebotomy, is typically almost exclusively maternal, and is also quite long—mostly above 500 bp. Therefore, blood samples that have sat around for more than a few minutes will contain a mixture of short (fetal+maternal) and longer (maternal) DNA. Performing a universal amplification with relatively short extension times on maternal plasma followed by targeted amplification will tend to increase the relative proportion of fetal DNA when compared to the plasma that has been amplified using targeted amplification alone. This can be seen in FIG. 19 which shows the measured fetal percent when the input is plasma DNA (vertical axis) vs. the measured fetal percent when the input DNA is plasma DNA that has had a library prepared using the ILLUMINA GAIIX library preparation

protocol. All the dots fall below the line, indicating that the library preparation step enriches the fraction of DNA that is of fetal origin. Two samples of plasma that were red, indicating hemolysis and therefore that there would be an increased amount of long maternal DNA present from cell lysis, show a particularly significant enrichment of fetal fraction when the library preparation is performed prior to targeted amplification. The method disclosed herein is particularly useful in cases where there is hemolysis or some other situation has occurred where cells comprising relatively long strands of contaminating DNA have lysed, contaminating the mixed sample of short DNA with the long DNA. Typically the relatively short annealing and extension times are between 30 seconds and 2 minutes, though they could be as short as 5 or 10 seconds or less, or as long as 5 or 10 minutes.

Example 7

[0849] The following protocol was used for 1,200-plex amplification of DNA isolated from maternal plasma from a euploid pregnancy and also genomic DNA from a triploidy 21 cell line using a direct PCR protocol, and also a semi-nested approach. Library preparation and amplification involved single tube blunt ending followed by A-tailing. Adaptor ligation was run using a modification of the ligation kit found in the AGILENT SURESELECT kit, and PCR was run for 7 cycles. In the targeted primer pool, there were 550 assays for SNPs from chromosome 21, and 325 assays for SNPs from each of chromosomes 1 and X. Both protocols involved 15 cycles of STA (95° C. for 30 s; 72° C. for 1 min; 60° C. for 4 min; 65° C. for 30 s; 72° C. for 30 s) using 16 nM primer concentration. The semi-nested PCR protocol involved a second amplification of 15 cycles of STA (95° C. for 30 s; 72° C. for 1 min; 60° C. for 4 min; 65° C. for 30 s; 72° C. for 30 s) using an inner forward tag concentration of 29 nM, and a reverse tag concentration of 1 uM or 0.1 uM. The DNA was then sequenced with an ILLUMINA IIGAX sequencer. For the direct PCR protocol, 73% of the reads map to the genome; for the semi-nested protocol, 97.2% of the sequence reads map to the genome. Therefore, the semi-nested protocol result in approximately 30% more information, presumably mostly due to the elimination of primers that are most likely to cause primer dimers.

[0850] The depth of read variability tends to be higher when using the semi-nested protocol than when the direct PCR protocol is used (see FIG. 20) where the diamonds refer to the depth of read for loci run with the semi-nested protocol, and the squares refer to the depth of read for loci run with no nesting. The SNPs are arranged by depth of read for the diamonds, so the diamonds all fall on a curved line, while the squares appear to be loosely correlated; the arrangements of the SNPs is arbitrary, and it is the height of the dot that denotes depth of read rather than its location left to right.

[0851] In some embodiments, the methods described herein can achieve excellent depth of read (DOR) variances. For example, in one version of this Example (FIG. 21) using a 1,200-plex direct PCR amplification of genomic DNA, of the 1,200 assays: 1186 assays had a DOR greater than 10; the average depth of read was 400; 1063 assays (88.6%) had a depth of read of between 200 and 800, and ideal window where the number of reads for each allele is high enough to give meaningful data, while the number of reads for each allele is not so high that the marginal use of those reads was

particularly small. Only 12 alleles had higher depth of read with the highest at 1035 reads. The standard deviation of the DOR was 290, the average DOR was 453, the coefficient of variance of the DOR was 64%, there were 950,000 total reads, and 63.1% of the reads mapped to the genome. In another experiment (FIG. 22) using a 1,200-plex semi-nested protocol, the DOR was higher. The standard deviation of the DOR was 583, the average DOR was 630, the coefficient of variance of the DOR was 93%, there were 870,000 total reads, and 96.3% of the reads mapped to the genome. Note, in both these cases, the SNPs are arranged by the depth of read for the mother, so the curved line represents the maternal depth of read. The differentiation between child and father is not significant; it is only the trend that is significant for the purpose of this explanation.

Example 8

[0852] In an experiment, the semi-nested 1,200-plex PCR protocol was used to amplify DNA from one cell and from three cells. This experiment is relevant to prenatal aneuploidy testing using fetal cells isolated from maternal blood, or for preimplantation genetic diagnosis using biopsied blastomeres or trophoctoderm samples. There were 3 replicates of 1 and 3 cells from 2 individuals (46 XY and 47 XX+21) per condition. Assays targeted chromosomes 1, 21 and X. Three different lysis methods were used: ARCTURUS, MPERv2 and Alkaline lysis. Sequencing was run multiplexing 48 samples in one sequencing lane. The algorithm returned correct ploidy calls for each of the three chromosomes, and for each of the replicates.

Example 9

[0853] In one experiment, four maternal plasma samples were prepared and amplified using a hemi-nested 9,600-plex protocol. The samples were prepared in the following way: Up to 40 mL of maternal blood were centrifuged to isolate the buffy coat and the plasma. The genomic DNA in the maternal sample was prepared from the buffy coat and paternal DNA was prepared from a blood sample or saliva sample. Cell-free DNA in the maternal plasma was isolated using the QIAGEN CIRCULATING NUCLEIC ACID kit and eluted in 45 uL TE buffer according to manufacturer's instructions. Universal ligation adapters were appended to the end of each molecule of 35 uL of purified plasma DNA and libraries were amplified for 7 cycles using adaptor specific primers. Libraries were purified with AGENCOURT AMPURE beads and eluted in 50 uL water.

[0854] 3 uL of the DNA was amplified with 15 cycles of STA (95° C. for 10 min for initial polymerase activation, then 15 cycles of 95° C. for 30 s; 72° C. for 10 s; 65° C. for 1 min; 60° C. for 8 min; 65° C. for 3 min and 72° C. for 30 s; and a final extension at 72° C. for 2 min) using 14.5 nM primer concentration of 9600 target-specific tagged reverse primers and one library adaptor specific forward primer at 500 nM.

[0855] The hemi-nested PCR protocol involved a second amplification of a dilution of the first STAs product for 15 cycles of STA (95° C. for 10 min for initial polymerase activation, then 15 cycles of 95° C. for 30 s; 65° C. for 1 min; 60° C. for 5 min; 65° C. for 5 min and 72° C. for 30 s; and a final extension at 72° C. for 2 min) using reverse tag concentration of 1000 nM, and a concentration of 16.6 uM for each of 9600 target-specific forward primers.

[0856] An aliquot of the STA products was then amplified by standard PCR for 10 cycles with 1 uM of tag-specific forward and barcoded reverse primers to generate barcoded sequencing libraries. An aliquot of each library was mixed with libraries of different barcodes and purified using a spin column.

[0857] In this way, 9,600 primers were used in the single-well reactions; the primers were designed to target SNPs found on chromosomes 1, 2, 13, 18, 21, X and Y. The amplicons were then sequenced using an ILLUMINA GAIIIX sequencer. Per sample, approximately 3.9 million reads were generated by the sequencer, with 3.7 million reads mapping to the genome (94%), and of those, 2.9 million reads (74%) mapped to targeted SNPs with an average depth of read of 344 and a median depth of read of 255. The fetal fraction for the four samples was found to be 9.9%, 18.9%, 16.3%, and 21.2%

[0858] Relevant maternal and paternal genomic DNA samples amplified using a semi-nested 9600-plex protocol and sequenced. The semi-nested protocol is different in that it applies 9,600 outer forward primers and tagged reverse primers at 7.3 nM in the first STA. Thermocycling conditions and composition of the second STA, and the barcoding PCR were the same as for the hemi-nested protocol.

[0859] The sequencing data was analyzed using informatics methods disclosed herein and the ploidy state was called at six chromosomes for the fetuses whose DNA was present in the 4 maternal plasma samples. The ploidy calls for all 28 chromosomes in the set were called correctly with confidences above 99.2% except for one chromosome that was called correctly, but with a confidence of 83%.

[0860] FIG. 23 shows the depth of read of the 9,600-plex hemi-nesting approach along with the depth of read of the 1,200-plex semi-nested approach described in Example 7, though the number of SNPs with a depth of read greater than 100, greater than 200 and greater than 400 was significantly higher than in the 1,200-plex protocol. The number of reads at the 90th percentile can be divided by the number of reads at the 10th percentile to give a dimensionless metric that is indicative of the uniformity of the depth of read; the smaller the number, the more uniform (narrow) the depth of read. The average 90th percentile/10th percentile ratio is 11.5 for the method run in Example 9, while it is 5.6 for the method run in Example 7. A narrower depth of read for a given protocol plexity is better for sequencing efficiency, as fewer sequence reads are necessary to ensure that a certain percentage of reads are above a read number threshold.

Example 10

[0861] In one experiment, four maternal plasma samples were prepared and amplified using a semi-nested 9,600-plex protocol. Details of Example 10 were very similar to Example 9, the exception being the nesting protocol, and including the identity of the four samples. The ploidy calls for all 28 chromosomes in the set were called correctly with confidences above 99.7%. 7.6 million (97%) of reads mapped to the genome, and 6.3 million (80%) of the reads mapped to the targeted SNPs. The average depth of read was 751, and the median depth of read was 396.

Example 11

[0862] In one experiment, three maternal plasma samples were split into five equal portions, and each portion was

amplified using either 2,400 multiplexed primers (four portions) or 1,200 multiplexed primers (one portion) and amplified using a semi-nested protocol, for a total of 10,800 primers. After amplification, the portions were pooled together for sequencing. Details of Example 11 were very similar to Example 9, the exception being the nesting protocol, and the split and pool approach. The ploidy calls for all 21 chromosomes in the set were called correctly with confidences above 99.7%, except for one missed call where the confidence was 83%. 3.4 million reads mapped to targeted SNPs, the average depth of read was 404 and the median depth of read was 258.

Example 12

[0863] In one experiment, four maternal plasma samples were split into four equal portions, and each portion was amplified using 2,400 multiplexed primers and amplified using a semi-nested protocol, for a total of 9,600 primers. After amplification, the portions were pooled together for sequencing. Details of Example 12 were very similar to Example 9, the exception being the nesting protocol, and the split and pool approach. The ploidy calls for all 28 chromosomes in the set were called correctly with confidences above 97%, except for one missed call where the confidence was 78%. 4.5 million reads mapped to targeted SNPs, the average depth of read was 535 and the median depth of read was 412.

Example 13

[0864] In one experiment, four maternal plasma samples were prepared and amplified using a 9,600-plex triply hemi-nested protocol, for a total of 9,600 primers. Details of Example 12 were very similar to Example 9, the exception being the nesting protocol which involved three rounds of amplification; the three rounds involved 15, 10 and 15 STA cycles respectively. The ploidy calls for 27 of 28 chromosomes in the set were called correctly with confidences above 99.9%, except for one that was called correctly with 94.6%, and one missed call with a confidence of 80.8%. 3.5 million reads mapped to targeted SNPs, the average depth of read was 414 and the median depth of read was 249.

Example 14

[0865] In one Example 45 sets of cells were amplified using a 1,200-plex semi-nested protocol, sequenced, and ploidy determinations were made at three chromosomes. Note that this experiment is meant to simulate the conditions of performing pre-implantation genetic diagnosis on single-cell biopsies from day 3 embryos, or trophoctoderm biopsies from day 5 embryos. 15 individual single cells and 30 sets of three cells were placed in 45 individual reaction tubes for a total of 45 reactions where each reaction contained cells from only one cell line, but the different reactions contained cells from different cell lines. The cells were prepared into 5 ul washing buffer and lysed the by adding 5 ul ARCTURUS PICOPURE lysis buffer (APPLIED BIOSYSTEMS) and incubating at 56° C. for 20 min, 95° C. for 10 min..

[0866] The DNA of the single/three cells was amplified with 25 cycles of STA (95° C. for 10 min for initial polymerase activation, then 25 cycles of 95° C. for 30 s; 72° C. for 10 s; 65° C. for 1 min; 60° C. for 8 min; 65° C. for 3 min and 72° C. for 30 s; and a final extension at 72° C. for

2 min) using 50 nM primer concentration of 1200 target-specific forward and tagged reverse primers.

[0867] The semi-nested PCR protocol involved three parallel second amplification of a dilution of the first STAs product for 20 cycles of STA (95° C. for 10 min for initial polymerase activation, then 15 cycles of 95° C. for 30 s; 65° C. for 1 min; 60° C. for 5 min; 65° C. for 5 min and 72° C. for 30 s; and a final extension at 72° C. for 2 min) using reverse tag specific primer concentration of 1000 nM, and a concentration of 60 nM for each of 400 target-specific nested forward primers. In the three parallel 400-plex reactions the total of 1200 targets amplified in the first STA were thus amplified.

[0868] An aliquot of the STA products was then amplified by standard PCR for 15 cycles with 1 uM of tag-specific forward and barcoded reverse primers to generate barcoded sequencing libraries. An aliquot of each library was mixed with libraries of different barcodes and purified using a spin column.

[0869] In this way, 1,200 primers were used in the single cell reactions; the primers were designed to target SNPs found on chromosomes 1, 21 and X. The amplicons were then sequenced using an ILLUMINA GAIIIX sequencer. Per sample, approximately 3.9 million reads were generated by the sequencer, with 500,000 to 800,000 million reads mapping to the genome (74% to 94% of all reads per sample).

[0870] Relevant maternal and paternal genomic DNA samples from cell lines were analyzed using the same semi-nested 1200-plex assay pool with a similar protocol with fewer cycles and 1200-plex second STA, and sequenced.

[0871] The sequencing data was analyzed using informatics methods disclosed herein and the ploidy state was called at the three chromosomes for the samples.

[0872] FIG. 24 shows normalized depth of read ratios (vertical axis) for six samples at three chromosomes (1=chrom 1; 2=chrom 21; 3=chrom X). The ratios were set to be equal to the number of reads mapping to that chromosome, normalized, and divided by the number of reads mapping to that chromosome averaged over three wells each comprising three 46XY cells. The three sets of data points corresponding to the 46XY reactions are expected to have ratios of 1:1. The three sets of data points corresponding to the 47XX+21 cells are expected to have ratios of 1:1 for chromosome 1, 1.5:1 for chromosome 21, and 2:1 for chromosome X.

