

Accurate whole human genome sequencing using reversible terminator chemistry

A list of authors and their affiliations appears at the end of the paper

DNA sequence information underpins genetic research, enabling discoveries of important biological or medical benefit. Sequencing projects have traditionally used long (400–800 base pair) reads, but the existence of reference sequences for the human and many other genomes makes it possible to develop new, fast approaches to re-sequencing, whereby shorter reads are compared to a reference to identify intraspecies genetic variation. Here we report an approach that generates several billion bases of accurate nucleotide sequence per experiment at low cost. Single molecules of DNA are attached to a flat surface, amplified *in situ* and used as templates for synthetic sequencing with fluorescent reversible terminator deoxyribonucleotides. Images of the surface are analysed to generate high-quality sequence. We demonstrate application of this approach to human genome sequencing on flow-sorted X chromosomes and then scale the approach to determine the genome sequence of a male Yoruba from Ibadan, Nigeria. We build an accurate consensus sequence from >30× average depth of paired 35-base reads. We characterize four million single-nucleotide polymorphisms and four hundred thousand structural variants, many of which were previously unknown. Our approach is effective for accurate, rapid and economical whole-genome re-sequencing and many other biomedical applications.

DNA sequencing yields an unrivalled resource of genetic information. We can characterize individual genomes, transcriptional states and genetic variation in populations and disease. Until recently, the scope of sequencing projects was limited by the cost and throughput of Sanger sequencing. The raw data for the three billion base (3 gigabase (Gb)) human genome sequence, completed in 2004 (ref. 1), was generated over several years for ~\$300 million using several hundred capillary sequencers. More recently an individual human genome sequence has been determined for ~\$10 million by capillary sequencing². Several new approaches at varying stages of development aim to increase sequencing throughput and reduce cost^{3–6}. They increase parallelization markedly by imaging many DNA molecules simultaneously. One instrument run produces typically thousands or millions of sequences that are shorter than capillary reads. Another human genome sequence was recently determined using one of these approaches⁷. However, much bigger improvements are necessary to enable routine whole human genome sequencing in genetic research.

We describe a massively parallel synthetic sequencing approach that transforms our ability to use DNA and RNA sequence information in biological systems. We demonstrate utility by re-sequencing an individual human genome to high accuracy. Our approach delivers data at very high throughput and low cost, and enables extraction of genetic information of high biological value, including single-nucleotide polymorphisms (SNPs) and structural variants.

DNA sequencing using reversible terminators

We generated high-density single-molecule arrays of genomic DNA fragments attached to the surface of the reaction chamber (the flow cell) and used isothermal 'bridging' amplification to form DNA 'clusters' from each fragment. We made the DNA in each cluster single-stranded and added a universal primer for sequencing. For paired read sequencing, we then converted the templates to double-stranded DNA and removed the original strands, leaving the complementary

strand as template for the second sequencing reaction (Fig. 1a–c). To obtain paired reads separated by larger distances, we circularized DNA fragments of the required length (for example, 2 ± 0.2 kb) and obtained short junction fragments for paired end sequencing (Fig. 1d).

We sequenced DNA templates by repeated cycles of polymerase-directed single base extension. To ensure base-by-base nucleotide incorporation in a stepwise manner, we used a set of four reversible terminators, 3'-O-azidomethyl 2'-deoxynucleoside triphosphates (A, C, G and T), each labelled with a different removable fluorophore (Supplementary Fig. 1a)⁸. The use of 3'-modified nucleotides allowed the incorporation to be driven essentially to completion without risk of over-incorporation. It also enabled addition of all four nucleotides simultaneously rather than sequentially, minimizing risk of misincorporation. We engineered the active site of 9°N DNA polymerase to improve the efficiency of incorporation of these unnatural nucleotides⁹. After each cycle of incorporation, we determined the identity of the inserted base by laser-induced excitation of the fluorophores and imaging. We added tris(2-carboxyethyl)phosphine (TCEP) to remove the fluorescent dye and side arm from a linker attached to the base and simultaneously regenerate a 3' hydroxyl group ready for the next cycle of nucleotide addition (Supplementary Fig. 1b). The Genome Analyzer (GA1) was designed to perform multiple cycles of sequencing chemistry and imaging to collect the sequence data automatically from each cluster on the surface of each lane of an eight-lane flow cell (Supplementary Fig. 2).

