

SYSTEMS AND METHODS FOR VISUALIZING STRUCTURAL VARIATION AND PHASING INFORMATION

RELATED APPLICATIONS

[0001] This application is related to United States Patent Application No. 62/120,873, entitled “Systems and Methods for Visualizing Structural Variation and Phasing Information,” filed February 25, 2015, which is hereby incorporated by reference herein in its entirety.

[0002] This application is also related to United States Patent Application No. 62/102,926, entitled “Systems and Methods for Visualizing Structural Variation and Phasing Information,” filed January 13, 2015, which is hereby incorporated by reference herein in its entirety.

TECHNICAL FIELD

[0003] This specification describes technologies relating to visualizing structural variation and phasing information in nucleic acid sequencing data.

BACKGROUND

[0004] Haplotype assembly from experimental data obtained from human genomes sequenced using massively parallelized sequencing methodologies has emerged as a prominent source of genetic data. Such data serves as a cost-effective way of implementing genetics based diagnostics as well as human disease study, detection, and personalized treatment.

[0005] The long-range information provided by such massively parallelized sequencing methodologies is disclosed, for example, in United States Patent Application No. 62/072,214, filed October 29, 2014, entitled “Analysis of Nucleic Acid Sequences.” Such techniques greatly facilitate the detection of large-scale structural variations of the genome, such as translocations, large deletions, or gene fusions. Other examples include, but are not limited to the sequencing-by-synthesis platform (ILLUMINA), Bentley *et al.*, 2008, “Accurate whole human genome sequencing using reversible terminator chemistry, *Nature* 456:53-59; sequencing-by-ligation platforms (POLONATOR; ABI SOLiD), Shendure *et al.*, 2005, “Accurate Multiplex Polony Sequencing of an Evolved bacterial Genome” *Science* 309:1728-1732; pyrosequencing platforms (ROCHE 454), Margulies *et al.*, 2005, “Genome sequencing in microfabricated high-density picoliter reactors,” *Nature* 437:376-380; and

single-molecule sequencing platforms (HELICOS HELISCAPE); Pushkarev *et al.*, 2009, “Single-molecule sequencing of an individual human genome,” *Nature Biotech* 17:847-850, (PACIFIC BIOSCIENCES) Eid *et al.*, “Real-time sequencing form single polymerase molecules,” *Science* 323:133-138, each of which is hereby incorporated by reference in its entirety.

[0006] The availability of haplotype data spanning large portions of the human genome, the need has arisen for ways in which to efficiently work with this data in order to advance the above stated objectives of diagnosis, discovery, and treatment, particularly as the cost of whole genome sequencing for a personal genome drops below \$1000. To computationally assemble haplotypes from such data, it is necessary to disentangle the reads from the two haplotypes present in the sample and infer a consensus sequence for both haplotypes. Such a problem has been shown to be NP-hard. See Lippert *et al.*, 2002, “Algorithmic strategies for the single nucleotide polymorphism haplotype assembly problem,” *Brief. Bioinform* 3:23-31, which is hereby incorporated by reference.

[0007] The assembly view Consed supports visualization of reads obtained from the above-identified sequencing methods. See Gordon 1998, “Consed: A graphical tool for sequencing finishing,” *Genome Research* 8:198-202.

[0008] Another visualization tool is EagleView. See Huang and Marth, 2008, “EagleView: A genome assembly viewer for next-generation sequencing technologies,” *Genome Research* 18:1538-1543.

[0009] Still another such viewer is HapEdit. See Kim *et al.*, “HapEdit: an accuracy assessment viewer for haplotype assembly using massively parallel DNA-sequencing technologies,” *Nucleic Acids Research*, 2011, 1-5. HapEdit provides tools for assessing the accuracy of Haplotype assemblies and permits a user to fit the composition rates of reads sequence by numerous different sequencing technologies.

[0010] While the above-disclosed programs are each significant advancements in their own right, they do not adequately address the need in the art for tools for visually assessing structural variants (*e.g.*, deletions, duplications, copy-number variants, insertions, inversions, translocations, long terminal repeats (LTRs), short tandem repeats (STRs), and a variety of other useful characterizations) in sequencing data.

SUMMARY

[0011] Technical solutions (e.g., computing systems, methods, and non-transitory computer readable storage mediums) for visually assessing structural variants are provided. With platforms such as those disclosed in United States Patent Application No. 62/072,214, filed October 29, 2014, entitled “Analysis of Nucleic Acid Sequences,” which is hereby incorporated by reference, the genome is fragmented and partitioned and barcoded prior to the target identification. Therefore the integrity of the barcode information is maintained across the genome. The barcode information is used to identify potential structural variation breakpoints by detecting regions of the genome that show significant barcode overlap. They are also used to obtain phasing information.

[0012] The following presents a summary of the invention in order to provide a basic understanding of some of the aspects of the invention. This summary is not an extensive overview of the invention. It is not intended to identify key/critical elements of the invention or to delineate the scope of the invention. Its sole purpose is to present some of the concepts of the invention in a simplified form as a prelude to the more detailed description that is presented later.

[0013] One aspect of the present disclosure is a system for providing structural variation or phasing information over a network connection to a remote client computer. The system comprises one or more microprocessors, a persistent memory and a non-persistent memory. The persistent memory (e.g., a hard drive) and the non-persistent memory (e.g., RAM memory) collectively store one or more nucleic acid sequence datasets. Each respective nucleic acid sequencing dataset in the one or more nucleic acid sequence datasets corresponds to at least one target nucleic acid in a respective sample in a plurality of samples. The respective sample is associated with a reference genome of a species that may serve as a benchmark for analysis of the respective sample in some embodiments. For instance, in some embodiments the respective sample is mapped to the reference genome and the reference genome is thereby used as a template (reference) to parse queries to visualize portions of the respective sample. For instance, in some embodiments a sample is from a human subject. In such instance, a human genome (as opposed to a genome from a different species) serves as the reference genome and the respective sample is mapped to the human genome. In this way, requests to visual sequences or sequence variations in certain human chromosomes, or portions thereof from the sample, can be interpreted and handled using the disclosed systems and methods, based on such mapping to the reference genome.

[0014] The respective nucleic acid sequencing dataset comprises (i) a header, (ii) a synopsis, and (iii) a data section. The data section comprises a plurality of aligned sequence reads from the sample and information about each variant call made. Advantageously, the data section is extensible and can store additional data. Each respective sequencing read in the plurality of sequencing reads comprises a first portion that corresponds to a subset of at least one target nucleic acid in the respective sample and a second portion that encodes a respective identifier for the respective sequencing read in a plurality of identifiers. Each respective identifier is independent of the sequence of the at least one target nucleic acid. Sequencing reads in the plurality of sequencing reads collectively include the plurality of identifiers.

[0015] The persistent memory and the non-persistent memory further collectively store one or more programs that use the one or more microprocessors to provide a haplotype visualization tool to a client for installation on the remote client computer. The system receives a request, sent from the client over a network connection (e.g., Internet), for structural variation or phasing information using a first dataset in the one or more datasets. Responsive to receiving the request, the request is automatically filtered by performing a method comprising loading the header and the synopsis of the first dataset into the non-persistent memory if not already loaded into the non-persistent memory while retaining the data section in persistent memory. In the method, the request is compared (analyzed against) the synopsis of the first dataset thereby identifying one or more portions of the data section of the first dataset. These one or more identified portions of the data section are, in turn, loaded into non-persistent memory. Structural variation or phasing information is formatted for display on the client computer using the first dataset. Then the formatted structural variation or phasing information is transmitted over the network connection to the client device for display on the client device.

[0016] In some embodiments, the header delineates a plurality of components in the respective nucleic acid sequencing dataset. In some embodiments the plurality of components comprises two or more components, three or more components, four or more components or five or more components selected from the group consisting of a summary, an index to variant call data, a phase block track, a refseq index track, a gene track, an exon track, an index to read data, a structural variant dataset track, an index to a target dataset, and an index to a fragment dataset.

[0017] In some embodiments, the plurality of components comprises the summary and this summary comprises two or more items, three or more items, four or more items, five or more items, or six or more items in the group consisting of: a percentage of known SNPs

phased in the respective nucleic acid sequencing dataset, a longest phase block in the respective nucleic acid sequencing dataset, a number of unique barcodes used in the respective nucleic acid sequencing dataset, an average fragment length in the respective nucleic acid sequencing dataset, a mean of the average fragment length in the respective nucleic acid sequencing dataset, a percentage of fragments greater than a lower threshold in the respective nucleic acid sequencing dataset, a fragment length histogram in the respective nucleic acid sequencing dataset, an N50 phase block size in the respective nucleic acid sequencing dataset, a phase block histogram in the respective nucleic acid sequencing dataset, a number of sequence reads represented by respective the nucleic acid sequencing dataset, a median insert size in the respective nucleic acid sequencing dataset, a median depth in the respective nucleic acid sequencing dataset, a percent of the target genome with zero coverage in the respective nucleic acid sequencing dataset, a mapped reads percentage for the respective nucleic acid sequencing dataset, a PCR duplication percentage for the respective nucleic acid sequencing dataset, a coverage histogram for the in the respective nucleic acid sequencing dataset, an identity of a test nucleic acid that forms the basis for the respective nucleic acid sequencing dataset, a genome source for the respective nucleic acid sequencing dataset, a sex of an organism that originated the at least one test nucleic acid of the respective nucleic acid sequencing dataset, a sex of the organism that originate the respective sample of the in the respective nucleic acid sequencing dataset, a dataset file format version of the in the respective nucleic acid sequencing dataset, and a pointer to a plurality of structural variant calls made for the respective nucleic acid sequencing dataset. Advantageously, as this non-limiting example of the list of information indicates, the disclosed nucleic acid sequencing datasets can contain arbitrary bits of metadata (e.g., annotation data) that might be of user interest in along with sequencing data.

[0018] In some embodiments, the plurality of components comprises the index to variant call data that provides a correspondence between respective ranges of the genome of the species to offsets in the data section where variant call data for the respective ranges is found.

[0019] In some embodiments, the plurality of components comprises the phase block track. The phase block track comprises (i) a dictionary and (ii) a track data section comprising phase information for one or more chromosomes in the genome of the species. In some embodiments, the dictionary comprises a plurality of names, and for each respective name in the plurality of names, an offset into the track data where records for the corresponding name

are found. In some embodiments, the track data section comprises a plurality of records and wherein each record in the plurality of records represents a phase block in the target nucleic acid. In some embodiments, the tract data section is in the JSON file format.

[0020] In some embodiments, each respective record in the plurality of records specifies (i) a chromosome number corresponding to the respective record, (ii) a position where the phase block starts on the chromosome, (iii) a position where the phase block ends, (iv) a unique name for the record, and (v) phasing information about the phase block.

[0021] In some embodiments, each respective record in the plurality of records is represented by a node in a plurality of nodes in a respective interval tree in a plurality of interval trees, and each interval tree in the plurality of interval trees represents a chromosome in a plurality of chromosomes for the species. In some such embodiments, a node in the plurality of nodes of a first interval tree in the plurality of interval trees stores a midpoint of the node, the midpoint of the node is a position of the midpoint, on the corresponding chromosome, of the phase block corresponding to the node, each respective node in the plurality of nodes of the first interval tree has a link to a left child node, which corresponds to the phase block immediately to the left of (i.e., numerically less than) the phase block represented by the respective node in the genome of the species, each respective node in the plurality of nodes of the first interval tree has a link to a right child node, which corresponds to the phase block immediately to the right of (i.e., numerically greater than) the phase block represented by the respective node in the genome of the species, each respective node in the plurality of nodes of the first interval tree has a sorted set of nodes that represent phase blocks that overlap the midpoint of the respective node sorted by left hand position of such phase block, and each respective node in the plurality of nodes of the first interval tree has a sorted set of nodes that represent phase blocks that overlap the midpoint of the respective node sorted by right hand position of such phase blocks. In some such embodiments, each respective node in the plurality of nodes of the first interval tree further includes a name, which is an offset in the track data section to the record in the plurality of records that contains phase information for the phase block corresponding to the respective node.

[0022] In some embodiments, the header further comprises the version of the dataset structure used by the nucleic acid sequencing dataset. In some embodiments, the plurality of components comprises the refseq index, and the refseq index comprises an index of a plurality of molecular variation identifiers that are called in the sample. In some such embodiments,

each respective molecular variation identifier in the plurality of molecular variation identifiers is dbSNP identifier.

[0023] In some embodiments, the plurality of components comprises the gene track. In such embodiments, the gene track comprises a plurality of genes and, for each respective gene in the plurality of genes, a number of single nucleotide polymorphisms in the respective gene.