[0873] FIG. 25A-25C show allele ratios plotted for three chromosomes (1, 21, X) for three reaction. The reaction in the lower left shows a reaction on three 46XY cells (FIG. 25B). The left region are the allele ratios for chromosome 1, the middle region are the allele ratios for chromosome 21, and the right region are the allele ratios for chromosome X. For the 46XY cells, for chromosome 1 we expect to see ratios of 1, 0.5 and 0, corresponding to AA, AB and BB SNP genotypes. For the 46XY cells, for chromosome 21 we expect to see ratios of 1, 0.5 and 0, corresponding to AA, AB and BB SNP genotypes. For the 46XY cells, for chromosome X we expect to see ratios of 1 and 0, corresponding to A, and B SNP genotypes. The reaction in the lower right shows a reaction on three 47XX+21 cells (FIG. 25C). The allele ratios are segregated by chromosome as in the lower left graph. For the 47XX+21 cells, for chromosome 1 we expect to see ratios of 1, 0.5 and 0, corresponding to AA, AB and BB SNP genotypes. For the 47XX+21 cells, for chro-

mosome 21 we expect to see ratios of 1, 0.67, 0.33 and 0, corresponding to AAA, AAB, ABB and BBB SNP genotypes. For the 47XX+21 cells, for chromosome X we expect to see ratios of 1, 0.5 and 0, corresponding to AA, AB, and BB SNP genotypes. The plot in the upper right was made on a reaction comprising 1 ng of genomic DNA from the 47XX+21 cell line (FIG. 26A). FIGs. 26A and 26B shows the same graphs as in FIG. 25A-25C, but for reactions performed on only one cell. The left graph was a reaction that contained a 47XX+21 cell (FIG. 26A), and the right graph was for a reaction that contained a 46XY cell (FIG. 26B). **[0874]** From the graphs shown in FIGS. 25A-25C and FIGS. 26A and 26B, it is visually apparent that there are two clusters of dots for chromosomes where we expect to see ratios of 1 and 0; three clusters of dots for chromosomes where we expect to see ratios of 1, 0.5, and 0, and four clusters of dots for chromosomes where we expect to see ratios of 1, 0.67, 0.33 and 0. The parental support algorithm was able to make correct calls on all of the three chromosomes for all of the 45 reactions.

Example 15

[0875] In one experiment, maternal plasma samples were prepared and amplified using a hemi-nested 19,488-plex protocol. The samples were prepared in the following way: up to 20 mL of maternal blood were centrifuged to isolate the buffy coat and the plasma. The genomic DNA in the maternal sample was prepared from the buffy coat and paternal DNA was prepared from a blood sample or saliva sample. Cell-free DNA in the maternal plasma was isolated using the QIAGEN CIRCULATING NUCLEIC ACID kit and eluted in 50 uL TE buffer according to manufacturer's instructions. Universal ligation adapters were appended to the end of each molecule of 40 uL of purified plasma DNA and libraries were amplified for 9 cycles using adaptor specific primers. Libraries were purified with AGENCOURT AMPURE beads and eluted in 50 uL DNA suspension buffer.

[0876] 6 uL of the DNA was amplified with 15 cycles of STAR 1 (95° C. for 10 min for initial polymerase activation, then 15 cycles of 96° C. for 30 s; 65° C. for 1 min; 58° C. for 6 min; 60° C. for 8 min; 65° C. for 4 min and 72° C. for 30 s; and a final extension at 72° C. for 2 min) using 7.5 nM primer concentration of 19,488 target-specific tagged reverse primers and one library adaptor specific forward primer at 500 nM.

[0877] The hemi-nested PCR protocol involved a second amplification of a dilution of the STAR 1 product for 15 cycles (STAR 2) (95° C. for 10 min for initial polymerase activation, then 15 cycles of 95° C. for 30 s; 65° C. for 1 min; 60° C. for 5 min; 65° C. for 5 min and 72° C. for 30 s; and a final extension at 72° C. for 2 min) using reverse tag concentration of 1000 nM, and a concentration of 20 nM for each of 19,488 target-specific forward primers.

[0878] An aliquot of the STAR 2 products was then amplified by standard PCR for 12 cycles with 1 uM of tag-specific forward and barcoded reverse primers to generate barcoded sequencing libraries. An aliquot of each library was mixed with libraries of different barcodes and purified using a spin column.

[0879] In this way, 19,488 primers were used in the single-well reactions; the primers were designed to target SNPs found on chromosomes 1, 2, 13, 18, 21, X and Y. The amplicons were then sequenced using an ILLUMINA

GAIIIX sequencer. For plasma samples, approximately 10 million reads were generated by the sequencer, with 9.4-9.6 million reads mapping to the genome (94-96%), and of those, 99.95% mapped to targeted SNPs with a mean depth of read of 460 and a median depth of read of 350. For comparison, a perfectly even distribution would be: 10M reads/19,488 targets=513 reads/target. For primer-dimers, 30,000 reads were from sequenced primer-dimers (0.3% of the reads generated by the sequencer). For genomic samples, 99.4-99.7% of the reads mapped to the genome, of those, 99.99% of the mapped to targeted SNPs, and 0.1% of the reads generated by the sequencer were primer-dimers.

[0880] For plasma samples with 10 million sequencing reads, typically at least 19,350 of the 19,488 targeted SNPs (99.3%) are amplified and sequenced. For DNA samples with 2M sequencing reads, typically at least 19,000 targeted SNPs (97.5%) are amplified and sequenced. The lower number may be due to sampling noise since the number of reads is lower and the sequencer misses some of the amplified products. If desired, the number of sequencing reads can be increased to increase the number of targeted SNPs that are amplified and sequenced.

[0881] Relevant maternal and paternal genomic DNA samples amplified using a semi-nested 19,488 outer forward primers and tagged reverse primers at 7.5 nM in the STAR 1. Thermocycling conditions and composition of STAR 2, and the barcoding PCR were the same as for the hemi-nested protocol.

[0882] The average fetal fraction for 407 samples was found to be 14.8%. The sequencing data was analyzed using informatics methods disclosed herein and the ploidy state was called at four chromosomes (13, 18, 21, Y) for the fetuses whose DNA was present in 378 of the 407 maternal plasma samples, and at chromosome X in 375 of the 407 maternal plasma samples. The ploidy calls for all 1,887 chromosomes in the set were called correctly with confidences above 90%. 1882 of the 1887 calls were above 95%; and 1,862 of the 1,887 calls were called with confidences above 99%.

[0883] A similar control experiment was performed using water instead of DNA extracted from plasma in the plasma PCR protocol. Based on six such trials of an experiment, 5-6% of the sequenced reads were primer-dimers. Other sequenced reads were due to background noise. This experiment demonstrates that even in the absence of a nucleic acid sample with target loci for the primers to hybridize to (rather than hybridizing to other primers and forming amplified primer dimers) few primer dimers are formed.

Example 16

[0884] The following Example illustrates an exemplary method for designing and selecting a library of primers that can be used in any of the multiplexed PCR methods of the invention. The goal is to select primers from an initial library of candidate primers that can be used to simultaneously amplify a large number of target loci (or a subset of target loci) in a single reaction. For an initial set of candidate target loci, primers did not have to be designed or selected for each target locus. Preferably, primers are designed and selected for a large portion of the most desirable target loci.

Step 1

[0885] A set of candidate target loci (such as SNPs) were selected based on publically available information about

desired parameters for the target loci, such as frequency of the SNPs within a target population or heterozygosity rate of the SNPs (worldwide web at ncbi.nlm.nih.gov/projects/SNP/; Sherry S T, Ward M H, Kholodov M, et al. dbSNP: the NCBI database of genetic variation. *Nucleic Acids Res.* 2001 Jan. 1; 29(1):308-11, which are each incorporated by reference in its entirety). For each candidate locus, one or more PCR primer pairs were designed using the Primer3 program (the worldwide web at primer3.sourceforge.net; libprimer3 release 2.2.3, which is hereby incorporated by reference in its entirety). If there were no feasible designs for PCR primers for a particular target locus, then that target locus was eliminated from further consideration. If desired, a "target locus score" (higher score representing higher desirability) can be calculated for most or all of the target loci, such as a target locus score calculated based on a weighted average of various desired parameters for the target loci. The parameters may be assigned different weights based on their importance for the particular application that the primers will be used for. Exemplary parameters include the heterozygosity rate of the target locus, the disease prevalence associated with a sequence (e.g., a polymorphism) at the target locus, the disease penetrance associated with a sequence (e.g., a polymorphism) at the target locus, the specificity of the candidate primer(s) used to amplify the target locus, the size of the candidate primer(s) used to amplify the target locus, and the size of the target amplicon. In some embodiments, the specificity of the candidate primer for the target locus includes the likelihood that the candidate primer will mis-prime by binding and amplifying a locus other than the target locus it was designed to amplify. In some embodiments, one or more or all the candidate primers that mis-prime are removed from the library.

Step 2

[0886] A thermodynamic interaction score was calculated between each primer and all primers for all other target loci from Step 1 (see, e.g., Allawi, H. T. & SantaLucia, J., Jr. (1998), "Thermodynamics of Internal C-T Mismatches in DNA", *Nucleic Acids Res.* 26, 2694-2701; Peyret, N., Senéviratne, P. A., Allawi, H. T. & SantaLucia, J., Jr. (1999), "Nearest-Neighbor Thermodynamics and NMR of DNA Sequences with Internal A-A, C-C, G-G, and T-T Mismatches", *Biochemistry* 38, 3468-3477; Allawi, H. T. & SantaLucia, J., Jr. (1998), "Nearest-Neighbor Thermodynamics of Internal A-C Mismatches in DNA: Sequence Dependence and pH Effects", *Biochemistry* 37, 9435-9444.; Allawi, H. T. & SantaLucia, J., Jr. (1998), "Nearest Neighbor Thermodynamic Parameters for Internal G-A Mismatches in DNA", *Biochemistry* 37, 2170-2179; and Allawi, H. T. & SantaLucia, J., Jr. (1997), "Thermodynamics and NMR of Internal G-T Mismatches in DNA", *Biochemistry* 36, 10581-10594; MultiPLX 2.1 (Kaplinski L, Andreson R, Puurand T, Remm M. MultiPLX: automatic grouping and evaluation of PCR primers. *Bioinformatics.* 2005 Apr. 15; 21(8):1701-2, which are each hereby incorporated by reference in its entirety). This step resulted in a 2D matrix of interaction scores. The interaction score predicted the likelihood of primer-dimers involving the two interacting primers. The score was calculated as follows:

$$\text{interaction_score} = \max(-\Delta G_{2,0.8}^* - \Delta G_{1,1})$$

[0887] where

[0888] ΔG_2 =Gibbs energy (energy required to break the dimer) for a dimer that is extensible by PCR on both ends, i.e., the 3' end of each primer anneals to the other primer; and

[0889] ΔG_1 =Gibbs energy for a dimer that is extensible by PCR on at least one end.

Step 3:

[0890] For each target locus, if there was more than one primer-pair design, then one design was selected using the following method:

[0891] For each primer-pair design for the locus, find the worst-case (highest) interaction score for the two primers in that design and all primers from all designs for all other target loci.

[0892] Pick the design with the best (lowest) worst-case interaction score.

Step 4

[0893] A graph was built such that each node represented one locus and its associated primer-pair design (e.g., a Maximal Clique problem). One edge was created between every pair of nodes. A weight was assigned to each edge equal to the worst-case (highest) interaction score between the primers associated with the two nodes connected by the edge.

Step 5

[0894] If desired, for every pair of designs for two different target loci where one of the primers from one design and one of the primers from the other design would anneal to overlapping target regions, an additional edge was added between the nodes for the two design. The weight of these edges was set equal to the highest weight assigned in Step 4. Thus, Step 5 prevents the library from having primers that would anneal to overlapping target regions, and thus interfere with each other during a multiplex PCR reaction.

Step 6

[0895] An initial interaction score threshold was calculated as follows:

$$\text{weight_threshold} = \max(\text{edge_weight}) - 0.05 * (\max(\text{edge_weight}) - \min(\text{edge_weight}))$$

[0896] where

[0897] $\max(\text{edge_weight})$ is the maximum edge weight in the graph; and

[0898] $\min(\text{edge_weight})$ is the minimum edge weight in the graph.

[0899] The initial bounds for the threshold were set as follows:

$$\max_weight_threshold = \max(\text{edge_weight})$$

$$\min_weight_threshold = \min(\text{edge_weight})$$

Step 7

[0900] A new graph was constructed consisting of the same set of nodes as the graph from Step 5, only including edges with weights that exceed weight_threshold . Thus, step ignores interactions with scores equal to or below weight_threshold .

Step 8

[0901] Nodes (and all of the edges connected to the removed nodes) were removed from the graph of Step 7 until there were no edges left. Nodes were removed by applying the following procedure repeatedly:

[0902] Find the node with the highest degree (highest number of edges). If there is more than one then pick one arbitrarily.

[0903] Define the set of nodes consisting of the node picked above and all of the nodes connected to it, but excluding any nodes that have degree less than the node picked above.

[0904] Choose the node from the set that has the lowest target locus score (lower score representing lower desirability) from Step 1. Remove that node from the graph.

Step 9

[0905] If the number of nodes remaining in the graph satisfies the required number of target loci for the multiplexed PCR pool (within an acceptable tolerance), then the method was continued at Step 10.

[0906] If there were too many or too few nodes remaining in the graph, then a binary search was performed to determine what threshold values would result in the desired number of nodes remaining in the graphs. If there were too many nodes in the graph then, the weight threshold bounds were adjusted as follows:

$$\max_weight_threshold = \text{weight_threshold}$$

[0907] Otherwise (if there are too few nodes in the graph), then the weight threshold bounds were adjusted as follows:

$$\min_weight_threshold = \text{weight_threshold}$$

[0908] Then, the weight threshold was adjusted follows:

$$\text{weight_threshold} = (\max_weight_threshold + \min_weight_threshold) / 2$$

[0909] Steps 7-9 were repeated.

Step 10

[0910] The primer-pair designs associated with the nodes remaining in the graph were selected for the library of primers. This primer library can be used in any of the methods of the invention.

[0911] If desired, this method of designing and selecting primers can be performed for primer libraries in which only one primer (instead of a primer pair) is used for amplification of a target locus. In this case, a node presents one primer per target locus (rather than a primer pair).

Example 17

[0912] FIG. 27 is a graph comparing two primer libraries designed using the methods of the invention. This graph shows the number of loci with a particular minor allele frequency that are targeted by each primer library. During the selection of the "new pool" library, more primers were retained. This library enables the amplification of more target loci, especially target loci with relatively large minor allele frequencies (which are the more informative alleles for some method of the invention, such as for detecting fetal chromosomal abnormalities).

[0913] These primer libraries were used in the following multiplex PCR method. Blood (20-40 mL) was collected

from each subject into two to four CELL-FREE™ DNA tubes (Streck). Plasma (a minimum of 7 mL) was isolated from each sample via a double centrifugation protocol of 2,000 g for 20 min, followed by 3,220 g for 30 min, with supernatant transfer following the first spin. cfDNA was isolated from 7-20 mL plasma using the QIAGEN QIAamp Circulating Nucleic Acid kit and eluted in 45 uL TE buffer. Pure maternal genomic DNA was isolated from the buffy coat obtained following the first centrifugation, and pure paternal genomic DNA was prepared similarly from a blood, saliva or buccal sample.