To determine the sequence from each cluster, we quantified the fluorescent signal from each cycle and applied a base-calling algorithm. We defined a quality (Q) value for each base call (scaled as by the phred algorithm¹⁰) that represents the likelihood of each call being correct (Supplementary Fig. 3). We used the Q-values in subsequent analyses to weight the contribution of each base to sequence alignment and detection of sequence variants (for example, SNP

calling). We discarded all reads from mixed clusters and used the remaining 'purity filtered' reads for analysis. Typically we generated 1–2 Gb of high-quality purity filtered sequence per flow cell from ~30–60-million single 35-base reads, or 2–4 Gb in a paired read experiment (Supplementary Table 1).

To demonstrate accurate sequencing of human DNA, we sequenced a human bacterial artificial chromosome (BAC) clone (bCX98J21) that contained 162,752 bp of the major histocompatibility complex on




Figure 1 | Preparation of samples. **a**, DNA fragments are generated, for example, by random shearing and joined to a pair of oligonucleotides in a forked adaptor configuration. The ligated products are amplified using two oligonucleotide primers, resulting in double-stranded blunt-ended material with a different adaptor sequence on either end. **b**, Formation of clonal single-molecule array. DNA fragments prepared as in **a** are denatured and single strands are annealed to complementary oligonucleotides on the flow-cell surface (hatched). A new strand (dotted) is copied from the original strand in an extension reaction that is primed from the 3' end of the surface-bound oligonucleotide; the original strand is then removed by denaturation. The adaptor sequence at the 3' end of each copied strand is annealed to a new surface-bound complementary oligonucleotide, forming a bridge and generating a new site for synthesis of a second strand (dotted). Multiple cycles of annealing, extension and denaturation in isothermal conditions result in growth of clusters, each ~1 μm in physical diameter. This follows the basic method outlined in ref. 33. **c**, The DNA in each cluster is linearized by cleavage within one adaptor sequence (gap marked by an asterisk) and denatured, generating single-stranded template for sequencing by synthesis to obtain a sequence read (read 1; the sequencing product is dotted). To perform paired-read sequencing, the products of read 1 are removed by denaturation, the template is used to generate a bridge, the second strand is re-synthesized (shown dotted), and the opposite strand is then cleaved (gap marked by an asterisk) to provide the template for the second read (read 2). **d**, Long-range paired-end sample preparation. To sequence the ends of a long (for example, >1 kb) DNA fragment, the ends of each fragment are tagged by incorporation of biotinylated (B) nucleotide and then circularized, forming a junction between the two ends. Circularized DNA is randomly fragmented and the biotinylated junction fragments are recovered and used as starting material in the standard sample preparation procedure illustrated in **a**. The orientation of the sequence reads relative to the DNA fragment is shown (magenta arrows). When aligned to the reference sequence, these reads are oriented with their 5' ends towards each other (in contrast to the short insert paired reads produced as shown in **a–c**). See Supplementary Fig. 17a for examples of both. Turquoise and blue lines represent oligonucleotides and red lines represent genomic DNA. All surface-bound

human chromosome 6 (accession AL662825.4, previously determined using capillary sequencing by the Wellcome Trust Sanger Institute). We developed a fast global alignment algorithm ELAND that aligns a read to the reference only if the read can be assigned a unique position with 0, 1 or 2 differences. We collected 0.17 Gb of aligned data for the BAC from one lane of a flow cell. Approximately 90% of the 35-base reads matched perfectly to the reference, demonstrating high raw read accuracy (Supplementary Fig. 4). To examine consensus coverage and accuracy, we used 5 Mb of 35-base purity filtered reads (30-fold average input depth of the BAC) and obtained 99.96% coverage of the reference. There was one consensus miscall, at a position of very low coverage (just above our cutoff threshold), yielding an overall consensus accuracy of >99.999%.