[0024] Another aspect of the present disclosure provides a system for processing program output over a network connection using a local computer, where the local computer comprises one or more microprocessors, and a memory that stores one or more programs. The one or more programs use the one or more microprocessors to execute a method in accordance with a first operating system running on the local computer. In the method a first instance of a first program is invoked. Then, there is obtained through the first instance of the first program from a user, a login and a password to a user account on a remote computer. This is used to log the user into the user account on the remote computer automatically (using the login and the password provided by the first instance of the first program) across a network connection between the local computer and the remote computer. Responsive to successful login on the remote computer, there automatically sent, without human intervention, a second instance of the first program configured to auto-install on the remote computer upon transmission to the remote computer when the remote computer does not already have the first program available in the users account. Next, there is received from the remote computer a request to open a panel within the first instance of the first program. The panel is originated by the second instance of the first program running on the remote computer. The panel solicits input from the user for controlling the second instance of the first program. Responsive to receiving input from the user for controlling the second instance of the first program in the panel on the local computer, the input is sent to the second instance of the first program on the remote computer across the network connection (e.g., wireless or wired connection). Next, there is received, from the remote computer across the network connection, output from the second instance of the first program responsive to the input. This output is displayed at the local computer.

[0025] Another aspect of the present disclosure provides a system for viewing nucleic acid sequencing data. The system comprises one or more microprocessors and a memory. The memory stores one or more programs that use the one or more microprocessors to obtain a nucleic acid sequencing dataset corresponding to at least one target nucleic acid in a sample. The nucleic acid sequencing dataset comprises a plurality of sequencing reads from the sample. Each respective sequencing read in the plurality of sequencing reads comprises a first portion

that corresponds to a subset of at least one target nucleic acid in the sample and a second portion that encodes a respective identifier (e.g., bar code) for the respective sequencing read in a plurality of identifiers. Each respective identifier is independent of the sequence of the at least one target nucleic acid. The plurality of sequencing reads collectively includes the plurality of identifiers. A visualization tool is displayed. A request is obtained from a user through the visualization tool. The request specifies a genomic region represented by the nucleic acid sequencing dataset. Responsive to obtaining the request, the request is parsed by obtaining a plurality of sequencing reads within the genomic region from the nucleic acid sequencing dataset. A scan window is run against the plurality of sequencing reads thereby creating a plurality of windows, each respective window of the plurality of windows corresponding to a different region of the genomic region and including an identity of each identifier of each sequencing read in the different region of the genomic region in the nucleic acid sequencing dataset. A two dimensional heat map that represents each possible window pair in the plurality of windows is displayed. Each respective window pair is displayed in the two dimensional heat map as a color selected from a color scheme based upon the number of identifiers in common in the respective window pair.

[0026] Various embodiments of systems, methods and devices within the scope of the appended claims each have several aspects, no single one of which is solely responsible for the desirable attributes described herein. Without limiting the scope of the appended claims, some prominent features are described herein. After considering this discussion, and particularly after reading the section entitled “Detailed Description” one will understand how the features of various embodiments are used.

INCORPORATION BY REFERENCE

[0027] All publications, patents, and patent applications mentioned in this specification are herein incorporated by reference in their entireties to the same extent as if each individual publication, patent, or patent application was specifically and individually indicated to be incorporated by reference.

BRIEF DESCRIPTION OF THE DRAWINGS

[0028] The implementations disclosed herein are illustrated by way of example, and not by way of limitation, in the figures of the accompanying drawings. Like reference numerals refer to corresponding parts throughout the drawings.

[0029] Figure 1 is an example block diagram illustrating a computing device in accordance with some implementations.

[0030] Figure 2 illustrates exemplary constructs in accordance with an embodiment of the present disclosure.

[0031] Figure 3 provides an overview of a nucleic acid sequencing dataset in accordance with an embodiment of the present disclosure.

[0032] Figure 4 illustrates the data structure of an example phase block track within a nucleic acid sequencing dataset in accordance with some embodiments.

[0033] Figure 5 illustrates an example phase block track in accordance with some embodiments.

[0034] Figure 6 illustrates the data structure of an example gene track in accordance with some embodiments.

[0035] Figure 7 illustrates an example gene track in accordance with some embodiments.

[0036] Figure 8 illustrates the data structure of an example structural variant dataset track within a nucleic acid sequencing dataset in accordance with some embodiments.

[0037] Figure 9 illustrates an example structural variant dataset track in accordance with some embodiments.

[0038] Figure 10 illustrates target, fragment and sequence read data within a nucleic acid sequencing dataset in accordance with some embodiments.

[0039] Figure 11 illustrates variant call data within a nucleic acid sequencing dataset in accordance with some embodiments.

[0040] Figure 12 illustrates a summarization module in a haplotype visualization tool in accordance with some embodiments.

[0041] Figure 13 illustrates a summarization module in a haplotype visualization tool in accordance with additional embodiments.

[0042] Figure 14A illustrates a screen shot of a phase visualization module in a haplotype visualization tool in accordance with some embodiments.

[0043] Figure 14B illustrates another screen shot of a phase visualization module in a haplotype visualization tool in accordance with some embodiments.

[0044] Figure 15 illustrates another screen shot of a phase visualization module in a haplotype visualization tool in accordance with some embodiments.

[0045] Figure 16 illustrates another screen shot of a phase visualization module in a haplotype visualization tool in accordance with some embodiments.

[0046] Figure 17 illustrates search function features of a haplotype visualization tool in accordance with some embodiments.

[0047] Figure 18 illustrates a screen shot of a structural variants module in a haplotype visualization tool in accordance with some embodiments.

[0048] Figure 19 illustrates another screen shot of a structural variants module in a haplotype visualization tool in accordance with some embodiments.

[0049] Figure 20 illustrates still another screen shot of a structural variants module in a haplotype visualization tool in accordance with some embodiments.

[0050] Figure 21 illustrates still an additional screen shot of a structural variants module in a haplotype visualization tool in accordance with some embodiments.

[0051] Figure 22 illustrates a screen shot of a read visualization module in a haplotype visualization tool in accordance with some embodiments.

[0052] Figure 23 illustrates another screen shot of a structural variants module in a haplotype visualization tool in accordance with some embodiments.

[0053] Figure 24 illustrates another screen shot of a structural variants module in a haplotype visualization tool in accordance with some embodiments.

[0054] Figure 25 illustrates another screen shot of a structural variants module in a haplotype visualization tool in accordance with some embodiments.

[0055] Figure 26 illustrates a phase visualization module in a haplotype visualization tool in accordance with some embodiments.

[0056] Figure 27 illustrates another aspect of a phase visualization module in a haplotype visualization tool in accordance with some embodiments.

[0057] Figure 28A illustrates another aspect of a phase visualization module in a haplotype visualization tool in accordance with some embodiments.

[0058] Figure 28B illustrates still another aspect of a phase visualization module in a haplotype visualization tool in accordance with some embodiments.

[0059] Figure 29 illustrates another aspect of a phase visualization module in a haplotype visualization tool in accordance with some embodiments.

[0060] Figure 30 illustrates another aspect of a phase visualization module in a haplotype visualization tool in accordance with some embodiments.

[0061] Figure 31 is an example block diagram illustrating a computing system in accordance with some implementations.

[0062] Figure 32 is an example of a credential challenge for remote initiation of an instance of a haplotype visualization tool in accordance with the disclosed embodiments.

[0063] Figure 33 illustrates a structural variants module in a haplotype visualization tool in accordance with some embodiments in which a sequence read filter is turned off.

[0064] Figure 34 illustrates a structural variants module in a haplotype visualization tool in accordance with some embodiments in which a sequence read filter is turned on.

DETAILED DESCRIPTION

[0065] Reference will now be made in detail to embodiments, examples of which are illustrated in the accompanying drawings. In the following detailed description, numerous specific details are set forth in order to provide a thorough understanding of the present disclosure. However, it will be apparent to one of ordinary skill in the art that the present disclosure may be practiced without these specific details. In other instances, well-known methods, procedures, components, circuits, and networks have not been described in detail so as not to unnecessarily obscure aspects of the embodiments.

[0066] It will also be understood that, although the terms first, second, etc. may be used herein to describe various elements, these elements should not be limited by these terms. These terms are only used to distinguish one element from another. For example, a first subject could be termed a second subject, and, similarly, a second subject could be termed a first subject, without departing from the scope of the present disclosure. The first subject and the second subject are both subjects, but they are not the same subject.

[0067] The terminology used in the present disclosure is for the purpose of describing particular embodiments only and is not intended to be limiting of the invention. As used in the description of the invention and the appended claims, the singular forms “a”, “an” and “the” are intended to include the plural forms as well, unless the context clearly indicates otherwise. It will also be understood that the term “and/or” as used herein refers to and encompasses any and all possible combinations of one or more of the associated listed items. It will be further understood that the terms “comprises” and/or “comprising,” when used in this specification, specify the presence of stated features, integers, steps, operations, elements, and/or components, but do not preclude the presence or addition of one or more other features, integers, steps, operations, elements, components, and/or groups thereof.

[0068] As used herein, the term “if” may be construed to mean “when” or “upon” or “in response to determining” or “in response to detecting,” depending on the context. Similarly, the phrase “if it is determined” or “if [a stated condition or event] is detected” may be construed to mean “upon determining” or “in response to determining” or “upon detecting (the stated condition or event)” or “in response to detecting (the stated condition or event),” depending on the context.

[0069] The implementations described herein provide various technical solutions to detect a structural variant (*e.g.*, deletions, duplications, copy-number variants, insertions, inversions, translocations, long terminal repeats (LTRs), short tandem repeats (STRs), and a variety of other useful characterizations) in sequencing data of a test nucleic acid obtained from a biological sample. Details of implementations are now described in relation to the Figures.

[0070] Figure 1 is a block diagram illustrating a structural variant and phasing visualization system 100 in accordance with some implementations. The device 100 in some implementations includes one or more processing units CPU(s) 102 (also referred to as processors), one or more network interfaces 104, a user interface 106, a memory 112, and one or more communication buses 114 for interconnecting these components. The communication

buses 114 optionally include circuitry (sometimes called a chipset) that interconnects and controls communications between system components. The memory 112 typically includes high-speed random access memory, such as DRAM, SRAM, DDR RAM, ROM, EEPROM, flash memory, CD-ROM, digital versatile disks (DVD) or other optical storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, other random access solid state memory devices, or any other medium which can be used to store desired information; and optionally includes non-volatile memory, such as one or more magnetic disk storage devices, optical disk storage devices, flash memory devices, or other non-volatile solid state storage devices. The memory 112 optionally includes one or more storage devices remotely located from the CPU(s) 102. The memory 112, or alternatively the non-volatile memory device(s) within the memory 112, comprises a non-transitory computer readable storage medium. In some implementations, the memory 112 or alternatively the non-transitory computer readable storage medium stores the following programs, modules and data structures, or a subset thereof:

- an optional operating system 116, which includes procedures for handling various basic system services and for performing hardware dependent tasks;
- an optional network communication module (or instructions) 118 for connecting the device 100 with other devices, or a communication network;
- an optional sequencing read processing module 120 for processing sequencing reads, including a structural variation determination sub-module 120 for identifying structural variations in a genetic sample from a single organism of a species and a phasing sub-module 124 for identifying the haplotype of each sequencing read of the genetic sample;
- one or more nucleic acid sequencing datasets 126, each such dataset obtained using a genetic sample from a single organism of a species;
- gene annotation data, optionally in the form of a gene track interval tree 128;
- exon annotation data, optionally in the form of an exon track interval tree 142;
- one or more additional sources of annotation data, optionally in the form of interval trees 146;
- a haplotype visualization tool 148 for visualizing structural variation and phasing information in nucleic acid sequencing data, including any combination of one or more

of a summarization module 150, a phase visualization module 152, a structural variants (visualization) module 154, and a read visualization module 156.

[0071] In some implementations, the user interface 106 includes an input device (e.g., a keyboard, a mouse, a touchpad, a track pad, and/or a touch screen) 100 for a user to interact with the system 100 and a display 108.

[0072] In some implementations, one or more of the above identified elements are stored in one or more of the previously mentioned memory devices, and correspond to a set of instructions for performing a function described above. The above identified modules or programs (e.g., sets of instructions) need not be implemented as separate software programs, procedures or modules, and thus various subsets of these modules may be combined or otherwise re-arranged in various implementations. In some implementations, the memory 112 optionally stores a subset of the modules and data structures identified above. Furthermore, in some embodiments, the memory stores additional modules and data structures not described above. In some embodiments, one or more of the above identified elements is stored in a computer system, other than that of system 100, that is addressable by system 100 so that system 100 may retrieve all or a portion of such data when needed.

[0073] Although Figure 1 shows a “structural variation and phasing visualization system 100,” the figure is intended more as functional description of the various features which may be present in computer systems than as a structural schematic of the implementations described herein. In practice, and as recognized by those of ordinary skill in the art, items shown separately could be combined and some items could be separated.