[0914] Maternal cfDNA, maternal genomic DNA, and paternal genomic DNA samples were pre-amplified for 15 cycles using 11,000 target-specific assays and an aliquot was transferred to a second PCR reaction of 15 cycles using nested primers. Finally, samples were prepared for sequencing by adding barcoded tags in a third 12-cycle round of PCR. Thus, 11,000 targets were amplified in a single reaction; the targets included SNPs found on chromosomes 13, 18, 21, X, and Y. The amplicons were then sequenced using an ILLUMINA GAIIX or HISEQ sequencer. Parental genotypes were sequenced at a lower read depth (~20% of cfDNA read depth) than the fetal genotypes.

Example 18

[0915] If desired, the size and quantity of the PCR products can be analyzed using standard methods, such as the use of the Agilent Technologies 2100 Bioanalyzer (FIG. 28A-M). For example, direct PCR methods described herein without nesting were used in 2,400-plex (FIGS. 28B-28G) and 19,488-plex experiments (FIGS. 28H to 28M). The amount of primer was 10 nM for FIGS. 28B-28D and 28H to 28J. The amount of primer was 1 nM for FIGS. 28E-28G and 28K to 28M. The amount of input DNA was 24 ng for FIGS. 28B, 28E, 28H, and 28K; 80 ng for FIGS. 28C, 28F, 28I, and 28L; and 250 ng for FIGS. 28D, 28G, 28J, and 28M. More input DNA resulted in a greater proportion of the desired 180 base pair product. The peak at 140 base pairs is a primer dimer product.

Example 19

[0916] A proof-of-principle study demonstrated the detection of T13, T18, T21, 45,X, and 47,XXY with equally high accuracies across all chromosomes.

Patients

[0917] Pregnant couples were enrolled at specific prenatal care centers under protocols approved by an Institutional Review Board pursuant to local laws. Inclusion criteria were at least 18 years of age, gestational age of at least nine weeks, singleton pregnancies, and signed informed consent. Blood samples were drawn from pregnant mothers, and a blood or buccal sample was collected from the father. Samples from 2 pregnancies with T13 (Patau Syndrome), 2 with T18 (Edwards Syndrome), 2 with T21 (Down's Syndrome), 2 with 45,X, 2 with 47,XXY, and 90 normal pregnancies were selected prior to testing from a cohort of ~500 women to test which chromosomal abnormalities the method detects. Normal fetal karyotype was confirmed by molecular karyotyping for the samples where post-birth child tissue was available. Euploid sample were drawn prior to invasive testing from low-risk women. Aneuploid samples were drawn at least 7 days after invasive testing and

aneuploidy was confirmed via cytogenetic karyotyping or fluorescence in situ hybridization at independent laboratories.

Sample Preparation and Multiplex PCR

[0918] For the data in FIGS. 30, 30D, 30E, 30G, and 30H, and 31A-31G, sample preparation and 19,488-plex-PCR were performed as described in Example 15. For the data in FIG. 30F, sample preparation and 11,000-plex-PCR were performed as described in Example 17.

Methodology and Data Analysis

[0919] The algorithm considers parental genotypes and crossover frequency data (such as data from the HapMap database) to calculate expected allele distributions for 19,488 polymorphic loci for a very large number possible fetal ploidy states, and at various fetal cfDNA fractions. (FIG. 29). Unlike allele ratio based-methods, it also takes into account linkage disequilibrium, and uses non-Gaussian data models to describe the expected distribution of allele measurements at a SNP given observed platform characteristics and amplification biases. It then compares the various predicted allele distributions to the actual allelic distributions as measured in the cfDNA sample (FIG. 29 step C), and calculates the likelihood of each hypothesis (monosomy, disomy, or trisomy, for which there are numerous hypotheses based on the various potential crossovers) based on the sequencing data. The algorithm sums the likelihoods of each individual monosomy, disomy, or trisomy hypotheses (FIG. 29 step D), and calls the ploidy state with the maximum overall likelihood as the copy number and fetal fraction (FIG. 29 step E). Although laboratory researchers were not blinded to sample karyotype, the algorithm called the ploidy states without human intervention and was blind to the truth.

Data Interpretation

Graphical Representations of the Generated Data

[0920] To determine the ploidy state of a chromosome of interest, the algorithm considers the distribution of sequence counts from each of two possible alleles at 3,000 to 4,000 SNPs per chromosome. It is important to note that the algorithm makes ploidy calls using an approach that does not lend itself to visualization. Thus, for the purposes of illustration, the data is displayed here in a simplified fashion as ratios of the two most likely alleles, labeled as A and B, so that the relevant trends can be more readily visualized. This simplified illustration does not take into account some of the features of the algorithm. For example, two important aspects of the algorithm that are not possible to illustrate with a method of visualization that displays allele ratios are: 1) the ability to leverage linkage disequilibrium, i.e. the influence that a measurement at one SNP has on the likely identity of a neighboring SNP, and 2) the use of non-Gaussian data models that describe the expected distribution of allele measurements at a SNP given platform characteristics and amplification biases. Also note that the algorithm only considers the two most common alleles at each SNP, ignoring other possible alleles.

[0921] The graphical representations in FIGS. 30, 30D-30H include samples for which two, one, or three fetal chromosomes are present. Generally, these indicate euploidy (FIG. 30) monosomy (FIG. 30D), and trisomy (FIGS. 30E-

30H), respectively. In all plots, each spot represents a single SNP, where the targeted SNPs are plotted sequentially from left to right for one chromosome along the horizontal axes. The vertical axes indicate the number of reads for the A allele as a fraction of the total number of reads for both the A and B alleles for that SNP. Note that the measurements are made on total cfDNA isolated from maternal blood, and the cfDNA includes both maternal and fetal cfDNA; thus, each spot represents the combination of the fetal and maternal DNA contribution for that SNP. Therefore, increasing the proportion of maternal cfDNA from 0% to 100% will gradually shift some spots up or down within the plots, depending on the maternal and fetal genotype. This is described in more detail below with the corresponding plots.

[0922] If desired to facilitate visualization, the spots may be color-coded according to maternal genotype, as maternal genotype contributes more to the localization of each spot and the majority of trisomies are maternally-inherited; this assists in visualizing ploidy states. Specifically, SNPs for which the maternal genotype is AA may be indicated in red, those for which the maternal genotype is AB may be indicated in green, and those for which the maternal genotype is BB may be indicated in blue.

[0923] In all cases, SNPs that are homozygous for the A allele (AA) in both the mother and the fetus are found tightly associated with the upper limit of the plots, as the fraction of A allele reads is high because there should be no B alleles present. Conversely, SNPs that are homozygous for the B allele in both the mother and the fetus are found tightly associated with the lower limit of the plots, as the fraction of A allele reads is low because there should be only B alleles. The spots that are not tightly associated with the upper and lower limits of the plots represent SNPs for which the mother, the fetus, or both are heterozygous; these spots are useful for identifying fetal ploidy, but can also be informative for determining paternal versus maternal inheritance. These spots segregate based on both maternal and fetal genotypes and fetal fraction, and as such the precise position of each individual spot along the y-axis depends on both stoichiometry and fetal fraction. For example, loci where the mother is AA and the fetus is AB are expected to have a different fraction of A allele reads, and thus different positioning along the y-axis, depending on the fetal fraction.

Two Chromosomes Present

[0924] FIG. 30 depict data that indicate the presence of two chromosomes when the sample is entirely maternal (no fetal cfDNA present, FIG. 30 (0% FF plot)), contains a moderate fetal cfDNA fraction (FIG. 30 (12% FF plot)), or contains a high fetal cfDNA fraction (FIG. 30 (26% FF plot)).

[0925] FIG. 30, 0% FF plot, shows data obtained from cfDNA isolated from the blood of a non-pregnant woman. When there is no fetal cfDNA present and the sample contains only maternal cfDNA, the plots represent purely the euploid maternal genotype; the hallmark pattern includes “clusters” of spots: a filled circle cluster tightly associated with the top of the plot (SNPs where the maternal genotype is AA), a filled square cluster tightly associated with the bottom of the plot (SNPs where the maternal genotype is BB), and a single, centered open triangle cluster (SNPs where the maternal genotype is AB).

[0926] When fetal cfDNA is present, the location of the spots shifts such that the clusters segregate into discrete

“bands”. Note that for samples with a fetal fraction of 0%, the groupings of spots are referred to as “clusters” (as in FIG. 30, 0% FF plot), and for all samples with a fetal fraction of >0%, the groupings of spots are referred to as “bands” (as in FIG. 30 (12% FF plot), 30 (26% FF plot), 30D-30J). If the fetal fraction is high enough, these discrete bands will be readily visible. Specifically, FIG. 30 12% and 26% FF plots demonstrate the characteristic pattern associated with two fetal chromosomes present at moderate and high fetal fractions, respectively. This pattern includes three central open triangle bands that correspond to SNPs that are heterozygous in the mother, and two “peripheral” bands each at both the top (filled circles) and bottom (filled square) of the plots that correspond to SNPs that are homozygous in the mother.

[0927] FIG. 30, 12% FF plot, shows data obtained from cfDNA isolated from a plasma sample from a woman carrying a euploid fetus and with a 12% fetal cfDNA fraction. Here, the clusters of spots tightly associated with the top and bottom of the plot segregate into two discrete bands each: one filled circles and one filled square external peripheral band that remains tightly associated with the upper or lower limit of the plots, and one filled circle and one filled square internal peripheral band that has separated from the limits of the plots. These internal peripheral bands, centered around 0.92 and 0.08, represent SNPs for which the maternal genotype is AA and the fetal genotype is AB (indicated in filled circles), and SNPs for which the maternal genotype is BB and the fetal genotype is AB (indicated in filled square), respectively. The center cluster of open triangle spots broadens, but at this fetal fraction the segregation into distinct bands is not readily visible.

[0928] At a high fetal cfDNA fraction, the typical pattern that indicates the presence of two chromosomes (a trio of open triangle bands as well as two filled circles and two filled square peripheral bands) is readily apparent. FIG. 30, 26% FF plot, displays data obtained from a plasma sample from a woman carrying a euploid fetus at a fetal cfDNA fraction of 26%. Here, the peripheral bands have separated such that the internal band has shifted towards the center of the plot due to the altered levels of B alleles from the increased fetal cfDNA fraction. Significantly, at higher fetal fractions, the separation of the center open triangle cluster into three distinct bands is now readily apparent. This central trio of bands, in this case clustering around 0.37, 0.50 and 0.63, corresponds to those SNPs where the maternal genotype is AB, and the fetal genotype is AA (top), AB (middle) and BB (bottom).

[0929] These hallmark patterns, namely three open triangle bands and four peripheral bands (two filled circles and two filled square), indicate the presence of two chromosomes, as in autosomal euploidy or for the X chromosome in a female (XX) fetus.

One Chromosome Present

[0930] When the fetus only inherits a single chromosome, and thus only inherits a single allele, heterozygosity of the fetus is not possible. As such, the only possible fetal SNP identities are A or B. Thus, maternally-inherited monosomic chromosomes have a characteristic pattern of two central open triangle bands that represent SNPs for which the mother is heterozygous, and only have single peripheral filled circles and filled square bands that represent SNPs for which the mother is homozygous, and which remain tightly

associated with the upper and lower limits of the plots (1 and 0), respectively (FIG. 30D). Note the absence of internal peripheral bands. This pattern indicates the presence of one chromosome, as in maternally-inherited autosomal monosomy, or for the X chromosome in a male (XY) fetus.

Three Chromosomes Present

[0931] Trisomic chromosomes have three characteristic patterns. The first pattern indicates maternally-inherited meiotic trisomy, a meiotic error where the fetus inherited two homologous, non-identical chromosomes from the mother (FIG. 30E); this pattern includes two central open triangle bands with two each of the peripheral filled circles and filled square bands. The second pattern indicates paternally-inherited meiotic trisomy, where the fetus inherited two homologous, non-identical chromosomes from the father (FIG. 30F); this pattern includes four central open triangle bands and three each of the peripheral filled circles and filled square bands. The third pattern indicates either maternally—(FIG. 30G) or paternally-inherited (FIG. 30H) mitotic trisomy, a mitotic error where the fetus inherited two identical chromosomes from either the mother or the father; this pattern includes four central open triangle bands with two each of the peripheral filled circles and filled square bands. Maternally- and paternally-inherited mitotic trisomies can be distinguished by the placement of the flanking filled circles and filled square bands, such that the filled circles and filled square internal peripheral bands (those not associated with the limits of the plots) are closer to the center in paternally-inherited mitotic trisomy. This is due to the paternal contribution of identical chromosomes. Note that our previous results indicate that at the blastomere stage, 66.7% of maternally-inherited trisomies are meiotic, and that only 10.2% of trisomies are paternally-inherited.

[0932] For the Y chromosome, the PS method considers a different set of hypotheses: zero, one, or two chromosomes present. As there is no maternal contribution to the sequence reads at each locus and because heterozygous loci are not possible (cases of two Y chromosomes necessarily involve two identical chromosomes), the bands remain tightly associated with the top (A alleles) or the bottom (B alleles) of the plot (data not shown), and analysis is greatly simplified, relying on quantitative allele count data. Note that since the method interrogates SNPs, it uses homologous non-recombinant SNPs from the Y chromosome, thus obtaining data on both X and Y for one probe pair.

Identifying Aneuploidies

[0933] Identification of autosomal aneuploidies using this plot-based visualization method is straightforward given a sufficient fetal fraction, and requires only identifying plots for which there are an abnormal number of chromosomes present, as described above. Combining the knowledge of copy number of the X and Y chromosomes identifies whether sex chromosome aneuploidies are present. Specifically, plots representing a fetus with a 47,XXX genotype will have a typical “three-chromosome” pattern, and plots representing a fetus with a 47,XXY genotype will have the typical “two-chromosome” pattern for the X chromosome, but will also have allele reads indicating the presence of one Y chromosome. The method is similarly able to call 47,YYY, where a “one chromosome” pattern indicates the presence of a single X chromosome, and allele reads indicate

the presence of two Y chromosomes. A fetus with a 45,X genotype will have the typical “one-chromosome” pattern for the X chromosome, and data indicating zero Y chromosomes.

Effects of Fetal Fraction

[0934] As discussed above, the number of sequence reads from the fetus contributes to the precise location of each spot along the y-axis in the plots. As fetal fraction will affect the proportion of reads that originate from the fetus and the mother, it will also affect the positioning of each spot. At a high fraction of fetal cfDNA (generally above ~20%), as in FIG. 30 (26% FF plot), 30D, 30E, 30G, and 30H, it is readily apparent that although the spots cluster based mainly on maternal genotype, the presence of fetal DNA from alleles whose genotype is distinct from the maternal genotype shift the clusters into multiple, distinct bands. However, as the fetal fraction decreases (as in FIGS. 30 (12% FF plot) and 30F), the spots regress towards the poles and center of the plot, resulting in tighter clusters. Specifically, the set of peripheral filled circles bands, where the maternal genotype is AA, regress towards the top of the plot; the set of peripheral filled square bands, where the maternal genotype is BB, regress towards the bottom; the set of central open triangle bands, where the mother is heterozygous, condense into a single cluster at the center of the plot (compare FIG. 30, 12% and 26% FF plots). Although aneuploidy is not readily apparent by eye using this visualization technique for low fetal fraction cases, the algorithm is able to identify ploidy states with a very low fetal fraction, such as 3% fetal fraction. It is able to do this because the statistical technique compares the observed data to very precise data models that predict the allele distributions for a given sample parameter set (including copy number, parental genotypes, and fetal fraction, for example). Data model precision is critical in low fetal fraction cases, as the differences between the allele distributions for different ploidy states are proportional to the fetal fraction. In addition, the algorithm is able to determine when a data set does not contain enough data to make a confident fetal ploidy determination.