Detecting genetic variation of the human X chromosome

For an initial study of genetic variation, we sequenced flow-sorted X chromosomes of a Caucasian female (sample NA07340 originating from the Centre d'Etude du Polymorphisme Humain (CEPH)). We generated 278-million paired 30–35-bp purity filtered reads and aligned them to the human genome reference sequence. We carried out separate analyses of the data using two alignment algorithms: ELAND (see above) or MAQ (Mapping and Assembly with Qualities)¹¹. Both algorithms place each read pair where it best matches the reference and assign a confidence score to the alignment. In cases where a read has two or more equally likely positions (that is, in an exact repeat), MAQ randomly assigns the read pair to one position and assigns a zero alignment quality score (these reads are excluded from SNP analysis). ELAND rejects all non-unique alignments, which are mostly in recently inserted retrotransposons (see Supplementary Fig. 5). MAQ therefore provides an opportunity to assess the properties of a data set aligned to the entire reference, whereas ELAND effectively excludes ambiguities from the short read alignment before further analysis.

We obtained comprehensive coverage of the X chromosome from both analyses. With MAQ, 204 million reads aligned to 99.94% of the X chromosome at an average depth of 43×. With ELAND, 192 million reads covered 91% of the reference sequence, showing what can be covered by unique best alignments. These results were obtained after excluding reads aligning to non-X sequence (impurities of flow sorting) and apparently duplicated read pairs (Supplementary Table 2). We reasoned that these duplicates (~10% of the total) arose during initial sample amplification.

The sampling of sequence fragments from the X chromosome is close to random. This is evident from the distribution of mapped read depth in the MAQ alignment in regions where the reference is unique (Fig. 2a): the variance of this distribution is only 2.26 times that of a Poisson distribution (the theoretical minimum). Half of this excess variance can be accounted for by a dependence on G+C content. However, the average mapped read depth only falls below 10× in regions with G+C content less than 4% or greater than 76%, comprising in total just 1% of unique chromosome sequence and 3% of coding sequence (Fig. 2b).

We identified 92,485 candidate SNPs in the X chromosome using ELAND (Supplementary Fig. 6). Most calls (85%) match previous entries in the public database dbSNP. Heterozygosity (π) in this data set is 4.3×10^{-4} (that is, one substitution per 2.3 kb), close to a previously published X chromosome estimate (4.7×10^{-4})¹². Using MAQ we obtained 104,567 SNPs, most of which were common to the results of the ELAND analysis. The differences between the two sets of SNP calls are largely the consequence of different properties of the alignments as described earlier. For example, most of the SNPs found only by the MAQ-based analysis were at positions of low or zero sequence depth in the ELAND alignment (Supplementary Fig. 6c).