[0074] Advantageously, because the nucleic acid sequence datasets 126 are large in typical embodiments (e.g., 1 gigabyte or greater, 5 gigabytes or greater, or 10 gigabytes or greater), in some embodiments the structural variation and phasing visualization system 100 is part of a system that includes one or more client devices 3102 that are in electronic communication with the structural variation and phasing visualization system 100 of Figure 1 across a communication network 3106. Such a network topology allows scientists and other users to use one of several network based technologies to run the haplotype visualization tool 148 on system 100, which in typical embodiments is a powerful server computer, but view the results on client device 3102, which can be, for example, a laptop computer. Any form of network technology for implementing this network topology is encompassed within the present disclosure. For instance X-windows session forwarding (not shown in Figure 31) is used in

some embodiments. In other embodiments, the Internet (web) is used. In particular, a browser application is run on the client device 3102.

[0075] The process of running a program on a remote computer (e.g., in system 3100, the structural variation and phasing visualization system 100 is considered remote) and viewing the results on a client device 3102 (e.g., desktop or laptop) is cumbersome. A user must generally (i) install certain parts of the program on their computer 3102 and other parts on the server 100, (ii) use SSH or firewall software to create an open network port connecting the two computers (system 3102 to client device 100), and (iii) independently start different parts of the program on different systems. For example, the URL blog.tracks.com/2014/05/17/ssh-tunnel-local-and-remote-port-forwardingexplained-with-examples.html, which is hereby incorporated by reference, explains one way of setting up forwarding. As another example, the URL itg.chem.indiana.edu/inc/wiki/software/openssh/200.html explains another way of setting up forwarding. The present disclosure incorporates such techniques. However, advantageously, in some embodiments, the present disclosure affords solutions to the above-disclosed networking techniques, which seeks to automate and improve upon the processes described above. Once a user has installed the haplotype visualization tool 148 on their client device 3102, they only need to provide the tool 148 with their credentials (e.g., user-name and password) for the remote computer (structural variation and phasing visualization system 100) that has the data and computational facilities to run the haplotype visualization tool 148. For instance, in some embodiments, referring to Figure 32, the user running the haplotype visualization tool 148 on client 3102 will be provided with the challenge 3200 that includes a query for the server name or address 3204, the user's name 3206, an optional SSH key file (to enable encrypted connection) 3208, an optional SSH key password 3210, and a work location 3212 on the server. The instance of the haplotype visualization tool 148 on their client device 3102 then connects to the remote computer 100 and authenticates as the user using the provided credentials. Using that connection, it installs the haplotype visualization tool 148 on the remote computer, starts it, and configures any necessary network port forwarding. Once the haplotype visualization tool has done this, it opens up a new window on the client device 3102 that is "connected" to the haplotype visualization tool running on the remote structural variation and phasing visualization system. Of particular note, in such embodiments, the haplotype visualization tool 148 on the client device 3102 includes in a copy of itself that is intended to run on the structural variation and phasing visualization system 100. In some

embodiments, the structural variation and phasing visualization system 100 is running a first operating system and the client device 3102 is running a second operating system. In some embodiments, the first operating system and the second operating system are the same. In some embodiments, the first operating system and the second operating system are different. In some embodiments, the first operating system is one of iOS, DARWIN, RTXC, LINUX, UNIX, OS X, or WINDOWS, and the second operating system is other than the first operating system and one of iOS, DARWIN, RTXC, LINUX, UNIX, OS X, or WINDOWS. In the disclosed embodiment, the haplotype visualization tool 148 running on the client device 3102 copies the archived copy of the haplotype visualization tool 148 to the structural variation and phasing system 100 and installs (if it has not been installed before) during the setup process. It will be appreciated that the system and method disclosed for remote initiation of the haplotype visualization tool 148 on a remote computer is applicable to a broad range of applications that require the computational resources of a remote server with the concomitant visual interface operating on a local computer in order to control such applications and to visualize data and computational results in real time or near real time.

[0076] Referring once again to Figures 1, 31, and 32, one aspect of the present disclosure provides a system 3100 for processing program output over a network connection 3106 (e.g., wired or wireless) using a local computer 3102. The local computer 3102 comprises one or more microprocessors (not shown), and a memory (not shown) that stores one or more programs (e.g., haplotype visualization tool 148). The one or more programs use the one or more microprocessors to execute a method in accordance with a first operating system running on the local computer. In the method, a first instance of a first program is invoked (e.g., a first instance of the haplotype visualization tool 148 is invoked on a client device 3102). Through the invoked first instance of the first program there is obtained, from a user, a login and a password to a user account on a remote computer (e.g., structural variation and phasing visualization system 100). The user is then logged into the user account on the remote computer automatically, using the login and the password provided by the first instance of the first program, across a network connection between the local computer and the remote computer (e.g., communication network 3106). Responsive to successful login on the remote computer 100, the method continues by automatically sending, without human intervention, a second instance of the first program 148 configured to auto-install on the remote computer 100 upon transmission to the remote computer. In some embodiments, the remote computer already has the second instance of the first program 148 installed and in some such

embodiments the second instance of the first program is therefore not transmitted to the remote computer for installation. Once the second instance of the first program is installed on the remote computer 100, there is received from the remote computer a request to open a panel (not shown). This panel is originated by the second instance of the first program running on the remote computer 100. The panel solicits input from the user for controlling the second instance of the first program. For instance, in some embodiments this panel is of the form illustrated in any one of Figure 12-21. In some embodiments, the panel is simpler, for instance containing a prompt for a dataset name or a search query for searching in a specified dataset. Responsive to receiving input from the user for controlling the second instance of the first program in the panel on the local computer, the input is sent to the second instance of the first program running on the remote computer 100 across the network connection. The remote computer receives across the network connection this input and, subsequently, output from the second instance of the first program responsive to the input is displayed as output on the local computer (e.g. within the first instance of the first program or in a separate web browser).

[0077] Referring to Figure 2, in accordance with the disclosed systems and methods, a plurality of sequencing reads (not shown in its entirety in Figure 2) is obtained using a test (target) nucleic acid 206 of a biological sample from a subject. In typical embodiments, the test (target) nucleic acid 206 is a fragment of the genome of the biological sample. In some embodiments, there is a single test (target) nucleic acid 206 (fragment) in a partition. In some embodiments, there are two or more test nucleic acids 206 (fragments) in a partition each corresponding to different portions of the genome of the species of the biological sample. In some embodiments, there are five or more nucleic acids 206 (fragments) in a partition each corresponding to different portions of the genome of the species of the biological sample. In some embodiments, there are ten or more nucleic acids 206 in a partition each corresponding to different portions of the genome of the species of the biological sample. In some embodiments, the biological sample is a mixture and includes nucleic data representing the genome of two or more individuals in a species. In some embodiments, the biological sample is a mixture and includes nucleic data representing the genome of two or more species. For instance, in some embodiments the biological sample is infected with a retrovirus. In another example, the biological sample contains metagenomes because the sample was taken from sand or dirt or some other location and the goal is to find all the different genomes that exist in the sample.

[0078] The sequencing reads ultimately form the basis of a nucleic acid sequencing dataset 126. Each respective sequencing read 202 in the plurality of sequencing reads

comprises a first portion that corresponds to a subset of a test nucleic acid and a second portion that encodes identification information for the respective sequencing read. The identification information is independent of the sequencing data of the test nucleic acid.

[0079] In some embodiments, sequencing read lengths have an N50 (where the sum of the sequence read lengths that are greater than the stated N50 number is 50% of the sum of all sequencing read lengths). In typical embodiments, sequencing reads are tens or hundreds of bases in length, which in turn, are aligned to form constructs of at least about 10kb, at least about 20kb, or at least about 50kb. In more preferred aspects, sequencing reads are tens or hundreds of bases in length, which in turn, are aligned to form constructs having at least about 100kb, at least about 150kb, at least about 200kb, and in many cases, at least about 250kb, at least about 300 kb, at least about 350 kb, at least about 400 kb, and in some cases, at least about 500 kb or more.

[0080] In some embodiments, to obtain the plurality of sequencing reads from a biological sample from a subject, a test nucleic acid 206 is fragmented and these fragments are compartmentalized, or partitioned into discrete compartments or partitions (referred to interchangeably herein as partitions). In some embodiments, the test nucleic acid is the genome of a multi-chromosomal organism such as a human. In typical embodiments, multiple sequencing reads are measured from each such compartment or partition with lengths that are tens or hundreds of bases in length. Sequencing reads from the same compartment or partition that have the same bar code can be aligned to form sequence constructs that are at least about 25kb, at least about 50kb, 100kb, at least about 150kb, at least about 200kb, and in many cases, at least about 250kb, at least about 300 kb, at least about 350 kb, at least about 400 kb, and in some cases, at least about 500 kb or more in length.

[0081] Each partition maintains separation of its own contents from the contents of other partitions. As used herein, the partitions refer to containers or vessels that may include a variety of different forms, e.g., wells, tubes, micro or nanowells, through holes, or the like. In preferred aspects, however, the partitions are flowable within fluid streams. In some embodiments, these vessels are comprised of, e.g., microcapsules or micro-vesicles that have an outer barrier surrounding an inner fluid center or core, or have a porous matrix that is capable of entraining and/or retaining materials within its matrix. In a preferred aspect, however, these partitions comprise droplets of aqueous fluid within a non-aqueous continuous phase, e.g., an oil phase. A variety of different vessels are described in, for example, U.S. Patent Application No. 13/966,150, filed August 13, 2013, which is hereby incorporated by

reference herein in its entirety. Likewise, emulsion systems for creating stable droplets in non-aqueous or oil continuous phases are described in detail in, e.g., Published U.S. Patent Application No. 2010-0105112, which is hereby incorporated by reference herein in its entirety. In certain embodiments, microfluidic channel networks are particularly suited for generating partitions as described herein. Examples of such microfluidic devices include those described in detail in Provisional U.S. Patent Application No. 61/977,804, filed April 4, 2014, as well as PCT/US15/025197, the full disclosures of which are incorporated herein by reference in their entirety for all purposes. Alternative mechanisms may also be employed in the partitioning of individual cells, including porous membranes through which aqueous mixtures of cells are extruded into non-aqueous fluids. Such systems are generally available from, e.g., NANOMI, Inc.

[0082] In the case of droplets in an emulsion, partitioning of the test nucleic acid fragments into discrete partitions may generally be accomplished by flowing an aqueous, sample containing stream, into a junction into which is also flowing a non-aqueous stream of partitioning fluid, e.g., a fluorinated oil, such that aqueous droplets are created within the flowing stream partitioning fluid, where such droplets include the sample materials. As described below, the partitions, e.g., droplets, also typically include co-partitioned barcode oligonucleotides.

[0083] The relative amount of sample materials within any particular partition may be adjusted by controlling a variety of different parameters of the system, including, for example, the concentration of test nucleic acid fragments in the aqueous stream, the flow rate of the aqueous stream and/or the non-aqueous stream, and the like. The partitions described herein are often characterized by having overall volumes that are less than 1000 pL, less than 900 pL, less than 800 pL, less than 700 pL, less than 600 pL, less than 500 pL, less than 400pL, less than 300 pL, less than 200 pL, less than 100pL, less than 50 pL, less than 20 pL, less than 10 pL, or even less than 1 pL. Where co-partitioned with beads, it will be appreciated that the sample fluid volume within the partitions may be less than 90% of the above described volumes, less than 80%, less than 70%, less than 60%, less than 50%, less than 40%, less than 30%, less than 20%, or even less than 10% the above described volumes. In some cases, the use of low reaction volume partitions is particularly advantageous in performing reactions with very small amounts of starting reagents, e.g., input test nucleic acid fragments. Methods and systems for analyzing samples with low input nucleic acids are presented in U.S. Provisional

Patent Application No. 62/017,580 June 26, 2014, the full disclosure of which is hereby incorporated by reference in its entirety.

[0084] Once the test nucleic acid fragments are introduced into their respective partitions, the test nucleic acid fragments within partitions are generally provided with unique identifiers such that, upon characterization of those test nucleic acid fragments, they may be attributed as having been derived from their respective partitions. Such unique identifiers may be previously, subsequently or concurrently delivered to the partitions that hold the compartmentalized or partitioned test nucleic acid fragments, in order to allow for the later attribution of the characteristics, e.g., nucleic acid sequence information, to the sample nucleic acids included within a particular compartment, and particularly to relatively long stretches of contiguous sample nucleic acids that may be originally deposited into the partitions.