Results

[0935] Sequencing reads that mapped to targeted SNPs were deemed to be informative and were used by the algorithm. More than 95% of targeted loci were observed in the sequencing results. The plots for visualizing key ploidy calls are depicted in FIG. 31A-31G. FIG. 31A indicates a euploid sample. Here, chromosomes 13, 18, and 21 have the typical “two chromosome” pattern (as described herein). This includes a trio of center open triangle bands, and two filled circles and two filled square peripheral bands. This, together with the two center open triangle bands for the X chromosome and the presence of Y chromosome bands along the plots’ peripheries, indicate a euploid XY genotype.

[0936] The most prevalent autosomal trisomies, T13, T18, and T21, are indicated by the plots in FIGS. 31B, 31C, and 31D, respectively. Specifically, FIG. 31B depicts a T13 sample. Here, chromosomes 18 and 21 display the typical “two chromosome” pattern, chromosome X displays the typical “one chromosome” pattern, and there are reads from the Y chromosome. Together, this indicates disomy at chromosomes 18 and 21, and identifies a fetal XY genotype. However, chromosome 13 depicts a typical “three chromo-

some” pattern—specifically. Similarly, FIG. 31C depicts a T18 sample, and FIG. 31D depicts a T21 sample.

[0937] The method is also able to detect sex chromosome aneuploidies, including 45,X (FIG. 31E), 47,XXY (FIG. 31F), and 47,XXX (FIG. 31G). Note that the method is calling copy number at chromosomes 13, 18, 21, X, and Y; the overall chromosome number is reported assuming disomy at the remaining chromosomes. The X chromosome regions of the plot depicting a 45,X sample reveals the presence of a single chromosome. However, the lack of reads from the Y chromosome, coupled with the “two chromosome” pattern for chromosomes 13, 18, and 21, indicate a 45,X genotype. Conversely, the 47,X×Y samples generate a plot revealing the presence of two X chromosomes. The data also revealed reads for alleles from the Y chromosome. Together with the presence of two copies of chromosomes 13, 18, and 21, this indicates a 47,XXY genotype. A 47,XXX genotype is indicated by the presence of a “one chromosome” pattern for the X chromosome, and reads indicating the presence of two Y chromosomes.

Discussion

[0938] This method detected T13, T18, T21, 45,X, 47,XXY, and 47,XXX non-invasively from maternal blood. This method interrogates cfDNA from maternal plasma by targeted multiplex PCR amplification and high-throughput sequencing of 19,488 SNPs. This, coupled with the method’s sophisticated informatics analyses that take into account parental genotypic information and numerous sample parameters, including fetal fraction and DNA quality, more robustly detects the fetal signal and makes highly accurate ploidy calls at all of the five chromosomes implicated in the seven most common types of at-birth aneuploidy (T13, T18, T21, 45,X, 47,XXX, 47,XXY, and 47,XXX). This method offers a number of clinical advantages over previous methods, including and most significantly greater clinical coverage and sample-specific calculated accuracies (analogous to a personalized risk score).

Increased Clinical Coverage

[0939] This method offers approximately a two-fold increase in aneuploidy coverage compared to clinically available NIPT methodologies, given its ability to accurately detect autosomal trisomies and sex chromosome aneuploidies. The method presented here is the only noninvasive test that calls ploidy at the sex chromosomes with high accuracy. Prior DNA mixing experiments and separate plasma samples analyzed in our experimental assays suggest that this method will detect a larger cohort of sex chromosome anomalies, including 47,XXX. The method presented here also detects aneuploidies at chromosomes 13, 18, and 21 with high sensitivities and specificities, and with appropriate primer design is expected to be able to detect copy number at the remaining chromosomes as well.

Sample-Specific Calculated Accuracies

[0940] Significantly, this method calculates a sample-specific accuracy for ploidy calls on each chromosome in each sample. Accuracies calculated by this method are expected to significantly lower the rate of incorrect calls by identifying and flagging individual samples that have poor quality DNA or low fetal fractions that are likely to result in a poor accuracy test result. By contrast, massively parallel shotgun

sequencing (MPSS)-based methods produce a positive or negative call using a single-hypothesis rejection test, and their accuracy estimate is based on a published study cohort rather than on the characteristics of the individual sample, which are assumed to have the same accuracy as the cohort. However, individual accuracies for samples with parameters in the tail of the cohort distribution may differ significantly. This is exacerbated at low fetal fractions, as in early gestational age, or for samples with low DNA quality. These samples are generally not identified and flagged for follow-up, which can result in missed calls. The present method, however, takes into account many parameters, including fetal fraction and a number of DNA quality metrics, to make each chromosome copy number call, calculating a sample-specific accuracy for that call. This allows the method to identify individual samples with low accuracy and flag them for follow-up. This is expected to nearly eliminate missed calls, especially at the early stages of pregnancy when fetal fractions are typically low. The presumption is that a no call is much preferred to a missed call, since a no call simply requires a redraw and reanalysis.

Converting Calculated Accuracies to Traditional Risk Scores

[0941] This method can offer an adjusted risk of aneuploidy for high-risk pregnant women, where the adjusted risk takes into account an a priori risk (Benn P, Cuckle H, Pergament E. Non-invasive prenatal diagnosis for Down syndrome: the paradigm will shift, but slowly. *Ultrasound Obstet Gynecol* 2012; 39:127-130, which is hereby incorporated by reference in its entirety). Although the present method offers each patient a customized calculated accuracy, for clinical use these accuracies can be converted to traditional risk scores, which also denote the risk of an aneuploid pregnancy but are expressed as fractions. Traditional risk scores take into account various parameters, including maternal age-related risk and serum levels of biochemical markers, to offer a risk score above which a mother is considered high-risk and for whom follow-up invasive diagnostic procedures are recommended. This method significantly refines this risk score, thus reducing both the false positive and false negative rates, and offering a more accurate assessment of individual maternal risk. A calculated accuracy as used here is the likelihood that the ploidy call is correct, and is expressed as a percentage, but the calculated accuracies used in Experiment 19 do not include an age-related risk. Because calculation of a risk score typically includes an age-related risk, the calculated accuracies and traditional risk scores are not interchangeable; they must be combined to convert into a traditional risk score. The formula to combine the age-related risk with the calculated accuracy is:

$$\frac{R_1 R_2}{R_1 R_2 + [1 - R_1][1 - R_2]}$$

where R_1 is the risk score as calculated by the present method and R_2 is the risk score as calculated by first trimester screening.

SNP-Based Methods Negate Issues with Amplification Variation

[0942] An inherent drawback to the counting methods used by some other methods is that they determine fetal ploidy state by measuring the ratio of the number of reads mapping to the chromosome of interest (e.g., chromosome 21) to those mapping to a reference chromosome. Chromosomes with high or low GC content, including chromosomes 13, X, and Y, amplify with high variability. This can result in signal variation that is comparable in magnitude to the fetal cfDNA signal, which can confound copy number calls by altering the ratio of allele reads from the chromosome-of-interest to those from the reference chromosome. This can result in low accuracy for chromosomes 13, X, and Y. Significantly, this problem is exacerbated at low fetal cfDNA fractions, as tends to be the case at early gestational ages.

[0943] In contrast, SNP-based methods do not rely on consistent amplification levels between chromosomes, and are thus expected to provide results that are equally accurate across all chromosomes. Because the present method looks, in part, at relative counts of different alleles at polymorphic loci, which by definition differ only by a single nucleotide, it does not require the use of reference chromosomes, and this obviates the problems with chromosome-to-chromosome amplification variation that are inherent to methods that rely on quantitating read counts. Unlike quantitative methods that require reference chromosomes that are euploid, the present method is expected to be able to detect triploidy as well as copy-number neutral anomalies like uniparental disomy.

The Importance of Early Detection

[0944] Significantly, the combined at-birth prevalence of sex chromosome aneuploidies is higher than that of the most common autosomal aneuploidies (FIG. 32). However, there are currently no routine non-invasive screening methods that reliably detect sex chromosome abnormalities. Thus, sex chromosome anomalies are generally detected prenatally as a side-effect of routine testing for Down syndrome or other autosomal aneuploidies; a large proportion of cases are missed entirely. Early and accurate detection is crucial for many of these disorders where early therapeutic intervention improves clinical outcomes. For example, Turner syndrome is often not diagnosed until adolescence, although its overall at-birth prevalence is 1 in 2,500 females. Growth hormone therapy is known to prevent short stature that results from the disorder, but treatments are significantly more effective when initiated prior to the age of 4. Additionally, estrogen replacement therapy can stimulate secondary sexual characteristics in patients with Turner syndrome, but again therapy must be initiated in the pre-teen years, before the syndrome is usually detected. Together, this underscores the importance of early, routine, and safe detection of sex chromosome aneuploidies. This method offers the first approach with the potential to serve as a routine screen for sex chromosome anomalies.

Example 20

[0945] The following experiment illustrates an exemplary method for designing and selecting primers that can be used in any of the multiplexed PCR methods of the invention. In some embodiments, primers from an initial library of can-

didate primers are selected so that they can be used to simultaneously amplify a large number of target loci (or a subset of target loci) in a single reaction. In some embodiments, primers from an initial library of candidate primers are selected to form multiple primer pools such that each pool can be used to simultaneously amplify a subset of target loci in a single reaction. Preferably, primers are designed and selected for a large portion or all of the most desirable target loci. Preferably, the minimum number of pools needed to amplify the target loci are created.

Step 1

[0946] Calculate a first score for each primer pair design using one or more of the following parameters: number of SNPs within the primers, location of SNPs within the primers, distance from an end of the amplicon to the target bases within the amplicon, number of target loci in an amplicon, heterozygosity rate of the target locus, disease prevalence associated with a sequence (e.g., a polymorphism) at the target locus, disease penetrance associated with a sequence (e.g., a polymorphism) at the target locus, specificity of the candidate primer for the target locus, size of the candidate primer, melting temperature of the candidate primer, melting temperature of the target amplicon, GC content of the target amplicon, GC content of the 3' end of the candidate primer, homopolymer length in the candidate primer, amplification efficiency of the target amplicon, and size of the target amplicon.

Step 2

[0947] Compare each primer pair to every other primer pair, and calculate a second score for the pair using one or more of the following parameters: likelihood of dimer formation, amplicon overlap, number of primer designs for a particular target locus, and distance between amplicons. In some embodiments, the score is infinite if amplicons overlap so that two different primer pairs that generate overlapping amplicons are not included in the same primer pool.

Step 3

[0948] Aggregate the first score and the second score together (such as by using a weighted average of the scores).

Step 4

[0949] If desired, order all target loci into one contiguous list based upon their genomic location in ascending order.

Step 5

[0950] Build a minimum priority queue data structure that prioritizes the pairs of designs (in which each design is one primer pair so that a pair of designs includes two primer pairs with a total of 4 primers) based on their score (such as the aggregate score from step 3). In some embodiments, the score for a pair of designs is the worse score (such as the worse aggregate score from step 3) out of the scores for all 4 primers in the pair of designs. The pair of designs with the best (most desirable) score is first in the queue, and the pair of designs with the worst (least desirable) score is last in the queue. If desired, pairs of designs with a score above a threshold (least desirable) are removed from the library of candidate primers such that they are not included in the final pool(s) (for example, these primers may be omitted from the

queue). In some embodiments, pairs of design with an interaction score above (worse than) 20, 18, 16, 14, 12, 10, 9, 8, 7, 6, 5, 4, 3, 2, or 1 kcal/mol are removed from the library of candidate primers. In some embodiments, pairs of design with a ΔG value below (worse than) -20, -18, -16, -14, -12, -10, -9, -8, -7, -6, -5, -4, -3, -2, or -1 kcal/mol are removed from the library of candidate primers.

[0951] Each design pair can be stored as a node of a doubly linked list with initial “next” and “previous” pointers set to NULL.

Step 6

[0952] Begin forming all pools simultaneously by doing the following steps. Take the design pair with the best (most desirable) score from the priority queue and add it to “the potential pools.” Begin storing designs in N number of doubly linked list data structures with the design pairs. N represents the current number of different primer pools. Initially, N=1, since there is only one primer pool. In some embodiments, a second pool is only created if necessary to include the desired target loci or the desired level of coverage of target loci. Check to see if the design pair removed from the queue is “connected” to any other existing design pair. By “connected” for purposes of this step is meant that a single design in one pair is the same as a single design in another pair. If two pairs are connected, then assign the appropriate next and previous pointers to one another. If two pairs are not connected, then add them to the “potential pools” In some embodiments, a design pair is only placed in a particular pool if it would be connected to at most two other design pairs in that pool (otherwise it can be assigned to a different pool).

[0953] Check to see if (i) any linked list spans from the first target to the last target (such that all the desired target loci are included) or (ii) if a pool meets the cutoff for the desired minimum pool level. If it does, that list now forms a pool and can be added to the “final pools” list.

Step 7

[0954] If desired, check to see if the desired level of coverage (such as all the bases in the target loci being included in amplicons from 4 different primer pairs) that is desired for each location. Repeat step 6 until achieving the desired level of coverage.

[0955] The resulting primer pool(s) can be used in any of the methods of the invention.

Example 21

[0956] The following Example illustrates an exemplary method for designing and selecting primers that can be used in any of the multiplexed PCR methods of the invention. In some embodiments, the primers are divided into different pools (e.g., 2, 3, 4, 5, 6, or more different pools) such that each pool is used to amplify target loci in a different reaction volume. Each pool is used to simultaneously amplify a large number of target loci (or a subset of target loci) in a single reaction volume. Preferably, primers are designed and selected for a large portion of the most desirable target loci or for all of the target loci. A set of candidate target loci can be selected as described in Examples 16 or 20 based on the particular polymorphisms or mutations of interest. In some embodiments, one or more of the following type of target loci are included: SNPs, short indels, long indels, exons, and

combinations thereof. In some embodiments for target loci that are short indels, the PCR primer or primer pair targets a sequence of adjacent base pairs; and the indel is completely covered by one sequencing read. In some embodiments for target loci that are large indels, two primer pairs are used to target a pair of breakpoints at the boundaries of the indel. In this case, the two primer pairs are designed such that when the deletion is present there is a PCR product and the two primer pairs are selected together for inclusion in the same pool (the four primers are treated by the algorithm as a single assay rather than two assays). In some embodiments for target loci that are exons, a set of primers pairs are designed to tile the full exon.