We assessed accuracy and completeness of SNP calling by compar-

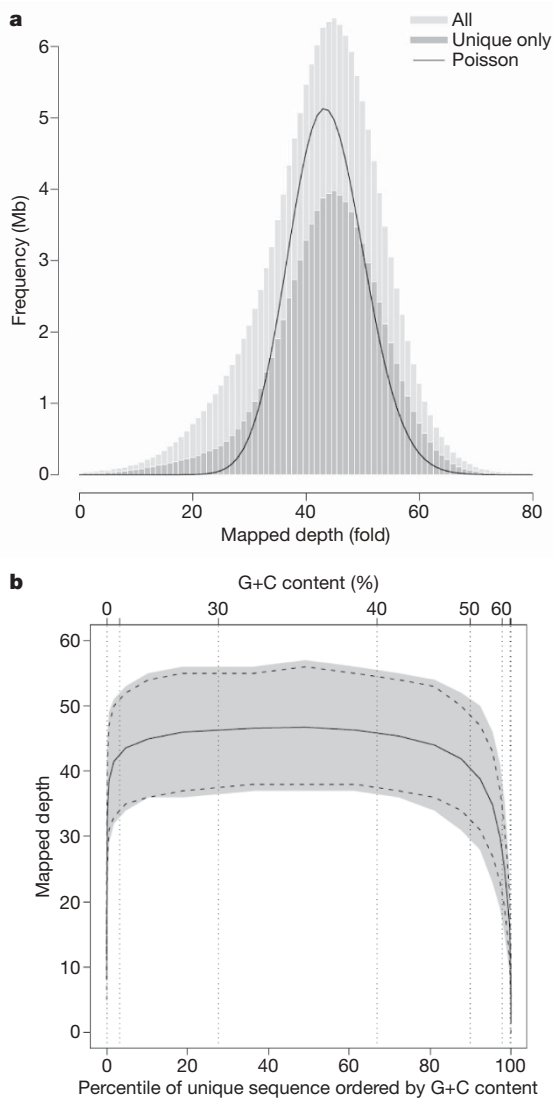


Figure 2 | X chromosome data. **a**, Distribution of mapped read depth in the X chromosome data set (NA07340), sampled at every 50th position along the chromosome and displayed as a histogram ('All'). An equivalent analysis of mapped read depth for the unique subset of these positions is also shown ('Unique only'). The solid line represents a Poisson distribution with the same mean. **b**, Distribution of X chromosome uniquely mapped reads as a function of G+C content. Note that the x axis is per cent G+C content and is scaled by percentile of unique sequence. The solid line is average mapped depth of unique sequence; the grey region is the central 80% of the data (10th to 90th centiles); the dashed lines are 10th and 90th centiles of a Poisson distribution with the same mean as the data.

agreement between sequence-based SNP calls and genotyping data (99.52% or 99.99% using ELAND or MAQ, respectively; Supplementary Table 3). There was complete concordance of all homozygous calls and a low level of 'under-calling' from the sequence data (denoted as 'GT>Seq' in Table 1) at a small number of the heterozygous sites, caused by inadequate sampling of one of the two alleles. The depth of input sequence influences the coverage and accuracy of SNP calling. We found that reducing the read depth to 15 \times still gives 97% coverage of genotype positions and only 1.27% of the heterozygous sites are under-called. We observed no other types of disagreement at any input depth (Supplementary Fig. 7).

We detected structural variants (defined as any variant other than a single base substitution) as follows. We found 9,747 short insertions (1.1 kb) (Supplementary Table 4). We found 1,600 deletions (1.1 kb) (Supplementary Table 5). We found 1,600 deletions (1.1 kb) (Supplementary Table 5). We found 1,600 deletions (1.1 kb) (Supplementary Table 5).

depth and/or anomalous read pair spacing, similar to previous approaches^{13–15}. We detected 115 indels in total, 77 of which were visible from anomalous read-pair spacing (see Supplementary Tables 4 and 5). We developed Resembl, an extension to the Ensembl browser¹⁶, to view all variants (Supplementary Fig. 9). Inversions can be detected when the orientation of one read in a pair is reversed (for example, see Supplementary Fig. 10). In general, inversions occur as the result of non-allelic homologous recombination, and are therefore flanked by repetitive sequence that can compromise alignments. We found partial evidence for other inversion events, but characterization of inversions from short read data is complex because of the repeats and requires further development.

Sequencing and analysis of a whole human genome

Our X chromosome study enabled us to develop an integrated set of methods for rapid sequencing and analysis of whole human genomes. We sequenced the genome of a male Yoruba from Ibadan, Nigeria (YRI, sample NA18507). This sample was originally collected for the HapMap project^{17,18} through a process of community engagement and informed consent¹⁹ and has also been studied in other projects^{20,21}. We were therefore able to compare our results with publicly available data from the same sample. We constructed two libraries: one of short inserts (~200 bp) with similar properties to the previous X chromosome library and one from long fragments (~2 kb) to provide longer-range read-pair information (see Supplementary Fig. 11 for size distributions). We generated 135 Gb of sequence (~4 billion paired 35-base reads; see Supplementary Table 6) over a period of 8 weeks (December 2007 to January 2008) on six GA1 instruments averaging 3.3 Gb per production run (see Supplementary Table 1 for example). The approximate consumables cost (based on full list price of reagents) was \$250,000. We aligned 97% of the reads using MAQ and found that 99.9% of the human reference (NCBI build 36.1) was covered with one or more reads at an average of 40.6-fold depth. Using ELAND, we aligned 91% of the reads over 93% of the reference sequence at sufficient depth to call a strong consensus (>three Q30 bases). The distribution of mapped read depth was close to random, with slight over-dispersion as seen for the X chromosome data. We observed comprehensive representation across a wide range of G+C content, dropping only at the very extreme ends, but with a different pattern of distribution compared to the X chromosome (see Supplementary Fig. 12).