[0085] Accordingly, the test nucleic acid fragments are typically co-partitioned with the unique identifiers (e.g., barcode sequences). In particularly preferred aspects, the unique identifiers are provided in the form of oligonucleotides that comprise nucleic acid barcode sequences that is attached to test nucleic acid fragments in the partitions. The oligonucleotides are partitioned such that as between oligonucleotides in a given partition, the nucleic acid barcode sequences contained therein are the same, but as between different partitions, the oligonucleotides can, and preferably have differing barcode sequences. In some embodiments, only one nucleic acid barcode sequence is associated with a given partition, although in some embodiments, two or more different barcode sequences are present in a given partition.

[0086] The nucleic acid barcode sequences will typically include from 6 to about 20 or more nucleotides within the sequence of the oligonucleotides. These nucleotides may be completely contiguous, i.e., in a single stretch of adjacent nucleotides, or they may be separated into two or more separate subsequences that are separated by one or more nucleotides. Typically, separated subsequences may typically be from about 4 to about 16 nucleotides in length.

[0087] The test nucleic acid is typically partitioned such that the nucleic acids are present in the partitions in relatively long fragments or stretches of contiguous nucleic acid molecules. These fragments typically represent a number of overlapping fragments of the overall test nucleic acid to be analyzed, e.g., an entire chromosome, exome, or other large genomic fragment. This test nucleic acid may include whole genomes, individual chromosomes, exomes, amplicons, or any of a variety of different nucleic acids of interest.

Typically, the fragments of the test nucleic acid that are partitioned are longer than 1 kb, longer than 5 kb, longer than 10 kb, longer than 15 kb, longer than 20 kb, longer than 30 kb, longer than 40 kb, longer than 50 kb, longer than 60 kb, longer than 70 kb, longer than 80 kb, longer than 90 kb or even longer than 100 kb.

[0088] The test nucleic acid is also typically partitioned at a level whereby a given partition has a very low probability of including two overlapping fragments of the starting test nucleic acid. This is typically accomplished by providing the test nucleic acid at a low input amount and/or concentration during the partitioning process. As a result, in preferred cases, a given partition includes a number of long, but non-overlapping fragments of the starting test nucleic acid. The nucleic acid fragments in the different partitions are then associated with unique identifiers, where for any given partition, nucleic acids contained therein possess the same unique identifier, but where different partitions include different unique identifiers. Moreover, because the partitioning step allocates the sample components into very small volume partitions or droplets, it will be appreciated that in order to achieve the desired allocation as set forth above, one need not conduct substantial dilution of the sample, as would be required in higher volume processes, e.g., in tubes, or wells of a multiwell plate. Further, because the systems described herein employ such high levels of barcode diversity, one can allocate diverse barcodes among higher numbers of genomic equivalents, as provided above. In some embodiments, in excess of 10,000, 100,000, 500,000, etc. diverse barcode types are used to achieve genome:(barcode type) ratios that are on the order of 1:50 or less, 1:100 or less, 1:1000 or less, or even smaller ratios, while also allowing for loading higher numbers of genomes (e.g., on the order of greater than 100 genomes per assay, greater than 500 genomes per assay, 1000 genomes per assay, or even more) while still providing for far improved barcode diversity per genome. Here, each such genome is an example of a test nucleic acid.

[0089] Referring to Figure 2, panels A and B, often the above-described partitioning is performed by combining the sample containing the test nucleic acid with a set of oligonucleotide tags (containing the barcodes) that are releasably-attached to beads 308 prior to the partitioning step. The oligonucleotides may comprise at least a primer region 216 and a barcode 214 region. Between oligonucleotides within a given partition, the barcode region 214 is substantially the same barcode sequence, but as between different partitions, the barcode region in most cases is a different barcode sequence. In some embodiments, the primer region 216 is an N-mer (either a random N-mer or an N-mer designed to target a particular sequence) that is used to prime the nucleic acids within the sample within the partitions. In some cases,

where the N-mer is designed to target a particular sequence, the primer region 216 is designed to target a particular chromosome (e.g., human chromosome 1, 13, 18, or 21), or region of a chromosome, e.g., an exon or other targeted region. In some cases, the N-mer is designed to target a particular gene or genetic region, such as a gene or region associated with a disease or disorder (e.g., cancer). In some cases, the N-mer is designed to target a particular structural variation. Within the partitions, an amplification reaction is conducted using the primer sequence 216 (e.g. N-mer) to prime the nucleic acid sample at different places along the length of the nucleic acid. As a result of the amplification, each partition contains amplified products of the nucleic acid 202 that are attached to an identical or near-identical barcode, and that represent overlapping, smaller fragments of the nucleic acids in each partition. The barcode 214 therefore serves as a marker that signifies that a set of nucleic acids originated from the same partition, and thus potentially also originated from the same strand of test nucleic acid. Following amplification, the nucleic acids are pooled, sequenced, and aligned using a sequencing algorithm. Because shorter sequence reads may, by virtue of their associated barcode sequences, be aligned and attributed to a single, long fragment of the test nucleic acid, all of the identified variants on that sequence can be attributed to a single originating fragment and single originating chromosome of the test nucleic acid. Further, by aligning multiple co-located variants across multiple long fragments, one can further characterize that chromosomal contribution. Accordingly, conclusions regarding the phasing of particular genetic variants may then be drawn. Such information may be useful for identifying haplotypes, which are generally a specified set of genetic variants that reside on the same nucleic acid strand or on different nucleic acid strands. Moreover, additionally or alternatively, structural variants are identified.

[0090] In some embodiments, the co-partitioned oligonucleotides also comprise functional sequences in addition to the barcode region 214 and the primer region 216 region of the nucleic acids within the sample within the partitions. See, for example, the disclosure on co-partitioning of oligonucleotides and associated barcodes and other functional sequences, along with sample materials as described in, for example, U.S. Patent Application Nos. 61/940,318, filed February 7, 2014, 61/991,018, Filed May 9, 2014, and U.S. Patent Application No. 14/316,383, (Attorney Docket No. 43487-708.201) filed on June 26, 2014, as well as U.S. Patent Application No. 14/175,935, filed February 7, 2014, the full disclosures of which is hereby incorporated by reference in their entireties.

[0091] In one exemplary process, beads are provided, where each such bead includes large numbers of the above described oligonucleotides releasably attached to the beads. In such embodiments, all of the oligonucleotides attached to a particular bead include the same nucleic acid barcode sequence, but a large number of diverse barcode sequences are represented across the population of beads used. Typically, the population of beads provides a diverse barcode sequence library that includes at least 1000 different barcode sequences, at least 10,000 different barcode sequences, at least 100,000 different barcode sequences, or in some cases, at least 1,000,000 different barcode sequences. Additionally, each bead typically is provided with large numbers of oligonucleotide molecules attached. In particular, the number of molecules of oligonucleotides including the barcode sequence on an individual bead may be at least about 10,000 oligonucleotides, at least 100,000 oligonucleotide molecules, at least 1,000,000 oligonucleotide molecules, at least 100,000,000 oligonucleotide molecules, and in some cases at least 1 billion oligonucleotide molecules.

[0092] In some embodiments, the oligonucleotides are releasable from the beads upon the application of a particular stimulus to the beads. In some cases, the stimulus may be a photo-stimulus, e.g., through cleavage of a photo-labile linkage that may release the oligonucleotides. In some cases, a thermal stimulus may be used, where elevation of the temperature of the beads environment may result in cleavage of a linkage or other release of the oligonucleotides from the beads. In some cases, a chemical stimulus may be used that cleaves a linkage of the oligonucleotides to the beads, or otherwise may result in release of the oligonucleotides from the beads.

[0093] In accordance with the methods and systems described herein, the beads including the attached oligonucleotides may be co-partitioned with the individual samples, such that a single bead and a single sample are contained within an individual partition. In some cases, where single bead partitions are desired, it may be desirable to control the relative flow rates of the fluids such that, on average, the partitions contain less than one bead per partition, in order to ensure that those partitions that are occupied, are primarily singly occupied. Likewise, one may wish to control the flow rate to provide that a higher percentage of partitions are occupied, e.g., allowing for only a small percentage of unoccupied partitions. In preferred aspects, the flows and channel architectures are controlled as to ensure a desired number of singly occupied partitions, less than a certain level of unoccupied partitions and less than a certain level of multiply occupied partitions.

[0094] Figure 3 of United States Patent Application No. 62/072,214, filed October 29, 2014, entitled “Analysis of Nucleic Acid Sequences,” which is hereby incorporated by reference and the portions of the specification therein describing Figure 3 provide a detailed example of one method for barcoding and subsequently sequencing a test nucleic acid (referred to in the reference as a “sample nucleic acid”) in accordance with one embodiment of the present disclosure. As noted above, while single bead occupancy may be the most desired state, it will be appreciated that multiply occupied partitions, or unoccupied partitions may often be present. Figure 4 of United States Patent Application No. 62/072,214, filed October 29, 2014, entitled “Analysis of Nucleic Acid Sequences,” which is hereby incorporated by reference and the portions of the specification describing Figure 4 therein provide a detailed example of a microfluidic channel structure for co-partitioning samples and beads comprising barcode oligonucleotides in accordance with one embodiment of the present disclosure.

[0095] Once co-partitioned, the oligonucleotides disposed upon the beads may be used to barcode and amplify the partitioned samples. One process for use of these barcode oligonucleotides in amplifying and barcoding samples is described in detail in U.S. Patent Application Nos. 61/940,318, filed February 7, 2014, 61/991,018, Filed May 9, 2014, and U.S. Patent Application No. 14/316,383, (Attorney Docket No. 43487-708.201) filed on June 26, 2014, the full disclosures of which are hereby incorporated by reference in their entireties. Briefly, in one aspect, the oligonucleotides present on the beads that are co-partitioned with the samples are released from their beads into the partition with the samples. The oligonucleotides typically include, along with the barcode sequence, a primer sequence at its 5' end. This primer sequence may be a random oligonucleotide sequence intended to randomly prime numerous different regions of the samples, or it may be a specific primer sequence targeted to prime upstream of a specific targeted region of the sample.

[0096] Once released, the primer portion of the oligonucleotide can anneal to a complementary region of the sample. Extension reaction reagents, e.g., DNA polymerase, nucleoside triphosphates, co-factors (e.g., Mg^{2+} or Mn^{2+} etc.), that are also co-partitioned with the samples and beads, then extend the primer sequence using the sample as a template, to produce a complementary fragment to the strand of the template to which the primer annealed, with complementary fragment that includes the oligonucleotide and its associated barcode sequence. Annealing and extension of multiple primers to different portions of the sample may result in a large pool of overlapping complementary fragments of the sample, each possessing its own barcode sequence indicative of the partition in which it was created. In some cases,

these complementary fragments may themselves be used as a template primed by the oligonucleotides present in the partition to produce a complement of the complement that again, includes the barcode sequence. In some cases, this replication process is configured such that when the first complement is duplicated, it produces two complementary sequences at or near its termini, to allow the formation of a hairpin structure or partial hairpin structure that reduces the ability of the molecule to be the basis for producing further iterative copies. A schematic illustration of one example of this is shown in Figure 2.

[0097] As Figure 2 shows, oligonucleotides 202 that include a barcode sequence 214 are co-partitioned in, e.g., a droplet 204 in an emulsion, along with a sample test nucleic acid fragment 206. In some embodiments, the oligonucleotides 202 are provided on a bead 208 that is co-partitioned with the test nucleic acid fragment 206, which oligonucleotides are preferably releasable from the bead 208, as shown in Figure 2, panel (A). As shown in Figure 2 panel (B), the oligonucleotides 202 includes a barcode sequence 214, in addition to one or more functional sequences, e.g., sequences 212, 214 and 216. For example, oligonucleotide 202 is shown as further comprising sequence 212 that may function as an attachment or immobilization sequence for a given sequencing system, e.g., a P5 sequence used for attachment in flow cells of an ILLUMINA, HISEQ or MISEQ system. In other words, attachment sequence 212 is used to reversibly attach oligonucleotide 202 to a bead 208 in some embodiments. As shown in Figure 2, panel B, the oligonucleotide 202 also includes a primer sequence 216, which may include a random or targeted N-mer (discussed above) for priming replication of portions of the sample test nucleic acid fragment 206. Also included within exemplary oligonucleotide 202 of Figure 2, panel B, is a sequence 210 which may provide a sequencing priming region, such as a “read1” or R1 priming region, that is used to prime polymerase mediated, template directed sequencing by synthesis reactions in sequencing systems. In many cases, the barcode sequence 214, immobilization (attachment) sequence 212 and exemplary R1 sequence 214 may be common to all of the oligonucleotides 202 attached to a given bead. The primer sequence 216 may vary for random N-mer primers, or may be common to the oligonucleotides on a given bead for certain targeted applications. Figures 3B through 3E and the specification describing these Figures in United States Prov. Application No. 62/113,693, entitled “Systems and Methods for Determining Structural Variation,” filed February 9, 2014 detail how oligonucleotides 202 form sequencing reads of the sample test nucleic acid, where each such sequencing read includes a first portion that is a sequencing read of the sample test nucleic acid and a second portion that is the oligonucleotide 202. Such

sequencing reads, and analysis of such sequencing reads, form the basis of the disclosed nucleic acid sequencing dataset 126.