[0957] For each candidate locus, one or more PCR primer pairs are designed using the Primer3 program (available at the worldwide web at primer3.sourceforge.net; libprimer3 release 2.2.3, which is hereby incorporated by reference in its entirety). If there are no feasible designs for PCR primers for a particular target locus, then that target locus is eliminated from further consideration. In some embodiments, each target base is covered by at least two independent PCR assays (such as two independent primer pairs that will amplify the target base) and preferably by four assays, although not all of the available assays for a target must be used. In some embodiments, no targets are omitted. Desirably, the algorithm produces as few pools as possible but may produce more than one pool. In some embodiments, two different primer pairs that are in close proximity in the genome (such as within 2 kbases or 1 kbase) and whose forward primers are on the same strand are not be assigned to the same pool. This constraint avoids primer interference in the extension-and-ligation amplification method. In some embodiments in which the PCR will be performed using a polymerase with low 5'→3' exonuclease and/or low strand displacement activity, different primer pairs that are in close proximity in the genome and whose forward primers are on the same strand can be assigned to the same pool since the with low 5'→3' exonuclease and/or low strand displacement activity of the polymerase will reduce or prevent primer interference and allow nearby or adjacent amplicons to be produced.

Step 1

[0958] Build an interaction graph. Each node represents one assay (such as one primer pair). Each edge represents a conflict between two assays. There are three types. Interaction edges represent a potential primer dimer and have a score indicating the interaction strength. Proximity edges represent physical proximity of the primer binding sites which may result in interference. Target edges represent redundant designs associated with the same target (a special case of a proximity edge).

Step 2

[0959] Select an initial value for the maximum interaction score (e.g., 95% of the maximum score).

Step 3

[0960] Compute a score such as a utility score for each assay as follows using steps 3A and 3B.

Step 3A

[0961] Calculate a score for each assay based on one or more of its intrinsic characteristics. For example, favor assays with amplicons close to the optimal length (such as 300 bp); favor assays with a shorter distance from the beginning of the amplicon to the target; and/or penalize assays with primers overlapping known SNPs. Any other parameter, such as the parameters disclosed herein can also be included.

Step 3B

[0962] Multiply the score for each assay by a factor that varies from 0 to 1 according to the current coverage of the assay's target bases. This factor gives lower weight to targets that are already covered by assays. At the beginning of the algorithm this factor is 1 for all assays because none have been covered. Calculate the factor as follows. For each base in the target, compute a coverage score as $1/(2^c)$ where c is the number of previously-selected assays (in other pools) that cover that base. For instance, if three assays cover the base then the coverage score is $1/(2^3)=0.125$. The factor for the target is the maximum value of the coverage score for all bases in the target. For instance, if the target contains 10 bases, 3 bases are covered by 1 target, and 7 bases are covered by 3 targets, then the factor is $\text{MAX}(1/(2^1), 1/(2^3))=0.5$. The score in step 3A is then multiplied by this factor.

Step 4

[0963] Use a single iteration of the algorithm in Example 16 to design a pool given the current maximum interaction score: Construct a new graph with the assays that have not been assigned to a pool yet and with the edges that have weights exceeding the maximum interaction score. Remove nodes (assays) according to the algorithm in Example 16 until there are no edges left. The assay utility scores come from step 3 in this Example rather than the calculation used for Example 16.

Step 5

[0964] Save the assays selected in step 4 as a new pool and remove them from consideration. Then repeat steps 3 and 4 with the remaining assays, and iterate until all targets have sufficient coverage.

Step 6

[0965] If desired, evaluate the result. If the total number of pools meets the design goal then reduce the maximum interaction score; otherwise increase the maximum interaction score. Then go back to step 3. Iterate, using a binary search strategy to find the lowest maximum interaction score that produces the desired number of pools.

Step 7

[0966] Output the pools from the final iteration. After the selection process, the primers remaining in the pools may be used in any of the methods of the invention.

Example 22

[0967] The following Example illustrates an exemplary method for designing and selecting primers that can be used

in any of the multiplexed PCR methods of the invention. In some embodiments, the primers are divided into different pools (e.g., 2, 3, 4, 5, 6, or more different pools) such that each pool is used to amplify target loci in a different reaction volume. Any of the embodiments listed in Example 21 can be used for this Example as well.

[0968] This method uses a graph coloring algorithm.

Step 1

[0969] Select 2, 3 or 4 of the best assays (such as primer pairs) for each target locus from all of the available assays.

Step 2

[0970] Select an initial maximum interaction score.

Step 3

[0971] Build an interaction graph containing only edges that exceed the maximum interaction score.

Step 4

[0972] Color the graph such that no adjacent nodes have the same color (this is a standard problem with many heuristic solutions). Each color represents a different pool.

Step 5

[0973] Go back to step 3 and iterate, refining the maximum interaction score until the desired number of pools is achieved. In some embodiments, after the primers are selected in step 1, the algorithm assumes all assays must be included in a pool.

[0974] After the primers are divided into different pools, the pools may be used in any of the methods of the invention.

Example 23

[0975] This example illustrates three exemplary methods for calculating the limit of detection for any of the methods of the invention. These methods were used to calculate the limit of detection for single nucleotide variants (SNVs) in a tumor biopsy (FIG. 38) and a plasma sample (FIG. 39).

[0976] The first method (denoted "LOD-mr5" in FIGS. 38 and 39) calculates the limit of detection based on a minimum of 5 reads being chosen as the minimum number of times a SNV is observed in the sequencing data to have sufficient confidence the SNV is actually present. The limit of detection is based on whether the observed depth of read (DOR) is above this minimum of 5. The thin lines (LOD-z5.0) in FIGS. 38 and 39 indicate SNVs for which the limit of detection is limited by the DOR. In these cases, not enough reads were measured to reach the error limit of the assay. If desired, the limit of detection can be improved (resulting in a lower numerical value) for these SNVs by increasing the DOR.

[0977] The second method (denoted "LOD-zs5.0" in FIGS. 38 and 39) calculates the limit of detection based on the z-score. The Z-score is the number of standard deviations an observed error percentage is away from the background mean error. If desired, outliers can be removed and the z-score can be recalculated and this process can be repeated. The final weighted mean and the standard deviation

tion of the error rate are used to calculate the z-score. The mean is weighted by the DOR since the accuracy is higher when the DOR is higher.

[0978] For the exemplary z-score calculation used for this example, the background mean error and standard deviation were calculated from all the other samples of the same sequencing run weighted by their depth of read, for each genomic loci and substitution type. Samples were not considered in the background distribution if they were 5 standard deviations away from the background mean. The dashed lines in FIGS. 38 and 39 indicate SNVs for which the limit of detection is limited by the error rate. For these SNV's enough reads were taken to reach the 5 read minimum, and the limit of detection was limited by the error rate. If desired, the limit of detection can be improved by optimizing the assay to reduce the error rate.

[0979] The third method (denoted "LOD-zs5.0-mr5" in FIGS. 38 and 39) calculates the limit of detection based on the maximum value of the above two metrics.

[0980] For the analysis of a tumor sample shown in FIG. 38, the mean limit of detection was 0.36%, and the median limit of detection was 0.28%. The number of DOR limited (thin lines) SNVs was 934. The number of error rate limited (dashed lines) SNVs was 738.

[0981] For the analysis of cDNA in a plasma sample shown in FIG. 39, the mean limit of detection was 0.24%, and the median limit of detection was 0.09%. The number of DOR limited (thin lines) SNVs was 732. The number of error rate limited (dashed lines) SNVs was 921.

Example 24

[0982] This example illustrates the detection of CNVs and SNVs from the same single cell. The following primer libraries were used a library of ~28,000 primers for detecting CNVs, a library of ~3,000 primers for detecting CNVs, and library of primers for detecting SNVs. For analysis of a single cell, cells were serially diluted until there were 3 or 4 cells per droplet. An individual cell was pipetted and placed into a PCR tube. The cell was lysed using Protease K, salt, and DTT using the following thermocycling conditions: 56° C. for 20 minutes, 95° C. for 10 minutes, and then a 4° C. hold. For analysis of genomic DNA, DNA from the same cell line as the analyzed single cell was either purchased or obtained by growing the cells and extracting the DNA.

[0983] For amplification with the library of ~28,000 primers, the following PCR conditions were used a 40 uL reaction volume, 7.5 nM of each primer, and 2x master mix (MM). In some embodiments QIAGEN Multiplex PCR Kit is used for the master mix (QIAGEN catalog No. 206143; see, e.g., information available at the world wide web at qiagen.com/products/catalog/assay-technologies/end-point-pcr-and-rt-pcr-reagents/qiagen-multiplex-pcr-kit, which is hereby incorporated by reference in its entirety). The kit includes 2x QIAGEN Multiplex PCR Master Mix (providing a final concentration of 3 mM MgCl₂, 3x0.85 ml), 5x Q-Solution (1x2.0 ml), and RNase-Free Water (2x1.7 ml). The QIAGEN Multiplex PCR Master Mix (MM) contains a combination of KCl and (NH₄)₂SO₄ as well as the PCR additive, Factor MP, which increases the local concentration of primers at the template. Factor MP stabilizes specifically bound primers, allowing efficient primer extension by, e.g., HotStarTaq DNA Polymerase. HotStarTaq DNA Polymerase is a modified form of Taq DNA polymerase and has no polymerase activity at ambient temperatures. The following

thermocycling conditions were used for the first round of PCR: 95° C. for 10 minutes; 25 cycles of 96° C. for 30 seconds, 65° C. for 29 minutes, and 72° C. for 30 seconds; and then 72° C. for 2 minutes, and a 4° C. hold. For the second round of PCR a 10 ul reaction volume, 1xMM, and 5 nM of each primer was used. The following thermocycling conditions were used: 95° C. for 15 minutes; 25 cycles of 94° C. for 30 seconds, 65° C. for 1 minute, 60° C. for 5 minutes, 65° C. for 5 minutes, and 72° C. for 30 seconds; and then 72° C. for 2 minutes, and a 4° C. hold.

[0984] For the library of ~3,000 primers, exemplary reaction conditions include a 10 ul reaction volume, 2x MM, 70 mM TMAC, and 2 nM primer of each primer. For the library of primers for detecting SNVs, exemplary reaction conditions include a 10 ul reaction volume, 2x MM, 4 mM EDTA, and 7.5 nM primer of each primer. Exemplary thermocycling conditions include 95° C. for 15 minutes, 20 cycles of 94° C. for 30 seconds, 65° C. for 15 minutes, and 72° C. for 30 seconds; and then 72° C. for 2 minutes, and a 4° C. hold.

[0985] The amplified products were barcoded. One run of sequencing was performed with an approximately equal number of reads per sample.

[0986] FIGS. 40A and 40B show results from analysis of genomic DNA (FIG. 40A) or DNA from a single cell (FIG. 40B) using a library of approximately 28,000 primers designed to detect CNVs. Approximately 4 million reads were measured per sample. The presence of two central bands instead of one central band indicates the presence of a CNV. For three samples of DNA from a single cell, the percent of mapped reads was 89.9%, 94.0%, and 93.4%, respectively. For two samples of genomic DNA the percent of mapped reads was 99.1% for each sample.

[0987] FIGS. 41A and 41B show results from analysis of genomic DNA (FIG. 41A) or DNA from a single cell (FIG. 41B) using a library of approximately 3,000 primers designed to detect CNVs. Approximately 1.2 million reads were measured per sample. The presence of two central bands instead of one central band indicates the presence of a CNV. For three samples of DNA from a single cell, the percent of mapped reads was 98.2%, 98.2%, and 97.9%, respectively. For two samples of genomic DNA the percent of mapped reads was 98.8% for each sample. FIG. 42 illustrates the uniformity in DOR for these ~3,000 loci.

[0988] For calling SNVs, the call percent for true positive mutations was similar for DNA from a single cell and genomic DNA. A graph of call percent for true positive mutations for single cells on the y-axis versus that for genomic DNA on the x-axis yielded a curve fit of $y=1.0076x-0.3088$ with $R^2=0.9834$. FIG. 43 shows similar error call metrics for genomic DNA and DNA from a single cell. FIG. 44 shows that the error rate for detecting transition mutations was greater than for detecting transversion mutations, indicating it may be desirable to select transversion mutations for detection rather than transition mutations when possible. In some embodiments, at least 50, 60, 70, 80, 90, 95, 96, 98, 99, or 100% of the SNVs tested for are transversion mutations rather than transition mutations.

Example 25

[0989] The following is an example of multiplex PCR conditions in which the annealing temperature is significantly higher than the average or maximum melting (T_m) of the primers in the library. A 3,168-plex reaction was performed with 3,168 primer pairs to 3,168 different target loci.

For the PCR amplification a 20 ul total volume was used with 2 nM of each primer (3,168 pairs of forward and reverse primers), 70 mM TMAC (tetra-methyl ammonium chloride), and 7 ul library DNA or genomic DNA. The following thermocycling conditions were used: 95° C. for 10 minutes and then 25 cycles of 96° C. for 30 seconds, 65° C. for 20 minutes (this annealing temperature is higher than the T_m of the primers, listed above), and 72° C. for 30 seconds. Then, 72° C. for 2 minutes and a 4° C. hold were used.

[0990] The minimum T_m (the lowest numerical value for the T_m for any of the primers) for this primer library is 54.0° C. The maximum T_m (the highest numerical value for the T_m for any of the primers) for this primer library is 60.36° C. The average T_m (average value of the T_m values of the primers) for this primer library is 55.25° C. These T_m values were calculated using the following exemplary method for calculating T_m values. This method is used by the Primer3 program (the worldwide web at primer3.sourceforge.net, which is hereby incorporated by reference in its entirety) to calculate T_m values. In some embodiments, one or more of the following conditions are assumed for this calculation: temperature: of 60.0° C., primer concentration of 100 nM, and/or salt concentration of 100 mM. In some embodiments, other conditions are assumed for this calculation, such as the conditions that will be used for multiplex PCR with the library.

$$T_m = \Delta H / (\Delta S + R \cdot \ln(C/4))$$

[0991] Below is documentation from the Primer3 program for its T_m calculations; PRIMER_TM_FORMULA (int; default 0) specifies details of melting temperature calculation. This is new in version 1.1.0, and added by Mado Remm and Triinu Koressaar (the world wide web at primer3.ut.ee/primer3web_help.htm # PRIMER_TM_FORMULA, which is hereby incorporated by reference in its entirety). A value of 0 directs primer3 to a backward compatible calculation (in other words, the only calculation available in previous version of primer3). This backward compatible calculation uses the table of thermodynamic parameters in the paper (Breslauer K J et al. (1986) "Predicting DNA duplex stability from the base sequence" Proc Natl Acad Sci 83:4746-50, dx.doi.org/10.1073/pnas.83.11.3746, which is hereby incorporated by reference in its entirety), and the method in the paper (Rychlik W, Spencer W J and Rhoads R E (1990) "Optimization of the annealing temperature for DNA amplification in vitro", Nucleic Acids Res 18:6409-12, dx.doi.org/10.1093/nar/18.21.6409, which is hereby incorporated by reference in its entirety).

[0992] A value of 1 (which is recommended) directs primer3 to use the table of thermodynamic values and the method for melting temperature calculation suggested in the following paper (SantaLucia JR (1998) "A unified view of polymer, dumbbell and oligonucleotide DNA nearest-neighbor thermodynamics", Proc Natl Acad Sci 95:1460-65, dx.doi.org/10.1073/pnas.95.4.1460, which is hereby incorporated by reference in its entirety). The tag PRIMER_SALT_CORRECTIONS can be used to specify the salt correction method for melting temperature calculation.

[0993] The following is an example of calculating the melting temperature of an oligo with PRIMER_TM_FORMULA=1 and PRIMER_SALT_CORRECTIONS=1 recommended values for primer=CGTGACGTGACGGACT.