We identified ~4 million SNPs, with 74% matching previous entries in dbSNP (Fig. 3). We found excellent agreement of our SNP calls with genotyping results: sequence-based SNP calls covered almost all of the 552,710 loci of HM550, with >99.5% concordance of sequencing versus genotyping calls (Table 1 and Supplementary Table 7a). The few disagreements were mostly under-calls of heterozygous positions (GT>Seq) in areas of low sequence depth, providing us with a false-negative rate of <0.35% from the ELAND analysis (see Table 1). The other disagreements (0.09% of all genotypes) included errors in genotyping plus apparent tri-allelic SNPs (Supplementary Table 7a). The main cause of genotype error (0.05% of all genotypes) is the existence of a second 'hidden' SNP close to the assayed locus that disrupts the genotyping assay, leading to loss of one allele and an erroneous homozygous genotype (Supplementary Figs 13 and 14).

To examine the accuracy of SNP calling in more detail, we compared our sequence-based SNP calls with 3.7 million genotypes (HM-All) generated for this sample during the HapMap project (Table 1 and Supplementary Table 7b)¹⁸ and found excellent concordance between the data sets. Disagreements included sequence-based under-calls of heterozygous positions in regions of low read depth. The slightly higher level of other disagreements (0.76% seen in this analysis compared to that of the HM550 data (0.09%) is in line with the higher level of sequencing errors (0.7% for the

Table 1 | Comparison of SNP calls made from sequence versus genotype data for the human genome (NA18507) and X chromosome (NA07340)

	ELAND			MAQ			
	X	Human	Human	X	Human	Human	Human
	HM550 (13,604 SNPs) (%)	HM550 (552,710 SNPs) (%)	HM-All (3,699,592 SNPs) (%)	HM550 (13,604 SNPs) (%)	HM550 (552,710 SNPs) (%)	HM-All (3,699,592 SNPs) (%)	Combined (530,750 SNPs) (%) (n)
Covered by sequence	99.77	99.60	99.24	99.91	99.74	99.29	99.78 529,589
Concordant calls	99.52	99.57	98.80	99.99	99.90	99.12	99.94 529,285
All disagreements	0.48	0.43	1.20	0.01	0.1	0.88	0.06 304
GT>Seq	0.48	0.35	0.46	0.01	0.03	0.15	0.02 130
Seq>GT	0	0.05	0.52	0	0.05	0.54	0.02 130
Other discordances	0	0.03	0.22	0	0.02	0.2	0.01 44

SNP panels referred to are HM550 (Illumina Infinium HumanHap550 BeadChip) and HM-All (complete data from phase 1 and phase 2 of the International HapMap Project). 'Combined' is a set of concordant genotypes from both sets (HM550 and HM-All; see text). GT>Seq denotes a heterozygous genotyping SNP call where there is a homozygous sequencing SNP call (one of the two alleles); Seq>GT denotes the converse (that is, a heterozygous sequencing SNP call where there is a homozygous genotyping call). Other discordances are differences in the two SNP calls that cannot be accounted for by one allele being missing from one call.