[0098] In some embodiments, the sequencing reads in a nucleic acid sequencing dataset 126 are processed in order to sequence the at least one target nucleic acid. In some embodiments conventional methods are used to process the nucleic acid sequence reads in order to establish a sequence for the at least one target nucleic acid. In some embodiments the novel methods disclosed in PCT application PCT/US2015/038175, entitled “Processes and Systems for Nucleic Acid Sequence Assembly,” filed June 26, 2015, which is hereby incorporated by reference, are used to process the nucleic acid sequence reads in order to establish a sequence for the at least one target nucleic acid. In some embodiments, such sequencing involves mapping the sequencing reads to a reference genome, such as the genome of the species from which the sample is taken. In some embodiments, the sample is expected, or suspected, of containing multiple genomes (e.g., the case in which a sample, such as a human sample, infected with a retrovirus). In such cases, multiple reference genomes, from different species may be concurrently used.

[0099] In some embodiments, the sequencing reads are processed by phasing them and by looking for structural variations. In some embodiments, conventional phasing methods and structural variation methods are used. In some embodiments, novel phasing methods and structural variation methods, such as those disclosed in United States Provisional Application No. 62,238,077, entitled “Systems and Method for Determining Structural Variation Using Probabilistic Models,” filed October 6, 2015, which is hereby incorporated by reference, are used. Although not disclosed in this reference, in some embodiments the teachings of the reference are extended to incorporate multiple reference genomes in instances where the sample potential contains nucleic acid from multiple reference genomes. For instance, in the case where the sample is human but it is possible that the sample is infected with a retrovirus, the genome of the retrovirus is treated as an additional chromosome. In this way, it is possible to extend the visualization methods disclosed in the present disclosure to identify insertion of nucleic acid constructs, such as retroviruses, into the genome of the sample under study.

[00100] So, for example, the disclosed techniques can use the bar codes to distinguish the following two scenarios. One is a human sample with HPV virus free floating in the sample but the virus hasn’t been inserted into the human DNA. They are a free floating molecule - separate molecules, separate virus, separate human DNA. In that case, the measured sequence reads are going to include reads that map to HPV as well as the human genome but there will

not be bar codes in common with the HPV and the human genome meaning that the human genome and the HPV are distinct. On the other hand, if the HPV molecule has been inserted into a human chromosome or two, what will be measured is are sequence reads that map to both a human chromosome and the HPV at the same time and share the same bar codes meaning that they exist in the same molecule as opposed to separate molecules (e.g., the HPV has been incorporated into a human chromosome). Moreover, the bar codes can be used to localize the precise location(s) of the HPV insertion into the human chromosome.

[00101] Figure 3 illustrates the data that is obtained from a biological sample of a subject (e.g., a particular human). This data is summarized in the form of a nucleic acid sequence dataset 126. In some instances, a full-genome run of the type described above produces 30-40 gigabytes worth of data. In accordance with some aspects of the present disclosure, such raw data is condensed into a nucleic acid sequence dataset 126 that is a fraction of the size of the raw data. In some embodiments, although the raw data is condensed to form the nucleic acid sequence dataset 126, the dataset 126 is still too large to load into the RAM of typical computers. For instance, in some embodiments, nucleic acid sequence dataset 126 is five gigabytes or larger, ten gigabytes or larger, or fifteen gigabytes or larger.

[00102] As illustrated in Figure 3, the exemplary nucleic acid sequencing dataset 126 is organized into three parts, a header 302, a synopsis 308, and a data section 340. The purpose of the header 302 is to delineate the components 304 of the dataset 126 as well as, optionally, provide the version 306 of the dataset 126 structure, e.g., version 1.7. In some embodiments, the header 302 is formatted as a JSON structure to facilitate loading using web based applications such as a web browser. See the URL json.org, which is hereby incorporated by reference. For instance, in some embodiments, the header is formatted as a JSON object: beginning with { (left brace) and ending with } (right brace), with each name is followed by : (colon) and the name/value pairs are separated by , (comma). In one exemplary embodiment, the header 302 that specifies that the sequencing dataset has 126 has the components: fragment tracks (e.g., the length, position, barcode, and phase of all the fragments in the dataset), targets track (the regions of the genome selected by the capture protocol used during processing), structural variation track (lists of all the structural variants called in the sample), an index to a target dataset, vcf_index (an index that relates ranges of the genome to a position in the dataset 126 file), marker, phase block summary (a description of the various phase blocks in the test nucleic acid 206), genetrack (a description of all human genes, tagged with the number of SNPs in each gene), BAM data (associates ranges of the genome to the position in the file containing

read information for that range), summary (high level metrics extracted from the sequencing data), and refseq index (an index that contains a list of dbSNP identifiers (RSIDs) of SNPs that are called in the sample, thereby associating the RSID with its position in the genome).

[00103] The synopsis section 308 contains data that is read by haplotype visualization tool 148 into volatile (e.g., random access) memory, typically in its entirety, when the dataset 126 is accessed. This data consists of indexes into the data section 340 as well as other data that is referenced frequently by visualization tool 148. As illustrated in Figure 3, the synopsis section 308 is split up into several components which correspond to the “index” array (e.g., component list 302) in the header section 302.

[00104] Summary 310 provides high level metrics extracted from the data. In some embodiments, summary 310 is used by summarization module 150 to provide summary data such as that illustrated in Figures 12 and 13. This includes the percentage of known SNPs (e.g., human SNPs) phased 1202, the longest phase block 1204, the effective barcode count 1206 (e.g., the number of unique barcodes used in the dataset 126), average fragment length 1208, mean of average fragment length 1210, percentage of fragments greater than a lower threshold (e.g., 20kb) 1212, fragment length histogram or other form of fragment length metric 1214, N50 phase block size 1216, phase block length histogram or other form of phase block length metric 1218, number of sequence reads represented by the dataset 1220, median insert size 1222, median depth 1224, percent of the target genome with zero coverage 1226, mapped reads percentage 1228, PCR duplication percentage 1230, on target bases (percent) 1232, coverage histogram or other form of coverage metric 1234, source of dataset in memory 112 (1234), identity of test nucleic acid (1236), genome source (1238), sex of donating organism (1240), dataset file format version 1242, and pointer to structural variant calls 1244 made for dataset 126 (1244).

[00105] Index to variant call data 312 is an example of an index found in the summary and it relates respective ranges 214 of the genome of the target nucleic acid to offsets 316 in the corresponding data section 340 where variant call data for the respective ranges is found.

[00106] In some embodiments, the phase block track 318 is stored in the synopsis section 308 of the nucleic acid sequencing dataset 126. More details of the architecture of an exemplary phase block track 318 are found in Figure 4. Referring to Figure 4, in some embodiments, the phase block track 318 includes a dictionary section 402 and a track data section 408. the track data section comprises a plurality of records 410. In some embodiments,

each record in the plurality of records comprises phase information for a corresponding chromosome. In some embodiments, each of the one or more data sections stores phase information for one or more corresponding chromosomes. In some embodiments, each of the one or more data sections stores phase information in an interval tree 422 format for a corresponding chromosome.

[00107] The dictionary 402 of the phase block track 318 comprises a plurality of names 404, and for each name 404, an offset 406 into the track data 408 where records for the corresponding name 404 are found. In some embodiments, the dictionary 402 for the phase block track 318 contains a single name, e.g., “phase_data”.

[00108] In some embodiments, the track data 408 is in JSON format. In some embodiments, each record 410 represents a phase block in the target nucleic acid. As such, in some embodiments, each record 410 specifies a chromosome number 412 that the phase block is on as well as the position where the phase block starts 414 on the chromosome 412 and a position where the phase block ends 416 on the chromosome 412. Moreover, there is a unique name 418 for each record and phasing information 420 about the phase block. In some embodiments, the purpose for the information 420 is to provide details of phasing information of the phase block. In some embodiments, a phase block includes information about two haplotypes corresponding to the two parents (e.g., respectively denoted haplotype “A” and haplotype “B”). Accordingly, in some embodiments, the phase information comprises PhaseASNP 422 (the number of counted single nucleotide polymorphisms on haplotype “A” in the phase block), Unphased SNP 424 (the number of counted single nucleotide polymorphisms of unknown haplotype in the phase block) and PhaseBSNP (the number of counted single nucleotide polymorphisms on haplotype “B” in the phase block). As such, the track data 408 holds certain phase block data (e.g., SNP counts) for the nucleic acid sequencing dataset 126. Techniques for phasing genomic data and phase blocks are described in Browning and Browning, “Haplotype phasing: Existing methods and new developments,” *Nat Rev Genet.*; 12(10): 703–714. doi:10.1038/nrg3054, which is hereby incorporated by reference in its entirety.

[00109] In some embodiments, the track data 408 is put into context by corresponding interval trees 422. As such, each record 410 is represented by a node 424 in an interval tree 422. Each such interval tree 422 is a ternary tree with each node 424 of the tree storing a midpoint of the node x_{med} 432. This midpoint 432 is the position of the midpoint, on the corresponding chromosome, of the phase block corresponding to the node. Each respective

node 424 has a link to a left child node 428, which corresponds to the phase block immediately to the left of the phase block represented by the respective node 424 in the genome of the species of the target (genetic source) organism. Each respective node 424 has a link to a right child node 430, which corresponds to the phase block immediately to the right of the phase block represented by the respective node 424. Each respective node 424 has a sorted set of nodes 425 that represent phase blocks that overlap the x_{med} 432 of the respective node 424 sorted by left hand position of such phase block. Each respective node 424 has a sorted set of nodes 436 that represent phase blocks that overlap the x_{med} 432 of the respective node 424 sorted by right hand position of such phase blocks. In some embodiments, sorted sets 425 and 436 are represented in a node 424 by arrays or linked lists. Each respective node 424 further includes a name 426, which is an offset in track data 410 to the record 410 that contains phase information 420 for the phase block corresponding to the respective node 424.

[00110] As illustrated in Figure 4, in some embodiments, there is a separate interval tree 422 for each chromosome in the phase block track. Such interval trees advantageously provide a quick way of identifying all records 410 pertaining to a user specified region of the of the target genome. An example of a phase block track 318 is found in Figure 5. In Figure 5, exemplary elements that correspond to the data structure of Figure 4 are illustrated.

[00111] Referring to Figure 3, in some embodiments, the synopsis 308 further comprises a refseq index 319, which is an index that contains the molecular variation (e.g., SNP) identifiers that are called in the sample corresponding to the nucleic acid sequencing dataset. The refseq index 319 associates each such identifier with its position in the genome of the target organism. In some embodiments, the refseq index 319 is stored as a JSON data structure. In some embodiments, each polymorphism identifier in the refseq index 319 is a dbSNP identifier found in the National Center for Biotechnology Information (NCBI) database. See Wheeler *et al.*, 2007, “Database resources of the National Center for Biotechnology Information,” Nucleic Acids Res. 35 (Database issue): D5–12, which is hereby incorporated by reference. Such dbSNP identifiers are termed reference SNP cluster IDs (RSIDs).

[00112] In some embodiments, the synopsis 308 further comprises a gene track 320, which provides a reference of human genes tagged with the number of SNPs found in each gene. More details of the architecture of an exemplary gene track 320 are found in Figure 6. Referring to Figure 6, in some embodiments, the gene track 320 includes a dictionary section 602, a track data section 608, and one or more data sections 628. In some embodiments, each

of the one or more data sections stores gene information for a corresponding chromosome. In some embodiments, each of the one or more data sections stores gene information for one or more corresponding chromosomes. In some embodiments, each of the one or more data sections stores gene information in an interval tree 628 format for a corresponding chromosome.

[00113] The dictionary 602 of the gene track 320 comprises a plurality of names 604, and for each name 604, an offset 606 into the track data 608 where records for the corresponding name 604 are found. In some embodiments, each name 604 in dictionary 602 is the name of a chromosome in the target genome.

[00114] In some embodiments, the track data 608 for gene track 320 comprises a plurality of gene records 610. In some embodiments, the track data 608 is in JSON format. In some embodiments, each gene record 610 represents a gene in the species of the target nucleic acid. As such, in some embodiments, each gene record 610 specifies a chromosome number 612 the corresponding gene is on, the position where the gene starts 614 on the chromosome 612 and a position where the gene ends 616 on the chromosome 612. Moreover, there is a unique name 618 for each gene record and gene information 620 about the gene. In some embodiments, the purpose for the information 620 is to provide genetic information about the gene, such as, for example, an alternative name 622 for the gene, a count of single nucleotide polymorphisms 624 on the gene, and a direction (e.g., plus or minus) 626 of the gene.