[0994] Using default salt and DNA concentrations gives

$$T_m = \Delta H / (\Delta S + R \cdot \ln(C/4))$$

[0995] where R is the gas constant (1.987 cal/K mol) and C is the DNA concentration.

$$\begin{aligned} \Delta H(\text{predicted}) &= \Delta H(CG) + \Delta H(GT) + \Delta H(TG) + \dots + \Delta H(CT) + \\ &\quad \Delta H(\text{init.w.term.GC}) + \Delta H(\text{init.w.term.AT}) \\ &= -10.6 + (-8.4) + (-8.5) + \dots + (-7.8) + 0.1 + 2.3 \\ &= -128.8 \text{ kcal/mol} \end{aligned}$$

where 'init.w.term GC' and 'init.w.term AT' are two initiation parameters for duplex formation: 'initiation with terminal GC' and 'initiation with terminal AT.'

$$\begin{aligned} \Delta S(\text{predicted}) &= \Delta S(CG) + \Delta S(GT) + \Delta S(TG) + \dots + \Delta S(CT) + \\ &\quad \Delta S(\text{init.w.term.GC}) + \Delta S(\text{init.w.term.AT}) \\ &= -27.2 + (-22.4) + (-22.7) + \dots + (-21.0) + \\ &\quad (-2.8) + 4.1 \\ &= -345.2 \text{ cal/k} \cdot \text{mol} \end{aligned}$$

$$\begin{aligned} \Delta S(\text{salt corrected}) &= \Delta S(\text{predicted}) + 0.368 \cdot 15(NN \text{ pairs}) \cdot \\ &\quad \ln(0.05M \text{ monovalent cations}) \\ &= -361.736 \end{aligned}$$

$$T_m = -128.800 / (-361.736 + 1.987 \cdot \ln((5 \cdot 10^{-8})/4)) = 323.704 \text{ K}$$

$$T_m(C) = 323.704 - 273.15 = 50.554 \text{ C}$$

Additional Applications

[0996] Because this method utilizes targeted amplification, it is uniquely poised to detect submicroscopic anomalies, such as microdeletions and microduplications. Although non-targeted methods like MPSS have been shown to detect the DiGeorge microdeletion syndrome, this required a sufficiently high level of genomic coverage so as to make the approach unfeasible. This is because non-targeted amplification will be several orders of magnitude less efficient on submicroscopic regions, as very small fraction of the sequencing reads will be informative. Additionally, the fact that the currently available methods have trouble accurately identifying ploidy state for the sex chromosomes suggests that they will also encounter variable amplification problems on smaller chromosomal segments.

[0997] Similarly, SNP based methods can detect UPD disorders, which are copy number-neutral anomalies that will not be detected by either the current noninvasive methods that rely on counting or the traditional invasive methods like amniocentesis and CVS that rely on cytogenetic karyotyping and/or fluorescence in situ hybridization. This is because SNP-based methods are uniquely able to distinguish individual haplotypes, whereas the clinically available MPSS-based and targeted methods amplify non-polymorphic loci and are thus unable to determine, for example, whether the chromosomes-of-interest originate from the same parent. This means that these microdeletion/microduplication and UPD syndromes, including Prader-Willi, Angelman, and Beckwith-Wiedemann syndromes, are generally not diagnosed prenatally, and are often initially misdiagnosed postnatally. This significantly delays therapeutic intervention. Additionally, because this method tar-

gets SNPs, this method will also facilitate parental haplotype reconstruction, allowing for detection of fetal inheritance of individual disease-linked loci (Kitzman J O, Snyder M W, Ventura M, et al. Noninvasive whole-genome sequencing of a human fetus. *Sci Transl Med* 2012; 4:137ra76, which is hereby incorporated by reference in its entirety).

[0998] The results presented here confirm the expanded scope of this method for identifying prenatal aneuploidy. Specifically, by amplifying and sequencing 19,488 SNPs, this method is able to determine copy number at chromosomes 13, 18, 21, X, and Y, and is uniquely expected to detect other chromosomal abnormalities, such as triploidy and UPD, that are not detected by any other clinically available non-invasive method. The increased clinical coverage and powerful sample-specific calculated accuracies suggest that this method may offer a viable adjunct to invasive testing for detecting fetal chromosomal aneuploidies.

Example 26

[0999] This example describes an exemplary method for detection of copy number variations in breast cancer samples using SNP-targeted massively multiplexed PCR. Evaluation of CNV in tumor tissues typically involves SNP microarray or aCGH. These methods have high whole-genome resolution, but require large amounts of input material, have high fixed costs, and do not work well on formaldehyde fixed-paraffin embedded (FFPE) samples. For this example, 28,000-plex SNP-targeted PCR with next generation sequencing (NGS) was used to target 1p, 1q, 2p, 2q, 4p16, 5p15, 7q11, 15q, 17p, 2211, 2213 and chromosomes 13, 18, 21 and X for detection of CNVs in breast cancer samples. Accuracy was validated on 96 samples with aneuploidies or microdeletions. Single-molecule sensitivity was established by analyzing single cells. Of 17 breast cancer samples (15 fresh frozen and 2 FFPE tumor tissues, 5 pairs of matched tumor and normal cell lines) analyzed, 16 (including both FFPEs) were observed with full or partial CNVs in one to 15 targets (average: 7.8); evidence of tumor heterogeneity was observed. The three tissues with one CNV all had a 1q duplication, the most frequent cytogenetic abnormality in breast carcinoma. The most frequent regions with CNVs were 1q, 7p, and 221. Only one tumor tissue (with 9 CNVs) had a region with LOH; this LOH was also detected in adjacent putatively normal tissue that lacked the other 8 CNVs. By contrast, 5 or more regions with LOH and a high total CNV incidence (average: 12.8) was detected in cell lines. Thus, massively multiplexed PCR offers an economical high-throughput approach to investigate CNVs in a targeted manner, and is applicable to difficult-to-analyze samples, such as FFPE tissues.

Example 27

[1000] This example further validates a massively multiplexed PCR methodology for chromosomal aneuploidy and CNV determination disclosed herein, sometimes referred to as CoNVERGe (Copy Number Variant Events Revealed Genotypically), in cancer diagnostics, and further illustrates the development and use of "PlasmArt" standards for PCR of ctDNA samples. PlasmArt standards include polynucleotides having sequence identity to regions of the genome known to exhibit CNV and a size distribution that reflects that of cfDNA fragments naturally found in plasma.

Sample Collection

[1001] Human breast cancer cell lines (HCC38, HCC1143, HCC1395, HCC1937, HCC1954, and HCC2218) and matched normal cell lines (HCC38BL, HCC1143BL, HCC1395BL, HCC1937BL, HCC1954BL, and HCC2218BL) were obtained from the American Type Culture Collection (ATCC). Trisomy 21 B-lymphocyte (AG16777) and paired father/child DiGeorge Syndrome (DGS) cell lines (GM10383 and GM10382, respectively) were from the Coriell Cell Repository (Camden, N.J.). GM10382 cells only have the paternal 2211.2 region.

[1002] We procured tumour tissues from 16 breast cancer patients, including 11 fresh frozen (FF) samples from Geneticist (Glendale, Calif.) and five formalin-fixed paraffin-embedded (FFPE) samples from North Shore-LIJ (Manhasset, N.Y.). We acquired matched buffy coat samples for eight patients and matched plasma samples for nine patients. FF tumour tissues and matched buffy coat and plasma samples from five ovarian cancer patients were from North Shore-LIJ. For eight breast tumour FF samples, tissue sub-sections were resected for analysis. Institutional review board approvals from Northshore/LIJ IRB and Kharkiv National Medical University Ethics Committee were obtained for sample collection and informed consent was obtained from all subjects.

[1003] Blood samples were collected into EDTA tubes. Circulating cell free DNA (containing ctDNA) was isolated from 1 mL plasma using the QIAamp Circulating Nucleic Acid Kit (Qiagen, Valencia, Calif.). Genomic DNA (gDNA) from FF tumor tissues, blood, and buccal samples was extracted using the DNeasy Blood and Tissue Kit (Qiagen).

[1004] To make the PlasmArt standards according to one exemplary method, first, 9×10^6 cells were lysed with hypotonic lysis buffer (20 mM Tris-Cl (pH 7.5), 10 mM NaCl, and 3 mM MgCl₂) for 15 min on ice. Then, 10% IGEPAL CA-630 (Sigma, St. Louis, Mo.) was added to a final concentration of 0.5%. After centrifugation at 3,000 g for 10 min at 4° C., pelleted nuclei were resuspended in 1× micrococcal nuclease (MNase) Buffer (New England BioLabs, Ipswich, Mass.) before adding 1000 U of MNase (New England BioLabs), and then incubated for 5 min at 37° C. Reactions were stopped by adding EDTA to a final concentration of 15 mM. Undigested chromatin was removed by centrifugation at 2,000 g for 1 min. Fragmented DNA was purified with the DNA Clean & Concentrator™ 500 kit (Zymo Research, Irvine, Calif.). Mononucleosomal DNA produced by MNase digestion was also purified and size-selected using AMPure XP magnetic beads (Beckman Coulter, Brea, Calif.). DNA fragments were sized and quantified with a Bioanalyzer DNA 1000 chip (Agilent, Santa Clara, Calif.).

[1005] To model ctDNA at different concentrations, different fractions of PlasmArts from HCC1954 and HCC2218 cancer cells were mixed with those from the corresponding matched normal cell line (HCC1954BL and HCC2218BL, respectively). Three samples at each concentration were analyzed. Similarly, to model allelic imbalances in plasma DNA in a focal 3.5 Mb region, we generated PlasmArts from DNA mixtures containing different ratios of DNA from a child with a maternal 2211.2 deletion and DNA from the father. Samples containing only the father's DNA were used as negative controls. Eight samples at each concentration were analyzed.

Massively Multiplexed PCR and DNA Sequencing

[1006] Massively multiplex PCR and DNA sequencing methods below were used to determine allele counts at a plurality of polymorphic loci with 3-6 million (M) reads/sample for cell lines, 1.5-7 M reads/sample for tumour tissues, 18 M reads/sample for FFPE-LCM samples, 6-7 M reads/sample for germline controls, and 18-25 M reads/sample for plasma. For two representative exemplary runs using the 3,168 SNP primer pair pool, an average of 20 million reads were used to obtain allele counts for plasma DNA libraries and 6 million reads were used to obtain allele counts for genomic DNA libraries from fresh-frozen human tumors. The percent of mapped reads (i.e. mapped to the human genome) on these two exemplary runs were 98% and 95%, respectively. The fraction of sequencing reads at a given locus with a particular allele (allele fraction) was the fractional abundance of the allele in a sample. These counts provided observed allele frequencies that were used by the data analysis methods provided immediately below in this Example to determine the ploidy state of a chromosome or chromosome segment of interest and/or to determine the average allelic imbalance of the sample.

[1007] Libraries were generated from the samples above. Adapters were ligated to DNA fragments and the fragments were amplified using the following protocol: 95° C., 2 min; 15×[95° C., 20 sec, 55° C., 20 sec, 68° C., 20 sec], 68° C. 2 min, 4° C. hold.

[1008] Multiplexed PCR allows simultaneous amplification of many targets in a single reaction. In this study, we targeted 3,168 SNPs, which were distributed across five chromosome arms as follows: 646 on 1p, 602 on 1q, 541 on 2p, 707 on 2q, and 672 on the 2211.2 focal region. These genomic regions were selected for convenience from SNP panels available in our laboratory. Target SNPs had at least 10% population minor allele frequency (1000 Genomes Project data; Apr. 30, 2012 release) to ensure that a sufficient fraction would be heterozygous in any given patient. For each SNP, multiple primers were designed to have a maximum amplicon length of 75 bp and a melting temperature between 54.0-60.5° C. To minimize the likelihood of primer dimer product formation, primer interaction scores for all possible combinations of primers were calculated, and primers with high scores were eliminated. The 3,186 SNP primer pair pool all had ΔG values greater than -4Kcal/mol. Candidate PCR assays were ranked and 3,168 assays were selected on the basis of target SNP minor-allele frequency, observed heterozygosity rate (from dbSNP), presence in HapMap, and amplicon length.

[1009] For PCR amplifications, 3,168 SNPs were amplified in a multiplex PCR reaction using one primer pair for each SNP, during 25 cycles, and sequencing barcodes were added in 12 additional cycles. Prior to sequencing, the barcoded products were pooled, purified with the QIAquick PCR Purification Kit (Qiagen), and quantified using the Qubit™ dsDNA BR Assay Kit (Life Technologies). Amplicons were sequenced using an Illumina HiSeq 2500 sequencer with 1.5-7 M reads/sample for tumor tissue DNA and 18-25 M reads/sample for plasma cfDNA.

[1010] For the 3,168 SNP multiplex PCR reaction, approximately 7 ul (approx. 1200 ng) of library DNA, such as DNA from a DNA library generated from plasma of a target individual, was used. The master mix included the following: 2× (twice manufacturer's recommended concentration) Qiagen master mix, 70 mM TMAC (tetramethyl-

ammonium chloride, Sigma), 2 nM each primer, and 7 ul nucleic acid library (~1200 ng total library input) (20 ul total volume).

[1011] The cycling conditions for the 3,168 SNP multiplex PCR reaction were as follows: 95° C., 15 min; 25×[96° C., 30 sec; 65° C., 20 min; 72° C., 30 sec]; 72° C., 2 min; 4° C. hold.

[1012] For the barcoding reaction, a 1× master mix was prepared that included the following: 1 uM forward primer (containing Illumina sequencing tag), 1 uM reverse primer (containing Illumina sequencing tag as well as internally-designed sequencing barcode), 1 ul of mmPCR product, diluted 1:2,000, and 1× Qiagen master mix. Barcoding cycling conditions were as follows: 95° C., 10 min; 12×[95° C., 30 sec; 70° C., 10 sec, 60° C., 30 sec; 65° C., 15 sec, 72° C., 15 sec]; 72° C., 2 min; 4° C. hold.

Data Analysis of Tumor Tissue Genomic DNA

[1013] For tumor tissue samples, CNVs were delineated by transitions between allele frequency distributions. Regions with at least 100 SNPs that had an allele ratio statistically different from 0.50 were considered to be of interest. More specifically, the analysis focused on regions with average allele ratios of <0.45 or >0.55 for loci that are heterozygous in the germline. A segmentation algorithm was used to exhaustively search DNA sequences in five chromosome arms as follows: 646 on 1p, 602 on 1q, 541 on 2p, 707 on 2q, and 672 on the 2211.2 for such regions, and iteratively selected them starting from the longest one until a region of 100 SNPs was reached. Once a ≥100 SNP region was determined to contain a CNV, it was further segmented by average allelic ratios with a minimum segment size of 50 SNPs if needed.

[1014] Fresh frozen tissue samples from three patients with breast cancer were also analyzed using Illumina CytoSNP-12 microarrays as previously described (Levy, B. et al. Genomic imbalance in products of conception: single-nucleotide polymorphism chromosomal microarray analysis. *Obstetrics and gynecology* 124, 202-209 (2014)).