concordant calls in both the HM550 and HM-All genotype data sets. Comparing the results of the MAQ analysis to this high confidence set (see Table 1), we found 130 heterozygote under-calls GT>Seq (that is, a false-negative rate of 0.025%). There were also 130 heterozygote over-calls Seq>GT, but most of these are probably genotype errors as 82 have a nearby 'hidden' SNP and 3 have a nearby indel. A further 41 are tri-allelic loci, leaving at most 4 potential wrong calls by sequencing (that is, false-positive rate of 4 per 529,589 positions). Finally we selected a subset of novel SNP calls from the sequence data and tested them by genotyping. We found 96.1% agreement between sequence and genotype calls (Supplementary Table 8). However, the 47 disagreements included 10 correct sequencing calls (genotyping under-calls owing to hidden SNPs) and 7 sequencing under-calls. On this basis, therefore, the false-positive discovery rate for the one million novel SNPs is 2.5% (30 out of 1,206). For the entire data set of four million SNPs detected in this analysis, the false-positive and -negative rates both average <1%.

This genome from a Yoruba individual contains significantly more polymorphism than a genome of European descent. The autosomal heterozygosity (π) of NA18507 is 9.94×10^{-4} (1 SNP per 1,006 bp), higher than previous values for Caucasians (7.6×10^{-4} , ref. 12). Heterozygosity in the pseudoautosomal region 1 (PAR1) is substantially higher (1.92×10^{-3}) than the autosomal value. PAR1 (2.7 Mb)

at the tip of the short arm of chromosomes X and Y undergoes obligatory recombination in male meiosis, which is equivalent to $20\times$ the autosome average. This illustrates a clear correlation between recombination and nucleotide diversity. By contrast, the 0.33-Mb PAR2 region has a much lower recombination rate than PAR1; we observed that heterozygosity in PAR2 is identical to that of the autosomes in NA18507. Heterozygosity in coding regions is lower (0.54×10^{-3}) than the total autosome average, consistent with the model that some coding changes are deleterious and are lost as the result of natural selection²². Nevertheless, the 26,140 coding SNPs (Supplementary Fig. 15) include 5,361 non-conservative amino acid substitutions plus 153 premature termination codons (Supplementary Table 9), many of which are expected to affect protein function.

We performed a genome-wide survey of structural variation in this individual and found excellent correlation with variants that had been reported in previous studies, as well as detecting many new variants. We found 0.4 million short indels (1–16 bp; Supplementary Fig. 16), most of which are length polymorphisms in homopolymeric tracts of A or T. Half of these events are corroborated by entries in dbSNP, and 95 of 100 examined were present in amplicons sequenced from this individual in ENCODE regions, confirming the high specificity of this method of short indel detection. For larger structural variants (detected by anomalously spaced paired ends) we found that some were detected by both long and short insert data sets (Supplementary Fig. 17a), but most were unique to one or other data set. We observed two reasons for this: first, small events (<400 bp) are within the normal size variance of the long insert data; second, nearby repetitive structures can prevent unique alignment of read pairs (see Supplementary Fig. 17b, c). In some cases, the high resolution of the short insert data permits detection of additional complexity in a structural rearrangement that is not revealed by the long insert data. For example, where the long insert data indicate a 1.3-kb deletion in NA18507 relative to the reference, the short insert data reveal an inversion accompanied by deletions at both breakpoints (Fig. 4). We carried out *de novo* assembly of reads in this region and constructed a single contig that defines the exact structure of the rearrangement (data not shown).

We discovered 5,704 structural variants ranging from 50 bp to >35 kb where there is sequence absent from the genome of NA18507 compared to the reference genome. We observed a steadily decreasing number of events of this type with increasing size, except for two peaks (Supplementary Fig. 18). Most of the events represented by the large peak at 300–350 bp contain a sequence of the AluY family. This is consistent with insertion of short interspersed nuclear elements (SINEs) that are present in the reference genome but missing from the genome of NA18507. Similarly, the second




Figure 3 | SNPs identified in the human genome sequence of NA18507. **a**, Number of SNPs detected by class and percentage in dbSNP (release 128). Results from ELAND and MAQ alignments are reported separately. **b**, Analysis of SNPs detected in each analysis reveals extensive overlap. The percentage of NA18507 SNP calls that match previous entries in dbSNP is lower than that of our X chromosome study (see Supplementary Fig. 6). We