[00115] In some embodiments, the track data 608 is put into context by the corresponding interval trees 628. Each gene record 610 forms a node 630 in an interval tree 628. Each interval tree 628 is a ternary tree with each node 630 storing a midpoint of the node x_{med} 642. This midpoint 642 is the position of the midpoint, on the corresponding chromosome, of the gene corresponding to the node. Each respective node 630 has a link to a left child node 638, which corresponds to the gene immediately to the left (lesser position on the chromosome) of the gene represented by the respective node 630 in the species of the target organism. Each respective node 630 has a link to a right child node 640, which corresponds to the gene immediately to the right of the gene (greater position on the chromosome) represented by the respective node 630 in the species of the target organism. Each respective node 620 has a sorted set of nodes 632 that respectively represent genes that overlap x_{med} 632 of the respective node 620 sorted by left hand position. Each respective node 630 has a sorted set of nodes 630 that respectively represent genes that overlap the x_{med} 642 of the respective node 630 sorted by right hand position. In some embodiments, sorted sets 632 and 644 are represented in

a node 630 by arrays or linked lists. Each respective node 630 further includes a name 636, which is an offset in track data 608 to the gene record 610 that contains genetic information 620 for the gene corresponding to the respective node 630.

[00116] As illustrated in Figure 6, in some embodiments, there is a separate interval tree 628 for each chromosome in the gene track 320. Such interval trees advantageously provide a quick way of identifying all records 610 pertaining to a user specified region of the of the target genome. An example of a gene track 320 is found in Figure 7. In Figure 7, exemplary elements that correspond to the data structure of Figure 6 are illustrated.

[00117] In some embodiments, the synopsis 308 further comprises an exon track 322. In some embodiments, the exon track 322 has the same architecture as the gene track 320, the exception being that whereas the gene track 320 represents genetic information for genes in the species of the target organism, the exon track 320 provides genetic information for exons in the species of the target organism.

[00118] In some embodiments, the synopsis 308 further comprises an index to read data 324. This index 324 provides an index into sequence / read data 1048 in the data section 340 of the nucleic acid sequencing set, which is described in more detail below with reference to Figure 10. Referring to Figure 3, the index 324 comprises a database which associates identifiers to the barcodes used in the dataset (not shown). The database (lookup table) which associates identifiers to the barcodes used in the dataset is a useful way to compress the size of read data 1048, because identifiers can be used instead of the longer actual barcodes. This is because not all theoretically possible bar codes, for a given degree of information content, are used in a given dataset 126.

[00119] The index 324 further comprises a per chromosome array of chromosome-offset --> file-offset associations 328 into read data 1048 as well as a length of each such data element which allow lookup of the corresponding data for a specific genomic range. In some embodiments the read data is stored as a blocked index, and each record 328 is a fixed bit record for each entry in a BAM file that was incorporated into the dataset 126. Each such entry in the BAM file is organized into chunks within the data section 340 of the file. The index 324 in the synopsis 308 helps to find the correct chunk within the data section 340 to read. Referring to Figure 10, the corresponding architecture of the sequence / read data 1048 indexed by index 324 is disclosed. For each chromosome, read data 1048 is stored in chunks 1050. In

some embodiments, each data chunk 1050 is an array of 64-bit structures 1052 in the following format:

```

6      5            4            3            2            1            0
[3210987654321098765432109876543210987654321098765432109876543210]
[0XLRIIIIIIIIIIIIIIIIIIIIIEEEEEEEEEEEEEEESSSSSSSSSSSSSSSSSSSSS]
```

where O is always O, X indicates the read quality is below a threshold value (e.g., below 60), L indicates the read is from parental haplotype A, R indicates the read is from parental haplotype B, I is a numerical identifier corresponding to the barcode in the read, E is the 'end' length of the read, and S is the 'start' position of this read, relative to the start of the chunk 1050. More generally, referring to Figure 10, each structure 1052 corresponds to a single read from the target nucleic acid for the single organism of a species and comprises a start (offset), a length, an indicator to a bar code and some flags. In some embodiments the start within structure 1052 is the real position on the chromosome minus the start value stored for the chunk 1050 in the chromosome offset field of record 328 of index 324. Advantageously, this allows for avoidance of larger repetition of genomic coordinates in the structures 1052. Such coordinates can be in the billions and thus would required 30 bits to store. Advantageously, by chunking, as disclosed in sequence / read data 1048, each chunk covers up to about one million base pairs and thus each start (offset) in each structure 1052 in a chunk only needs 20 bits, since the range for any given chunk is specified by the chromosome offset / length portions of the corresponding record 328 in the index 324 stored in the synopsis 308. Similarly, as outlined above, in preferred embodiments, the barcode field in structure 1052 doesn't store the actual barcode. In some embodiments, the barcode indicator in structure 1052 is a 24-bit index into a barcode table that is stored in the index 324. So, when the actual barcode associated with a particular read is needed, the structure 1052 corresponding to the read is accessed, and the 24-bit bar code indicator in the structure 1052 is queried against the barcode table in the index 324 to obtain the bar code. In this way, 30 bit bar codes in the structures 1052 are avoided. In some embodiments, the bar code is greater than 30 bits (e.g., 32 bits, 34 bits, 36 bits or larger) and the indicator to the bar code in structure 1052 is greater than 20 bits (e.g., 22 bits, 24 bits, 26 bits or larger). In some embodiments, the bar code is less than 30 bits (e.g., 28 bits, 26 bits, 24 bits or smaller) and the indicator to the bar code in structure 1052 is less than 20 bits (e.g., 18 bits, 16 bit, 14 bits or smaller). In some embodiments, each data chunk 1050 is an array of structures 1052 having the same predetermined size (e.g., 128 bits, 64 bits, 32 bits, or some other fixed bit size).

[00120] In some embodiments, the synopsis 308 further comprises a structural variant dataset track 330. In some embodiments, the structural variants dataset track 330 comprises a listing of the called structural variants in the sample represented by the dataset 126. More details of the architecture of an exemplary structural variant dataset track 330 are found in Figure 8. Referring to Figure 8, in some embodiments, the structural variant dataset 330 includes a dictionary section 802, a track data section 808, and one or more data sections 840. In some embodiments, each of the one or more data sections 840 stores structural variant call information for a corresponding chromosome. In some embodiments, each of the one or more data sections 840 stores structural variant call information for one or more corresponding chromosomes. In some embodiments, each of the one or more data sections 840 stores gene information in an interval tree format for a corresponding chromosome.

[00121] The dictionary 802 of the structural variant dataset track 330 comprises a plurality of names 804, and for each name 804, an offset 606 into the track data 808 where records for the corresponding name 804 are found. In some embodiments, each name 804 in dictionary 802 is the name of a chromosome in the target genome.

[00122] In some embodiments, the track data 808 for structural variant dataset track 330 comprises a plurality of structural variant records 810. In some embodiments, the track data 808 is in JSON format. In some embodiments, each structural variant record 810 represents a structural variant call made for the target nucleic acid of the single organism represented by the dataset 126. As such, in some embodiments, each structural variant record 810 specifies a chromosome number 812, a start position 814 represented by the structural variation, a stop position 816 represented by the structural variation on the chromosome 812, a unique name 818 for the structural variation, and information 820 about the structural variation. In some embodiments, the structural variant dataset track 330 includes information analogous, corresponding to, or in a BEDPE format to advantageously concisely describe disjoint genome features, such as structural variations or paired-end sequence alignments. See the URL bedtools.readthedocs.org/en/latest/content/general-usage.html, which is hereby incorporated herein by reference. Accordingly, in some embodiments, the information section 820 in each structural variant record 810 includes a chromosome 1 name 822, which is the name of the chromosome on which the first end of the feature exists. In some embodiments chromosome 1 name 822 is in string format, for example, “chr1”, “III”, “myChrom”, or “contig1112.23.”

[00123] In some embodiments, the information section 820 in each record 810 further comprises a start 1 position 830, which is a zero-based starting position of the first end of the feature on chromosome 1 name 822.

[00124] In some embodiments, the information section 820 in each record 810 further comprises stop 1 (end 1) position 826, which is the one-based ending position of the first end of the feature (e.g., structural variation) represented by record 810 on chromosome 1 name 822.

[00125] In some embodiments, the information section 820 in each record 810 further comprises chromosome 2 name 836, which is the name of the chromosome on which the second end of the feature represented by record 810 exists. In some embodiments chromosome 2 name 836 is in string format, for example, “chr1”, “III”, “myChrom”, or “contig1112.23.”

[00126] In some embodiments, the information section 820 in each record 810 further comprises a start 2 position 828, which is the zero-based starting position of the second end of the feature represented by record 810 on chromosome 2 name 836.

[00127] In some embodiments, the information section 820 in each record 810 further comprises a stop 2 (end 2) position 824, which is the one-based ending position of the second end of the feature (e.g., structural variation) represented by record 810 on chromosome 2 name 836.

[00128] In some embodiments, the information section 820 in each record 810 further comprises a name of the structural variant field 834, which is the name of the feature (e.g., structural variation) represented by record 810. In some embodiments, the name of the structural variant 834 is in string format, for example, “LINE”, “Exon3”, “HWIEAS_0001:3:1:0:266#0/1”, or “my_Feature”.

[00129] In some embodiments, the information section 820 in each record 810 further comprises a quality (score) field 832, which is any metric the scores the quality of the feature (e.g., structural variation) represented by record 810. In some embodiments, quality 832 is in string format thereby permitting the expression of quality of the feature in any scientific metric, e.g., p-values, mean enrichment values, *etc.*

[00130] In some embodiments, the information section 820 in each record 810 further comprises further information 838 on the feature represented by the record 81, such as edit distance for each end of an alignment, or “deletion”, “inversion”, *etc.*

[00131] Continuing to refer to Figure 8, in some embodiments, the track data 808 is put into context by the corresponding interval trees 840. Each record 810 forms a node 842 in an interval tree 840. Each interval tree 840 is a ternary tree with each node 842 storing a midpoint of the node x_{med} 852. This midpoint 852 is the position of the midpoint, on the corresponding chromosome, of the feature (e.g., structural variant) corresponding to the node and represented by the corresponding record 810. Each respective node 842 has a link to a left child node 848, which corresponds to the feature (e.g., structural variant) immediately to the left (lesser position on the chromosome) of the feature represented by the respective node 842 in the dataset 126. Each respective node 842 has a link to a right child node 850, which corresponds to the feature (e.g., structural variant) immediately to the right (greater position on the chromosome) of the feature represented by the respective node 842 in the dataset 126. Each respective node 842 has a sorted set of nodes 854 that respectively represent features (e.g., structural variant) that overlap x_{med} 852 of the respective node 842 sorted by left hand position. Each respective node 842 has a sorted set of nodes 844 that respectively represent features that overlap the x_{med} 852 of the respective node 842 sorted by right hand position. In some embodiments, sorted sets 844 and 854 are represented in a node 840 by arrays or linked lists. Each respective node 840 further includes a name 846, which is an offset in track data 808 to the record 810 that contains information 820 for the feature (e.g., structural variation) corresponding to the respective node 840.

[00132] As illustrated in Figure 8, in some embodiments, there is a separate interval tree 840 for each chromosome in the structural variant dataset track 330. Such interval trees advantageously provide a quick way of identifying all records 810 pertaining to a user specified region of the of the target genome. An example of a portion of a structural variant dataset track 330 is found in Figure 9. In Figure 9, exemplary elements that correspond to the data structure of Figure 8 are illustrated.

[00133] Referring to Figure 3, in some embodiments, the synopsis 308 further comprises an index 332 to the target dataset 342. The target dataset 342 comprises the regions of the at least one target nucleic acid in the sample that were selected for sequencing in the nucleic acid sequencing dataset. In some embodiments index 332 and target dataset 342 are stored in a blocked JSON index. The blocked JSON index includes a single JSON object in the synopsis section (the index 332) and multiple JSON objects in the data section (the target dataset 342). The index 332 is used to calculate which data components must be read to fulfill a particular query. In some embodiments, the index 332 is split up by chromosome. For each

chromosome, the index 332 stores an array (record) 334 associating ranges on that chromosome with the offset at which specific data for that range may be found in the target dataset. In some embodiments, the target dataset 342 contains many independent arrays. Each array contains all of the ranges (and associated data) for one contiguous range of the genome. Each array in the target dataset 342 corresponds to a single array (entry) 334 in the index 332. In some embodiments, each such array in the target dataset is sized to contain about 1,000 entries. Because it is possible for a specific range to overlap multiple “chunks”, the same data may be written into multiple consecutive arrays. Referring to Figure 3, in some embodiments, the synopsis 308 further comprises an index 336 to the fragment dataset 344. The fragment dataset 344 comprises the length, position, barcode, and phase of all the fragments in the nucleic acid sequencing dataset. A fragment is the nucleic acid from a single partition, as described above. In some embodiments index 336 and fragment dataset 344 are stored in a blocked JSON index. The blocked JSON index includes a single JSON object in the synopsis section (the index 336) and multiple JSON objects in the data section (the fragment dataset 344). The index 336 is used to calculate which data components must be read to fulfill a particular query. In some embodiments, the index 336 of is split up by chromosome. For each chromosome, the index 336 stores an array 338 associating ranges on that chromosome with the offset at which specific data for that range may be found in the fragment dataset 344. An example of a data chunk in the fragment dataset 344 is:

```
{
  "Chromosome" : "chr1",
  "Name" : "19002" ,
  "Info" : {
    "h0" : "0.100000017888" ,
    "h1" : "0.899999982112",
    "hmix" : "0.0\n",
    "phsae_set" : "107163622",
    "ps_start" : "7163622",
    "be" : "CGTICCGTGGTATA-1",
    "ps_end" : "7276533"
    "Stop" : 7235518,
    "Start" : 7213929
  }
}
```

[00134] Thus, as the above provides, the disclosed nucleic acid sequencing datasets 126 of the present disclosure provide a streamlined file format that combines several forms of data that is conventionally found in separate files along with data that is of only secondary value. Advantageously, the nucleic acid sequencing dataset 126 file format is self-contained and has all the data required to support the features of haplotype visualization tool 148.