[1015] Data Analysis of Circulating Tumor DNA

[1016] CNVs were identified by a maximum likelihood algorithm that searched for plasma CNVs in regions where the tumor sample from the same individual also had CNVs, using haplotype information deduced from the tumor sample. This algorithm modeled expected allelic frequencies across a set of average allelic imbalances at 0.025% intervals for three sets of hypotheses: (1) all cells are normal (no allelic imbalance), (2) some/all cells have a homolog 1 deletion or homolog 2 amplification, or (3) some/all cells have a homolog 2 deletion or homolog 1 amplification. For at least some of the analysis, modeling was performed up to 15% average allelic imbalance, although for the vast majority of samples AAI was less than or equal to 5%. The likelihood of each hypothesis was determined at each SNP using a Bayesian classifier based on a beta binomial model of expected and observed allele frequencies at all heterozygous SNPs, and then the joint likelihood across multiple SNPs was calculated taking linkage of the SNP loci into consideration. The maximum likelihood hypothesis from the comparison of expected to observed allele frequencies was then selected.

[1017] This algorithm also calculates the confidence of each CNV call by comparing the likelihoods of different

hypotheses. A confidence threshold of 99.9% was used in plasma samples to minimize false positive results.

[1018] For dimorphic SNPs that have alleles arbitrarily designated 'A' and 'B', the allele ratio of the A allele is $n_A/(n_A+n_B)$, where n_A and n_B are the number of sequencing reads for alleles A and B, respectively. Allelic imbalance is the difference between the allele ratios of A and B for loci that are heterozygous in the germline. This definition is analogous to that for SNVs, where the proportion of abnormal DNA is typically measured using mutant allele frequency, or $n_m/(n_m+n_r)$, where n_m and n_r are the number of sequencing reads for the mutant allele and the reference allele, respectively.

[1019] Allele frequency data was corrected for errors before it was used to generate individual probabilities. Errors that were corrected included allele amplification bias, ambient contamination, genotype contamination, and sequencing error. Ambient contamination refers to the contamination error across all SNPs in addition to sequencing errors, and genotype contamination refers to the additional contamination at some SNPs due to contamination from another sample. Ambient contamination and genotype contamination were determined on the same run as the on-test sample analysis by analyzing homozygous alleles in the sample. The ploidy status of a chromosomal segment was estimated using heterozygous loci for a test individual.

[1020] Best hypothesis was defined to be the one with the highest likelihood across all polymorphic loci. Likelihood at each locus was calculated using a beta binomial model of observed allele frequencies at each of the polymorphic loci, and the likelihood across a set of polymorphic loci was computed using the phase information deduced from the corresponding tumor sample.

[1021] A linear regression model was used to compare either expected AAI or tumor input DNA percentage and observed AAI determined by the CNV detection algorithm. $P<0.05$ was considered statistically significant. SigmaPlot 12.5 (Systat Software, San Jose, Calif.) and Matlab 7.12.0 R2011a (MathWorks, Natick, Mass.) were used.

[1022] Accordingly, to evaluate the sensitivity and reproducibility of CoNVERGe, especially when the proportion of abnormal DNA for a CNV, or average allelic imbalance (AAI), is low, we used it to detect CNVs in DNA mixtures comprised of a previously characterized abnormal sample titrated into a matched normal sample. The mixtures consisted of artificial cfDNA, termed "PlasmArt", with fragment size distribution approximating natural cfDNA (see above). In the first pair, a son's tumor DNA sample having a 3 Mb Focal CNV deletion of the 2211.2 region was titrated into a matched normal sample from the father at between 0-1.5% total cfDNA. CoNVERGe reproducibly identified CNVs corresponding to the known abnormality with estimated AAI of $>0.35\%$ in mixtures of $>0.5\% \pm 0.2\%$ AAI, failed to detect the CNV in 6/8 replicates at 0.25% abnormal DNA, and reported a value of $<0.05\%$ for all eight negative control samples. The AAI values estimated by CoNVERGe showed high linearity ($R^2=0.940$) and reproducibility (error variance=0.087). The assay was sensitive to different levels of amplification within the same sample. Based on these data a conservative detection threshold of 0.45% AAI could be used for subsequent analyses.

[1023] Two additional PlasmArt titrations, prepared from pairs of matched tumor and normal cell line samples and having CNVs on chromosome 1 or chromosome 2, were also

evaluated. Among negative controls, all values were $<0.45\%$, and high linearity ($R^2=0.952$ for HCC1954 1p, $R^2=0.993$ for HCC1954 1q, $R^2=0.977$ for HCC2218 2p, $R^2=0.967$ for HCC2218 2q) and reproducibility (error variance=0.190 for HCC1954 1p, 0.029 for HCC1954 1q, 0.250 for HCC2218 2p, and 0.350 for HCC2218 2q) were observed between the known input DNA amount and that calculated by CoNVERGe. The difference in the slopes of the regressions for regions 1p and 1q of one sample pair correlates with the relative difference in copy number observed in the B-allelic frequencies (BAFs) of regions p and 1q of the same sample, demonstrating the relative precision of the AAI estimate calculated by CoNVERGe.

[1024] CoNVERGe has application to a variety of sample sources including FFPE, Fresh Frozen, Single Cell, Germ-line control and cfDNA. We applied CoNVERGe to six human breast cancer cell lines and matched normal cell lines to assess whether it can detect somatic CNVs. Arm-level and focal CNVs were present in all six tumour cell lines, but were absent from their matched normal cell lines, with the exception of chromosome 2 in HCC1143 in which the normal cell line exhibits a deviation from the 1:1 homolog ratio. To validate these results on a different platform, we performed CytoSNP-12 microarray analyses, which produced consistent results for all samples. Moreover, the maximum homolog ratios for CNVs identified by CoNVERGe and CytoSNP-12 microarrays exhibited a strong linear correlation ($R^2=0.987$, $P<0.001$).

[1025] We next applied CoNVERGe to fresh-frozen (FF) and formalin-fixed, paraffin-embedded (FFPE) breast tumour tissue samples. In both sample types, several arm-level and focal CNVs were present; however, no CNVs were detected in DNA from matched buffy coat samples. CoNVERGe results were highly correlated with those from microarray analyses of the same samples ($R^2=0.909$, $P<0.001$ for CytoSNP-12 on FF; $R^2=0.992$, $P<0.001$ for OncoScan on FFPE). CoNVERGe also produces consistent results on small quantities of DNA extracted from laser capture microdissection (LCM) samples, for which microarray methods are not suitable.

Detection of CNVs in Single Cells with CoNVERGe

[1026] To test the limits of the applicability of this mmPCR approach, we isolated single cells from the six aforementioned cancer cell lines and from a B-lymphocyte cell line that had no CNVs in the target regions. The CNV profiles from these single-cell experiments were consistent between three replicates and with those from genomic DNA (gDNA) extracted from a bulk sample of about 20,000 cells. On the basis of the number of SNPs with no sequencing reads, the average assay drop-out rate for bulk samples was 0.48% (range: 0.41-0.60%), which is attributable to either synthesis or assay design failure. For single cells, the additional average assay drop-out rate observed was 0.39% (range: 0.19-0.67%). For single cell assays that did not fail (i.e. no assay drop-out occurred), the average single ADO rate calculated using heterozygous SNPs only was 0.05% (range: 0.00-0.43%). Additionally, the percentage of SNPs with high confidence genotypes (i.e. SNP genotypes determined with at least 98% confidence) was similar for both single cell and bulk samples and the genotype in the single cell samples matched those in the bulk sample (average 99.52%, range: 92.63-100.00%).

[1027] In single cells, allele frequencies are expected to directly reflect chromosome copy numbers, unlike in tumour

samples where this may be confounded by TH and non-tumour cell contamination. BAFs of $1/n$ and $(n-1)/n$ indicate n chromosome copies in a region. Chromosome copy numbers are indicated on the allele frequency plots for both single cells and matched gDNA samples.

Application of CoNVERGe to Plasma Samples

[1028] To investigate the ability of CoNVERGe to detect CNVs in real plasma samples, we applied our approach to cfDNA paired with a matched tumour biopsy from each of two stage II breast cancer patients and five late-stage ovarian cancer. In all seven patients, CNVs were detected in both FF tumour tissues and in the corresponding plasma samples. A total of 32 CNVs, at a level of $>0.45\%$ AAI, were detected in the seven plasma samples (range: $0.48\text{--}12.99\%$ AAI) over the five regions assayed, which represent about 20% of the genome. Note that the presence of CNVs in plasma cannot be confirmed due to the lack of alternative orthogonal methods.

[1029] Although AAI estimates may appear correlated with BAFs in tumour, direct proportionality should not necessarily be expected due to tumour heterogeneity. For example, in sample BC5, regions that have BAFs compatible with $N=11$ were detecting; combining this with the AAI calculation from the plasma sample leads to estimates for c of 2.33% and 2.67% for the two regions. Estimating c using the other regions in the sample give values between 4.46% and 9.53% , which clearly demonstrates the presence of tumour heterogeneity.

[1030] Further CNV analyses of tumor tissue samples

[1031] We applied our mmPCR-NGS method described herein to plasma samples from four stage II breast cancer patients (BC1-BC4), and analyzed the concordance between CNVs detected in plasma and those detected in multiple tumor samples from each corresponding patient. Thus, we analyzed 4-6 tissue subsections from a tumor from each of four patients with breast cancer using mmPCR-NGS. All subsections for each patient had a CNV detected in at least one of the five targeted genomic regions (1p, 1q, 2p, 2q, and 2211.2). A CNV was identified in at least one tumor subsection in 18/20 (90%) genomic regions. Among these 18 CNV-positive regions, 11 (61%) had a CNV detected in that particular region in all subsections.

[1032] Interestingly, different patterns of AAIs across these five chromosomal regions were observed among different tumor subsections. In patient BC1, for instance, a similar pattern of CNVs was observed for regions 2p, 2q, and 2211.2 in all four subsections, suggesting that these CNVs are clonal mutations. In contrast, only two of the four subsections had CNVs observed in the 1p region, and three of the four subsections had CNVs observed in the 1q region, suggesting that those CNVs were subclonal mutations. Similar patterns of possible clonal and subclonal CNVs were observed in patients BC3 and BC4, whereas BC2 appeared to be more homogenous.

[1033] In addition, even when a CNV was detected in all subsections for a particular patient, such as in the 1q region for patient BC3, the AAI often varied between subsections. Overall AAI patterns also differed between patients. Taken together, these findings suggest that mmPCR-NGS can be used to elucidate both intra- and inter-tumor clonal heterogeneity.

Concordance of CNVs in Tumors and Plasma cfDNA

[1034] To quantify the amount of overlap between CNVs detected in plasma cfDNA and those detected in tumor tissue gDNA, we used mmPCR-NGS to interrogate CNVs in tumor tissue samples and matching plasma samples from patients BC1-BC4. Seven of the 18 (39%) CNV-positive genomic regions identified in tumor subsections were also detected in the plasma ($0.77\text{--}5.80\%$ AAI). Considering only the 11 clonal CNVs—those that were detected in all tumor subsections—a CNV was detected in four (36%) of the patient-matched plasma samples (estimated AAI: $0.77\text{--}5.80\%$).

[1035] Among the seven subclonal CNVs—those that were not observed in all subsections—we detected a CNV in 3/7 (43%) of the regions (estimated AA: $1.24\text{--}3.36\%$) in the corresponding cfDNA. Of note, in these three regions (BC1, chromosome 1p; BC1, chromosome 1q; and BC4, chromosome 2p), a CNV was detected in 10/14 (71%) of the matched tumor subsections. In contrast, in the other four genomic regions that did not have a CNV detected in the corresponding plasma samples (BC3, chromosomal regions p, 2p, 2q, and 2211.2), we only detected a CNV in 7/24 (29%) of the tissue subsections. These data suggest that the more prevalent a subclonal CNV is within a tumor, the more likely it is to be observed in cfDNA.

[1036] In the 150 genomic regions assayed in 30 negative controls, there were no CNVs with AAIs $>0.45\%$ and confidence $>99.9\%$, which suggests that mmPCR-NGS has a low false-positive rate.

[1037] These data demonstrate that CNVs can be detected in plasma in a substantial fraction of samples, and suggest that the more prevalent a CNV is within a tumour, the more likely it is to be observed in cfDNA. Furthermore, CoNVERGe detected CNVs from a liquid biopsy that may have otherwise gone unobserved in a traditional tumour biopsy.

Example 28

[1038] This example provides details regarding certain exemplary sample preparation methods used for analysis of different types of samples. The sample preparation methods disclosed in this example, were used in other Examples provided herein, to generate nucleic acid templates spanning a plurality of SNP sites for next generation sequencing reactions. From these NGS reactions, allele counts were generated at a plurality of polymorphic loci. These counts were then used by the analytical methods provided herein, to determine the ploidy state of a chromosome or chromosome segment of interest and/or to determine the average allelic imbalance of a sample.

Single Cell CNV Protocol for 28,000-Plex PCR

[1039] Multiplexed PCR allows simultaneous amplification of many targets in a single reaction. Target SNPs were identified in each genomic region with 10% minimum population minor allele frequency (1000 Genomes Project data; Apr. 30, 2012 release). For each SNP, multiple primers, semi-nested, were designed to have an amplicon length of a maximum length of 75 bp and a melting temperature between $54\text{--}60.5^\circ\text{C}$. Primer interaction scores for all possible combinations of primers were calculated; primers with high scores were eliminated to reduce the likelihood of primer dimer product formation. Candidate PCR assays were ranked and selected on the basis of target SNP minor

allele frequency, observed heterozygosity rate (from dbSNP), presence in HapMap, and amplicon length.

[1040] In certain experiments, single cell samples were prepared and amplified using a mmPCR 28,000-plex protocol. The samples were prepared in the following way: For analysis of a single cell, cells were serially diluted until there were 3 or 4 cells per droplet. An individual cell was pipetted and placed into a PCR tube. The cell was lysed using Protease K, salt, and DTT using the following conditions: 56° C. for 20 minutes, 95° C. for 10 minutes, and then a 4° C. hold. For analysis of genomic DNA, DNA from the same cell line as the analyzed single cell was either purchased or obtained by growing the cells and extracting the DNA. The DNA was amplified in a 40 uL reaction volume containing Qiagen mp-PCR master mix (2×MM final conc), 7.5 nM primer conc. for 28K primer pairs having hemi-nested Rev primers under the following conditions: 95 C 10 min, 25×[96 C 30 sec, 65 C 29 min, 72 C 30 sec], 72 C 2 min, 4 C hold. The amplification product was diluted 1:200 in water and 2 ul added to STAR 2 (10 ul reaction volume) 1×MM, 5 nM primer conc. and PCR was performed using hemi-nested inner Fwd primer and tag specific Rev primer: 95 C 15 min, 25×[94 C 30 sec, 65 C 1 min, 60 C 5 min, 65 C 5 min, 72 C 30 sec], 72 C 2 min, 4 C hold.

[1041] Full sequence tags and barcodes were attached to the amplification products and amplified for 9 cycles using adaptor specific primers. Prior to sequencing, the barcoded library product were pooled, purified with the QIAquick PCR Purification Kit (Qiagen), and quantified using the Qubit[®] dsDNA BR Assay Kit (Life Technologies). Amplicons were sequenced using an Illumina HiSeq 2500 sequencer.