Figure 4 | Homozygous complex rearrangement detected by anomalous paired reads. The rearrangement involves an inversion of 369 bp (blue–turquoise bar in the schematic diagram) flanked by deletions (red bars) of 1,206 and 164 bp, respectively, at the left- and right-hand breakpoints. **a**, Summary tracks in the Resembl browser, denoting scale, simulated alignability of reads to reference (blue plot), actual aligned depth of coverage by NA18507 reads (green plot), density of anomalous reads indicating structural variants (red plot; peaks denote ‘hotspots’) and density of singleton reads (pink plot). **b**, Anomalous long-insert read pairs (orange lines denote DNA fragment; blocks at either end denote each read); the data indicate loss of ~ 1.3 kb in NA18507 relative to the reference. **c**, Anomalous short-insert pairs of two types (red and pink) indicate an inverted sequence flanked by two deletions. **d**, Normal short-insert read-pair alignments (each green line denotes the extent of the reference that is covered by the short fragment, including the two reads). **e**, The schematic diagram depicts the arrangement of normal and anomalous read pairs relative to the rearrangement. Top line, structure of NA18507; second line, structure of reference sequence. Green bars denote sequence that is collinear in the reference and NA18507 genomes. The turquoise–blue bar illustrates the inverted segment. Red bars indicate the sequences present in the reference but absent in NA18507. Arrows denote orientation of reads when aligned to the reference. The display in **a–d** is a composite of screen shots of the same window, overlapped for display purposes.

Supplementary Fig. 20. The ‘singleton’ reads on either side of the event, which have partners that do not align to the reference, form part of a *de novo* assembly that precisely defines the novel sequence and breakpoint (Supplementary Fig. 21).

Effect of sequence depth on coverage and accuracy

We investigated the impact of varying input read depth (and hence cost) on SNP calling using chromosome 2 as a model. SNP discovery increases with increasing depth: essentially all homozygous positions are detected at $15\times$, whereas heterozygous positions accumulate more gradually to $33\times$ (Fig. 5a). This effect is influenced by the stringency of the SNP caller. To call each allele in this analysis we required the equivalent of two high-quality Q30 bases (as opposed to three used in full depth analyses). Homozygotes could be detected at read depth of $2\times$ or higher, whereas heterozygote detection required at least double this depth for sampling of both alleles. Missing calls (not covered by sequence) and discordances between sequence-based SNP calls and genotype loci (mostly under-calls of heterozygotes due to low depth) progressively reduced with increasing depth (Fig. 5b). We observed very few other types of discordance at any depth; many of these are genotyping errors as described above.

Concluding remarks

Reversible terminator chemistry is a defining feature of this sequencing approach, enabling each cycle to be driven to completion while minimizing misincorporation. The result is a system that generates accurate data at very high throughput and low cost. We determined an accurate whole human genome sequence in 8 weeks to an average depth of $\sim 40\times$. We built a consensus sequence, optimized methods for analysis, assessed accuracy and characterized the genetic variation of this individual in detail.

We assessed accuracy relative to genotype data over the entire fraction of the human sequence where SNP calling was possible ($>90\%$). We established very low false-positive and -negative rates for the \sim four million SNPs detected ($<1\%$ over-calls and under-calls). This compares favourably with previous individual genome analyses which reported a 24% under-calling of heterozygous positions^{2,7}.

Paired reads were very powerful in all areas of the analysis. They provided very accurate read alignment and thus improved the accuracy and coverage of consensus sequence and SNP calling. They were essential for developing our short indel caller, and for detecting larger structural variants. Our short indel caller and detection of structural

many cases. We found good correspondence between our results and the data of ref. 23, which reported 148 deletions of <100 kb in this individual on the basis of abnormal fosmid paired-end spacing. We found supporting evidence for 111 of these events. We detected a further 2,245 indels in the range 60–160 bp, which are sequences

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts



Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research



With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips



Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.