[00135] Figures 12 - 30 illustrate an embodiment of the haplotype visualization tool 148 that reads nucleic acid sequencing datasets 126. In some embodiments, the haplotype visualization tool 148 is a variant oriented and haplotype aware genome browser. To produce such views, the haplotype visualization tool 148 overlays data from several sources as tracks into a single unified nucleic acid sequencing dataset 126 for display that can be scrolled and zoomed. In some embodiments, the tracks that are stored includes phased variant calls, phase blocks, genes, exons, structural variant breakpoints and read count (coverage) as tracks. One such embodiment for how such information is stored is disclosed in Figure 3 and described above. Advantageously the disparate information in the nucleic acid sequencing set can be displayed in a single display. The haplotype visualization tool 148 is distinguished from other genome browsers by its ability to show phasing information. Referring to Figures 12 and 13, from the summarization module displayed in Figures 12 and 13, a user can advantageously use the search prompt 1250 to select regions of the nucleic acid sequencing dataset for further analysis. In some embodiments, through search prompt 1250, the haplotype visualization tool 148 supports a broad range of valid search syntaxes such as chr1:1000000 (select the first million nucleotides of chromosome 1), chr1:1000000-2000000 (select the second million nucleotides of chromosome 1), BRCA1, BRCA2 (select BRCA1 and BRCA2), and chr1:1000000-2000000, chr2:5000000-6000000 (select the second million nucleotides of chromosome 1 and the fifth million nucleotides of chromosome 2). In some embodiments, the user provides a symbolic name of a gene and the haplotype visualization tool 148 converts this symbolic name to the appropriate genomic coordinates by using one or more lookup tables that convert symbolic names to genomic coordinates. Advantageously, a user can provide in a single search a mix of absolute coordinate ranges and gene names. In some embodiments, a user provides a single search query that includes multiple loci. Responsive to such a query, the haplotype visualization tool 148 parses the multiple loci and provides results for each such query. In some embodiments, the user provides a search query of syntax $X_1:N_1-N_2$, where X_1 is an identity of a selected first chromosome or a selected first contig sequence, N_1 is a selected start position within the first chromosome or the selected first contig sequence, and N_2 is a selected end position within the first chromosome or the selected first contig sequence. As used in this context, the term “contig” means any “contig” from a reference genome which could correspond to an isolated molecule of interest that isn't a chromosome or an incompletely assembled part of a chromosome. In some embodiments, the user provides a search query of syntax $X_1:N_1-N_2$, where X_1 is an identity within a selected first chromosome or a selected first contig sequence, N_1 is a selected start position within the first chromosome or the selected first

contig sequence, and N_2 is a selected end position within the first chromosome or the selected first contig sequence. In some embodiments, the user provides a search query of syntax $X_1:N_1$, where X_1 is an identity of a selected first chromosome or a selected first contig sequence, and N_1 is a number of nucleotides, beginning at the origin of the first chromosome or the selected first contig sequence.

[00136] In some embodiments, a user provides a search query of syntax $Y_1, Y_2, \dots, Y_N$, where each Y_i in $Y_1, Y_2, \dots, Y_N$ is either an alphanumeric identification of a selected gene, a selection of a chromosomal region, or selection of a region of a contig sequence. In some such embodiments, a first Y_i in $Y_1, Y_2, \dots, Y_N$ is an identity of a first chromosome or a first contig sequence having the syntax $X_1:N_1-N_2$, where X_1 is an identity of the first chromosome or the first contig sequence, N_1 is a selected start position within the first chromosome or the first contig sequence, and N_2 is a selected end position within the first chromosome or the first contig sequence, and a second Y_i in $Y_1, Y_2, \dots, Y_N$ is an alphanumeric identification of a selected gene. In other such embodiments, a first Y_i in $Y_1, Y_2, \dots, Y_N$ is an identity of a first chromosome or a first contig sequence having the syntax $X_1:N_1-N_2$, where X_1 is an identity of the first chromosome or the first contig sequence, N_1 is a selected start position within the first chromosome or the first contig sequence, and N_2 is a selected end position within the first chromosome or the first contig sequence, and a second Y_i in $Y_1, Y_2, \dots, Y_N$ is an alphanumeric identification of a selected gene. In some embodiments, the request is converted, without human intervention, to genomic coordinates by comparison of the request against one or more lookup tables that match alphanumeric entries of genes to genomic coordinates. In some embodiments, the request comprises one or more gene names, one or more genomic coordinates, or a combination thereof.

[00137] Advantageously, the haplotype visualization tool 148 can be invoked in a variety of different system topologies. For instance, referring to Figure 31, in some embodiments, the haplotype visualization tool 148 operates on a client computer 3102 and accesses the nucleic acid sequence dataset remotely across a network connection. For instance, referring to Figure 31, in some embodiments, the haplotype visualization tool 148 tool is on a client computer system 3102 that communicates with the structural variation and phasing visualization system 100 across a network connection 3106. One such embodiment of the present disclosure provides a system 3100 for providing structural variation or phasing information 3100 over a network connection to a remote client computer 3102. Referring to Figures 1 and 32, the system 3100 comprises a server 100 having one or more microprocessors

102, a persistent memory (e.g., hard drive) and a non-persistent memory (e.g., random access memory). One of skill in the art will appreciate that persistent memory is memory that stores information even when system 100 is powered down whereas non-persistent memory is not able to store information when system 100 is powered down. Moreover, one of skill in the art will appreciate that access times to data stored in persistent memory is slower than access times to data stored in non-persistent memory. Further still, non-persistent memory is more expensive than persistent memory. As such, the disclosed nucleic acid datasets 126, which are large, are typically relegated to storage in persistent memory. In some embodiments, a nucleic acid sequencing dataset is 1 gigabyte or larger, 5 gigabytes or larger, or 10 gigabytes or larger.

[00138] In some embodiments, the persistent memory and the non-persistent memory, collectively referenced as memory 112 in Figure 1, store one or more nucleic acid sequence datasets 126. Each respective nucleic acid sequencing dataset 126 in the one or more nucleic acid sequence datasets corresponds to at least one target nucleic acid in a respective sample in a plurality of samples. The respective sample is associated with a genome of a species. Referring to Figure 3, the respective nucleic acid sequencing dataset 126 comprises (i) a header 302, (ii) a synopsis 308, and (iii) a data section 340.

[00139] The data section 340 comprises a plurality of sequencing reads and is the largest component of the dataset 126. Each respective sequencing read in the plurality of sequencing reads comprises a first portion that corresponds to a subset of at least one target nucleic acid in the respective sample and a second portion that encodes a respective identifier for the respective sequencing read in a plurality of identifiers. Each respective identifier is independent of the sequence of the at least one target nucleic acid. The plurality of sequencing reads collectively includes the plurality of identifiers.

[00140] The persistent memory and the non-persistent memory further collectively store one or more programs that use the one or more microprocessors 102 to provide a haplotype visualization tool 148 to the client for installation on the remote client computer. In turn, a request, sent from the client over the network connection, is received for structural variation or phasing information using a first dataset 126 in the one or more datasets. Responsive to receiving the request, the request is automatically filtered by loading the header 302 and the synopsis 308 of the first dataset into the non-persistent memory if not already loaded into the non-persistent memory while retaining the data section 340 in persistent memory. In this way, the amount of non-persistent memory is minimized. The request is compared to the synopsis 308 of the first dataset thereby identifying one or more portions of the data section of the first

dataset. In particular, the various components of the synopsis 308, as described in further detail below, are used to identify which portions of the data 340 are needed to fulfill the request. In some embodiments, the request identifies a particular dataset 126 and a region of a genome. In some embodiments, the request identifies a particular dataset 126 and one or more genes. In some embodiments, the request identifies a particular dataset 126 and one or more exons. Once the portions of the data section that are needed to fulfill the request are identified, they are loaded into non-persistent memory and the requested structural variation or phasing information is formatted for display on the client computer 3102 using the first dataset. This formatted structural variation or phasing information is then sent over the network connection 3106 to the client device for display on the client device. In some embodiments, as disclosed in Figure 1, a client computer is not used and the haplotype visualization tool is resident on the structural variation and phasing visualization system 100.

[00141] Now that advantages of splitting up the nucleic acid sequence dataset 126 have been explained, graphical user interface features of the haplotype visualization tool 148, and its component modules (e.g., summarization module 150, phase visualization module 152, structural variations module 154, etc.) will be described in further detail. Turning to Figure 12, once a user has entered a query in panel 1250 phase visualization module 152 may be used to view the phase of the query as illustrated in Figures 14 through 16. For instance, upon entering the query chr1+100000000-chr1+105000000 (or chr1:100000000-chr1:105000000), the selected region is illustrated in the genome browser (phase visualization module 152) illustrated in Figure 14A. Here, the selected region of the genome is advantageously shown in a way that reflects the actual physical structure of the selected region: there are two copies of the genome, and this is reflected by showing two tracks, one for each haplotype - haplotype 1 (1402) and haplotype 2 (1404), and a middle area 1406 where the parental haplotype has not been determined. Small insertions and deletions are mapped to each haplotype based on phasing algorithms. Portions of the selected region that have been phased to the first haplotype are shown as bars in the corresponding portion of the first haplotype 1 region 1402, portions of the selected region that have been phased to the second haplotype are shown as bars in the corresponding portion of the second haplotype 1 region 1404, and portions of the selected region that have not been phased to a haplotype are shown as bars in the middle area 1406.

[00142] In the haplotype view, phased portions of the selected region are enclosed in black rectangular boxes 1440. The entire region illustrated in Figure 14A is in a single phase block 1440-1. This also the case for Figures 14B, Figure 15, and chromosomes 1 and 2 of

Figure 16. However, the displayed region of chromosome 4 in Figure 16 includes five different phase blocks, each demarked by a black rectangular box. These boxes demarcate phased blocks, a contiguous phased region of the chromosome as determined by phasing algorithms.

[00143] Vertical bars in the haplotype 1 (1402), haplotype 2 (1404), and middle area 1406 represent single nucleotide polymorphisms, small insertions and deletions. In some embodiments, these bars are color coded with a first color (e.g. grey) representing the reference genotype, and a second color (e.g., green) representing the alternative genotype.

[00144] A homozygous SNP will have a vertical bar spanning the two haplotype tracks and the middle area (unphased track) since homozygous variants cannot be phased. This is illustrated as element 2602 in Figure 26.

[00145] Phased heterozygous SNPs are placed on the haplotype tracks 1402 / 1404. This is illustrated as element 2604 in Figure 26.

[00146] Heterozygous SNPs are placed in the middle area 1405 (unphased track) sandwiched in between the haplotype tracks 1402 / 1404 when they are not phased. This is illustrated as element 2606 in Figure 26.

[00147] Finally, if both phased single nucleotide polymorphisms are of alternative genotype, two vertical bars of the second color (e.g., green) will be displayed in the haplotype tracks 1402 / 1404, one for each track. This is illustrated as element 2608 in Figure 26.

[00148] Dark regions, such as region 2710 of Figure 27, of the haplotype track represent areas with high SNP density. Clicking on a region 2710 zooms into individual SNPs within the region 2710. Furthermore, in some embodiments, when this is done, a pop-up box 2712 will appear with a link allowing the user to zoom in on the SNP group. In general, the box 2712 provides additional information on the SNP, such as position, the reference genotype, observed genotypes of haplotype 1 and 2 in the sample, the gene where SNP is found (if associated with a gene), phasing quality, and allele counts of the two observed genotypes. The box 2712 can be dismissed by clicking on an X on a corner of the box. In some embodiments, the phasing quality provided for the SNP is a Phred-like score used to quantify the phasing quality of a SNP.