Extraction of DNA from a Blood/Plasma Sample

[1042] Blood samples were collected into EDTA tubes. The whole blood sample was centrifuged and separated into three layers: the upper layer, 55% of the blood sample, was plasma and contains cell-free DNA (cfDNA); the buffy coat middle layer contained leucocytes having DNA, <1% of total; and the bottom layer, 45% of the collected blood sample, contained erythrocytes, no DNA was present in this fraction as erythrocytes are enucleated. Circulating tumor DNA was isolated from at least 1 mL plasma using the QIAamp Circulating Nucleic Acid Kit, Qia-Amp (Qiagen, Valencia, Calif.), according to the manufacturer's protocol. In certain experiments genomic DNA (gDNA) from FF tumor tissues, blood, and buccal samples was extracted using the DNeasy Blood and Tissue Kit (Qiagen).

Plasma CNV Protocol for 3,168-Plex for Chromosomes 1p, 1q, 2p, 2q, and 22q11

[1043] Plasma DNA libraries were prepared and amplified using a mmPCR 3,168-plex protocol. The samples were prepared in the following way: Up to 20 mL of blood was centrifuged to isolate the buffy coat and the plasma. Plasma extraction of cfDNA and library preparation was performed. DNA was eluted in 50 uL TE buffer. The input for mmPCR was 6.7 uL of amplified and purified Natera plasma library at an input amount of approximately 1200 ng. The plasma DNA was amplified in a 20 uL reaction volume containing Qiagen mp-PCR master mix (2×MM final conc), 2 nM tagged primer conc. (total 12.7 uM) in 3,168-plex primer pools and PCR amplified: 95 C 10 min, 25×[96 C 30 sec, 65 C 20 min, 72 C 30 sec], 72 C 2 min, 4 C hold. The amplification product was diluted 1:2,000 in water and 1 ul added to the Barcoding-PCR in a 10 uL reaction volume. The barcodes were attached to the amplification products via PCR amplification for 12 cycles using tag specific primers. Products of multiple samples were pooled and then purified with QIAquick PCR Purification Kit (Qiagen) and eluted in

50 ul DNA suspension buffer. Samples were sequenced by NGS as described for the Single Cell CNV Protocol for 28,000-plex PCR.

Breast Cancer Feasibility SNV Panel from Plasma

[1044] cfDNA from breast cancer patient blood samples was prepared and amplified using 336 primer pairs that were distributed into four 84-plex pools. Natera plasma libraries were prepared as described for Plasma CNV Protocol for 3,168-plex for Chromosomes 1p, 1q, 2p, 2q, and 22q11. DNA was eluted in 50 uL TE buffer. The input for mPCR was 2.5 uL of amplified and purified Natera plasma library at an input amount of approximately 600 ng. SNPs were selected from the 1000 Genomes map for Humans, Group 19 and dbSNP to pick targets, but only SNPs from the 1000 Genomes were used to screen for minor allele frequencies. The plasma DNA was amplified in four parallel reactions of 84-plex primer pools, a 10 uL reaction volume containing Qiagen mp-PCR master mix (2×MM final conc.), 4 mM EDTA, 7.5 nM primer concentration (total 1.26 uM) and PCR amplified: 95 C 15 min, 25×[94 C 30 sec, 65 C 15 min, 72 C 30 sec], 72 C 2 min, 4 C hold. The amplification product of the 4 subpools were each diluted 1:200 in water and 1 ul added to the Barcoding-PCR reaction in a 10 uL reaction volume containing Q5 HS HF master mix (1×final), and 1 uM each barcoding primer and each of the pools were amplified in the following reaction: 98 C 1 min, 25×[98 C 10 sec, 70 C 10 sec, 60 C 30 sec, 65 C 15 sec, 72 C 15 sec], 72 C 2 min, 4 C hold. Libraries were purified with QIAquick PCR Purification Kit (Qiagen) and eluted in 50 ul DNA suspension buffer. Samples were sequenced by paired end sequencing.

Example 29

[1045] This example demonstrates that by using low primer concentrations such that primer amount is the limiting reactant in multiplex PCR in a workflow that is followed by next generation sequencing, uniformity of density of reads, and therefore limits of detection, across a pool of amplification reactions is improved. Some experiments were carried out for plasma CNV using the 3,168-plex panel according to Example 28 above except that the total reaction volume was 10 uL instead of 20 uL. Furthermore, PCR was carried out for 15, 20, or 25 cycles. Other experiments were carried out using the four 84-plex pools on breast cancer samples according to the protocol of Example 28 except that primer concentrations were 2 nM and PCR amplification was carried out for 15, 20, or 25 cycles.

[1046] Not to be limited by theory, it is believed that primer limited multiplex PCR provides improved depth of read uniformity for multiplex PCR before multi-read sequencing, such as high throughput or massively parallel sequencing, such as sequencing on an Illumina HiSeq or MiSeq system or an Ion Torrent PGM or Proton system, based on the following considerations: If some of the amplifications in a multiplex PCR have lower efficiencies than others, then with normal multiplex PCR we will end up with a wide range of depth of read ("DOR") values. However, if the amount of primer is limited, and the multiplex PCR is cycled more times than what it takes to exhaust the primers, then the more efficient amplifications will stop doubling (because they have no more primers to use) and the less efficient ones will continue to double; this will result in a more similar amount of amplification product for all of the amplification products. This will translate into a much more uniform distribution of the DOR.

[1047] The following calculations are used to determine the number of cycles that would exact a given amount of primer and starting nucleic acid template:

[1048] assume a given starting DNA input level: 100 k copies of each target (10^5 ; this is easily achieved with using amplified library).

[1049] assume we use 2 nM of each primer as an exemplary concentration, although other concentrations such as, for example, 0.2, 0.5, 1, 1.5, 2, 2.5, 5, or 10 nM could work too.

[1050] calculate the number of primer molecules for each primer: $2 \times 10^{-4} \times 9$ (molar concentration, 2 nM) $\times 10 \times 10^{-4} \times 6$ (reaction volume, 10 μ l) $\times 6 \times 10^{23}$ (number of molecules per mole, Avogadro's number) $= 12 \times 10^{-4} \times 9$

[1051] calculate the amplification fold needed to consume all primers: $12 \times 10^{-4} \times 9$ (number of primer molecules) $/ 10^{-5}$ (number of copies of each target) $= 12 \times 10^{-4}$

[1052] calculate the number of cycles needed to achieve this amplification fold, assuming 100% efficiency at each cycle: $\log_2(12 \times 10^{-4}) = 17$ cycles. (this is $\log_2$ because at each cycle, the number of copies doubles).

[1053] So for these conditions (100 k copies input, 2 nM primers, 10 μ l reaction volume, assuming 100% PCR efficiency at each cycle), the primers would be consumed after 17 PCR cycles.

[1054] However, the key assumption is that some of the products DO NOT have 100% efficiency, so without measuring their efficiencies (which is only practicable for a small number of them anyway), it would take more than 17 cycles to consume them.

[1055] For each of four 84-plex SNV PCR primer pools we observed improved DOR efficiency with increasing cycles from 15 to 20 to 25. Similar results were obtained for experiments using the 3,168-plex panel. The limit of detection decreased (i.e. SNV sensitivity increased) with increasing depth of read. Furthermore, the sensitivity was consistently better when detecting transversion mutations than transition mutations. It is likely that additional increases in DOR efficiency can be obtained with additional cycles when using primer-limiting multiplex PCR before multi-read sequencing.

[1056] Accordingly, in one aspect provided herein is a method of amplifying a plurality of target loci in a nucleic acid sample that includes (i) contacting the nucleic acid sample with a library of primers and other primer extension reaction components to provide a reaction mixture, wherein the relative amount of each primer in the reaction mixture compared to the other primer extension reaction components creates a reaction wherein the primers are present at a limiting concentration, and wherein the primers hybridize to a plurality of different target loci; and (ii) subjecting the reaction mixture to primer extension reaction conditions for sufficient number of cycles to consume or exhaust the primers in the library of primers, to produce amplified products that include target amplicons. For example, the plurality of different target loci can include at least 2, 3, 5, 10, 25, 50, 100, 200, 250, 500, 1,000, 2,000, 5,000, 7,500, 10,000, 20,000, 25,000, 30,000, 40,000, 50,000, 75,000, or 100,000 different target loci, and at most, 50, 100, 200, 250, 500, 1,000, 2,000, 5,000, 7,500, 10,000, 20,000, 25,000, 30,000, 40,000, 50,000, 75,000, 100,000, 200,000, 250,000, 500,000, and 1,000,000 different target loci to produce a reaction mixture.

[1057] The method in illustrative embodiments, includes determining an amount of primer that will be a rate limiting amount. This calculation typically includes estimating and/or determining the number of target molecules and involves analyzing and/or determining the number of amplification cycles performed. For example, in illustrative embodiments, the concentration of each primer is less than 100, 75, 50, 25, 10, 5, 2, 1, 0.5, 0.25, 0.2 or 0.1 nM. In various embodiments, the GC content of the primers is between 30 to 80%, such

as between 40 to 70% or 50 to 60%, inclusive. In some embodiments, the range of GC content (e.g., the maximum GC content minus minimum GC content, such as 80%-60%=a range of 20%) of the primers is less than 30, 20, 10, or 5%. In some embodiments, the melting temperature (T_m) of the primers is between 40 to 80° C., such as 50 to 70° C., 55 to 65° C., or 57 to 60.5° C., inclusive. In some embodiments, the range of melting temperatures of the primers is less than 20, 15, 10, 5, 3, or 1° C. In some embodiments, the length of the primers is between 15 to 100 nucleotides, such as between 15 to 75 nucleotides, 15 to 40 nucleotides, 17 to 35 nucleotides, 18 to 30 nucleotides, 20 to 65 nucleotides, inclusive. In some embodiments, the primers include a tag that is not target specific, such as a tag that forms an internal loop structure. In some embodiments, the tag is between two DNA binding regions. In various embodiments, the primers include a 5' region that is specific for a target locus, an internal region that is not specific for the target locus and forms a loop structure, and a 3' region that is specific for the target locus. In various embodiments, the length of the 3' region is at least 7 nucleotides. In some embodiments, the length of the 3' region is between 7 and 20 nucleotides, such as between 7 to 15 nucleotides, or 7 to 10 nucleotides, inclusive. In various embodiments, the test primers include a 5' region that is not specific for a target locus (such as a tag or a universal primer binding site) followed by a region that is specific for a target locus, an internal region that is not specific for the target locus and forms a loop structure, and a 3' region that is specific for the target locus. In some embodiments, the range of the length of the primers is less than 50, 40, 30, 20, 10, or 5 nucleotides. In some embodiments, the length of the target amplicons is between 50 and 100 nucleotides, such as between 60 and 80 nucleotides, or 60 to 75 nucleotides, inclusive. In some embodiments, the range of the length of the target amplicons is less than 100, 75, 50, 25, 15, 10, or 5 nucleotides.

[1058] In various embodiments of any of the aspects of the invention, the primer extension reaction conditions are polymerase chain reaction conditions (PCR). In various embodiments, the length of the annealing step is greater than 3, 5, 8, 10, or 15 minutes but less than 240, 120, 60, or 30 minutes. In various embodiments, the length of the extension step is greater than 3, 5, 8, 10, or 15 minutes but less than 240, 120, 60 or 30 minutes.

[1059] All patents, patent applications, and published references cited herein are hereby incorporated by reference in their entirety. While the methods of the present disclosure have been described in connection with the specific embodiments thereof, it will be understood that it is capable of further modification. Furthermore, this application is intended to cover any variations, uses, or adaptations of the methods of the present disclosure, including such departures from the present disclosure as come within known or customary practice in the art to which the methods of the present disclosure pertain, and as fall within the scope of the appended claims. For example, any of the methods disclosed herein for DNA can be readily adapted for RNA by including a reverse transcription step to convert the RNA into DNA. Examples that use polymorphic loci for illustration can be readily adapted for the amplification of nonpolymorphic loci if desired. Any of the methods disclosed herein can be used with a low level of multiplexing if desired (such as with less than 1,000, 750, 500, 250, 100, 75, 50, 25, or 10 different primers or different primer pairs to different target loci).

SEQUENCE LISTING

The patent application contains a lengthy "Sequence Listing" section. A copy of the "Sequence Listing" is available in electronic form from the USPTO web site (<https://seqdata.uspto.gov/?pageRequest=docDetail&DocID=US20200347454A1>). An electronic copy of the "Sequence Listing" will also be available from the USPTO upon request and payment of the fee set forth in 37 CFR 1.19(b)(3).

What is claimed is:

1. A method for amplifying and sequencing DNA, comprising:

tagging isolated cell free DNA with one or more universal tail adaptors to generate tagged products, wherein the isolated cell-free DNA is isolated from a blood sample collected from a subject;

amplifying the tagged products one or more times to generate final amplification products, wherein one of the amplifying steps introduces a barcode and one or more sequencing tags; and

sequencing a plurality of loci on the cell free DNA by conducting massively parallel sequencing on the final amplification products, wherein the plurality of loci comprises 25-2,000 loci.

2. The method of claim 1, wherein the tagging the cell free DNA comprises ligating the one or more universal tail adaptors to the cell free DNA.

3. The method of claim 2, wherein the one or more universal tail adaptors each comprise a first strand and a second strand, wherein a first end of each of the universal tail adaptors comprises a double-stranded section comprising the 5' portion of the first strand and the 3' portion of the second strand, wherein the first end is ligated to the cell free DNA.

4. The method of claim 3, wherein the amplifying the tagged products comprises a first amplifying and a second amplifying, wherein the first amplifying comprises a first target-specific primer that specifically anneals to a target sequence and a first adaptor primer having a nucleotide sequence identical to a first portion of the first strand to generate a first amplification product.

5. The method of claim 4, wherein the second amplifying comprises a second target-specific primer that specifically

anneals to the first amplification product and a second adaptor primer having a nucleotide sequence identical to a second portion of the first strand to generate the final amplification product.

6. The method of claim 5, wherein the second adaptor primer is nested relative to the first adaptor primer.

7. The method of claim 6, wherein the second target-specific primer comprises an index tag.

8. The method of claim 7, wherein the second amplifying further comprises an index primer comprising a sequence complementary to the index tag.

9. The method of claim 8, wherein the index primer comprises the barcode and a first sequencing tag.

10. The method of claim 1, wherein the one or more universal tail adaptors comprise a second barcode.

11. The method of claim 1, wherein the one or more universal tail adaptors comprise a second sequencing tag.

12. The method of claim 1, wherein the one or more universal tail adaptors comprise a first universal tail adaptor and a second universal tail adaptor.

13. The method of claim 12, wherein the tagging the cell free DNA comprises amplifying the cell free DNA with a first primer comprising the first universal tail adaptor and a second primer comprising the second universal tail adaptor.

14. The method of claim 12, wherein the amplifying the tagged products comprises a single amplifying.

15. The method of claim 14, wherein the amplifying the tagged products comprises a third primer and a fourth primer, wherein the third primer comprises a first sequencing tag and wherein the fourth primer comprises a second sequencing tag.

* * * * *