[00149] Referring to Figure 28A, when a user clicks on one of the alleles for a variant, a rectangular box (e.g., rectangular box 2802) highlights that variant. The number 2804 displayed next to the highlighted variant represents the number of barcodes that are associated with the selected allele for that variant. For instance, in Figure 28A, the number “31” is

displayed next to box 2802 indicating that the number of barcodes that are associated with the selected allele for that variant is 31. There are also numbers displayed on the top and/or bottom of variants adjacent to box 2802. Each such number represents the number of barcodes that overlap between the selected allele and one of the two alleles of the adjacent variants. Numbers displayed in a first color (e.g., black) agree with the phasing call of the variant 2802, while numbers displayed in a second color (e.g., red) disagree with the call. The greater the barcode overlap there is between neighboring variants, the more confidence there is in the phasing of the variant. As an example, for the reference call at Chr7: 117,216,030 of Figure 28A, there is a 31 (2804) on the top of the haplotype 1 panel 1402, indicating there are 31 barcodes associated with the reference allele at that position. Referring to Figure 28B, when the variant SNV at the same position 2802 is selected, 13 barcodes support the phasing and the labeled neighboring SNVs change as seen in Figure 28B.

[00150] In some embodiments the genome browser further provides a chromosome map 1424 and the location 1426 on the chromosome that is being displayed. Referring to Figure 14A, at the top of the browser, a miniature chromosome 1424 with the centromere marked by a dark rectangle is shown with chromosome bands marked by light rectangles. A triangle 1426 indicates the location currently in zoom, giving the user an overall view of the region selected using search bar 1250 with respect to the rest of the chromosome.

[00151] The disclosed genome browser further provides a graphic representation 1408 of each gene that is in the displayed genomic region. This genes track 1408 displays annotated reference genes. Multiple genes can be displayed using the search bar 1250 by entering the genes of interest. The direction of each gene is indicated with arrows. Although not illustrated in Figure 14A, exons are highlighted with dark shades. This feature is illustrated in Figures 26-28. In some embodiments, overlapping genes are shown on a maximum of three tracks in the genes track 1408 but many genes may be displayed using the search bar.

[00152] The disclosed genome browser further provides a graphic representation 1410 of exons that are in the displayed genomic region.

[00153] The disclosed genome browser further provides a coverage track 1412 for the coverage in the displayed genomic region. Aligned sequence reads are shown on the coverage track. Each vertical bar in the coverage track 1412 shows the average coverage-per-base for the area of the genome under the bar. The height is scaled such that maximum height is four times the median coverage. In some embodiments, when a user clicks on a portion of the

coverage track 1412, the mean reads per base pair and total number of reads is displayed in a coverage details pop-up black box for that portion of the coverage track.

[00154] The disclosed genome browser further provides a breakpoints track 1414 in the displayed region. Structural variants including inter-chromosomal translocations, gene fusions, inversions and deletions are highlighted in the breakpoints track 1414. Structural variants are arbitrarily numbered in the display. Structural variant call are indicated in a first color (e.g., orange) in the breakpoints track 1414 and structural variant candidate are specified in a second color (e.g., grey) in the breakpoints track 1414. To display structural variant breakpoint pairs, a user can click on the structural variant displayed for the gene, as illustrated in Figure 29. The structural variant is displayed in the details box 2902. By selecting “Zoom in on this breakpoint” 2094 in details box 2902, the other side of the breakpoint is brought up as an additional haplotype track, zoomed to the breakpoints, as illustrated in Figure 30.

[00155] Advantageously, what is not shown in some embodiments of the display mode of the disclosed genome browser, illustrated in Figure 14A, are base calls, error rates, specific reads, and alignments. Rather, the disclosed genome browser operate at a higher level in order to provide a more conceptual indication of what is going on in the selected region and to provide this information in a way that is easy to understand. For this reason, some embodiments of the disclosed browser provide a display mode, such as the display mode illustrated in Figure 14A, in which all of the sequence read data is not shown.

[00156] Referring to Figure 14A, zoom affordance 1420 can be used zoom into a subset of the region identified by search bar 1250 and zoom affordance 1422 can be used to zoom out of the region. In addition, a user can zoom in to a specific gene by clicking on the icon in region 1408 representing the specific gene.

[00157] In some embodiments, the search bar 1250 of the disclosed genome browser provides intelligent auto complete features. For instance, when a user starts typing a gene name in the search bar 1250, the genome browser auto completes on the genes. In some embodiments, the genome browser accomplishes this by comparing partial search queries that the user enters against genomic information stored in the nucleic acid sequencing dataset such as the names of genes in the gene track. Advantageously, in such embodiments the search bar 1250 auto completes on gene names. For instance, referring to Figure 17, when a user enters the expression “atp” into the search bar, several possible matches 1702-1 through 1702-10 found within the nucleic acid sequence dataset 126 are displayed.

[00158] As illustrated in Figures 12 through 30, the haplotype visualization tool 148 provides structural variation or phasing (e.g. haplotype) information for a nucleic acid sequence dataset.

[00159] In particular, referring to Figures 12 and 13, selection of the phasing / haplotypes toggle 1252 of the haplotype visualization tool 148 invokes the phase visualization module 152 as illustrated in Figures 14-17 and Figures 26-30. As illustrated in Figures 14-17 and Figures 26-30, visually separated tracks for haplotypes as well as a virtual track for variants that could not be assigned to either haplotype is provided. Phased variants can have a wide number of classifications including: unphased, homozygous, and/or heterozygous-with-no-reference-reads, heterozygous-with-reference-reads. The haplotype visualization tool 148 applies visually distinct stylings to these different configurations so that a user can quickly tell them apart. The haplotype visualization tool 148 can display the amount of barcode evidence used in assigning a variant to a particular phase block. In some embodiments, when the user “clicks” on a variant, every other visible variant is decorated with the count of barcodes that overlapped with the selected variant. Data that contradicts the called haplotype is highlighted. The haplotype visualization tool 148 also allows the user to view multiple regions at once. This is displayed as separate haplotype in different areas of the screen. In this mode “counts” are shared between each displayed region allowing the user to view barcodes overlaps between distant regions of the genome.

[00160] Again referring to Figures 12 and 13, selection of the structural variants toggle 1254 of the haplotype visualization tool 148 invokes the structural variants module 154 as illustrated in Figures 23-25 and 33-34. The matrix view provided by the structural variants module 154 encompasses a method for visualizing candidate structural variants. The visualization works by quantifying two (possibly overlapping) regions of the genome (test nucleic acid data) into chunks of between 100 and 10,000 base pairs per chunk. The number of shared barcodes between the reads in every pair of chunks is computed. The resulting matrix (with the chunks from one region as the rows and the other region as the columns) can be displayed as a two dimensional image (heat map), as illustrated in Figures 23-25 and 33-34. In some embodiments, the color of a pixel corresponds to number of distinct overlapping barcodes between a specific chunk (e.g. window) of each region. For example, consider two regions with consecutive chunks with the following barcodes:

[00161] (1) AAA, ACA ACA, AGT GTG

[00162] (2) GTG, AAA CCC ACA ,AAA

[00163] There are nine pairs of chunks between region (1) and region (2) which can be placed in a matrix such as the one set forth below in Table 1.

[00164] Table 1 – matrix of pairs of chunks between region (1) and region (2).

	(1)		
(2)	AAA,ACA vs GTG,AAA	AAA,ACA vs CCC	AAA,ACA vs ACA,AAA
	ACA,AGT vs GTG,AAA	ACA,AGT vs CCC	ACA, AGT vs ACA,AAA
	GTG vs GTG,AAA	GTG vs CCC	GTG vs ACA,AAA

Computing the overlap between the two sets of barcodes in each cell yields the values set forth in Table 2.

[00165] Table 2 – matrix values between region (1) and region (2).

	(1)		
(2)	1	0	2
	0	0	1
	1	0	0

[00166] Table 2 can be displayed by the structural variants module 154 as a heat map which efficiently shows areas of low and high barcode correlation to the user. In some embodiments, the structural variants module 154 provides additional information, such as gene and exon boundaries overlaid with the matrix to allow easy alignment of the data to known places of interest. In some embodiments, the structural variants module 154 also allows a textual copy of the matrix to be downloaded for analysis with other computer programs. In some embodiments, the user may adjust the region of the genome that is visualized in the structural variants module 154 by scrolling or zooming in real time. In some embodiments, the user can adjust the resolution (chunk size / window size) to avoid aliases or overload when looking at very small or very large areas of the genome.

[00167] Some embodiments of the present disclosure provide a system 100 for viewing nucleic acid sequencing data (e.g., information obtained from nucleic acid sequencing datasets

126). The system 100 comprises one or more microprocessors 102 and a memory 112. The memory stores a nucleic acid sequence dataset 126 corresponding to at least one target nucleic acid in a sample. The memory further stores one or more programs (e.g., the haplotype visualization tool 148) that use the one or more microprocessors to obtain the nucleic acid sequencing dataset that comprises a plurality of sequencing reads from a sample. Then, a request is obtained from a user (e.g., through search bar 1250 of the haplotype visualization tool 148 illustrated in Figures 12 and 13) that specifies a genomic region represented by the nucleic acid sequencing dataset. Advantageously, this request can be in any of the syntaxes disclosed in the present disclosure. In some embodiments, the genomic region in the request is an entire chromosome. In some embodiments, the genomic region in the request is between 100 and 10000 bases of the chromosome. In some embodiments, the genomic region in the request is between 10 and 1×10^5 bases of the chromosome. In some embodiments, the genomic region in the request is between 10 and 1×10^6 bases of the chromosome. In some embodiments, the genomic region in the request is between 10 and 1×10^7 bases of the chromosome. In some embodiments the request is for a gene in the genome of the sample. Responsive to obtaining the request, the request is parsed by obtaining a plurality of sequencing reads 1048 within the genomic region of the request from the nucleic acid sequencing dataset 126. Next, a scan window is run against the plurality of sequencing reads thereby creating a plurality of windows, each respective window of the plurality of windows corresponding to a different region of the genomic region in the request and including an identity of each identifier (e.g., bar code) of each sequencing read in the different region of the genomic region in the nucleic acid sequencing dataset. Further, referring for example to Figure 34, a two dimensional heat map 3312 that represents each possible window pair in the plurality of windows is displayed. Each respective window pair is displayed in the two dimensional heat map as a color selected from a color scheme based upon the number of identifiers in common in the respective window pair. It will be appreciated that window size will depend on the amount of the genome the user has requested to visualize. In some embodiments, when the user has requested to visualize a small region of the genome, smaller windows sizes are used and when the user has requested to visualize a larger region of the genome, larger window sizes are used.

[00168] Referring to Figures 33 and 34, affordances 3302 and 3304 provide unique tools to clarify the displayed information. First, selection of the “hide expected overlap” affordance 3302 causes the bar code overlap signal that is expected from the genome being in a normal state, where bar codes associated with reads that are next to each other because they are

supposed to be, to be hidden. Compare Figure 33, with affordance 3302 not selected, with Figure 34, with affordance 3302 selected. The view provided when affordance 3302 is selected is intended to emphasize those parts of the genome that are now touching each other that are unexpected. For instance, this view highlights a structural variation, a trans location from one chromosome to another that, based on the reference genome, you wouldn't expect to be there but suddenly the bar codes now shows the association. As such, affordance 3302 activates a filter that hides the normal signal and highlights the unexpected signals. In other words, the number of identifiers in common in respective window pairs is down-weighted to remove bar code signals arising from bar codes that are expected to be proximate to each other based on the reference genome sequence. In some embodiments, the filter associated with affordance 3302 considers the mean length of the fragments of the target nucleic acid that were sequenced (e.g. 50 kb). Bar codes that are within this threshold distance of the mean length of fragments do not contribute to the heat map when affordance 3302 is activated. In some embodiments, the filter is enabled by taking the entire set of bar codes in the nucleic acid sequencing dataset 126 that have been aligned against a reference genome. Then, only those regions along the reference genome that exhibit a gap that is greater than the mean fragment length displayed. As such, the affordance 3302 filter act to filter out the expected and highlights the differences between the bar code data and a reference genome.

[00169] Referring to affordance 3304, each respective sequence read 1048 is mapped to a location on a reference genome with a confidence value that represents a probability that the respective sequence read was correctly mapped. The default is to only show data for sequence reads when this confidence value satisfies a stringent (high) threshold value so that misleading information is not displayed. But sometimes a user still wants to see information for sequence reads that do not satisfy the stringent threshold confidence value. For instance, sometimes, when too much data is filtered out based on the confidence threshold unusual artifacts may appear in the heat map. For instance, regions of the heat map will appear to have no data. In reality, such regions may be just regions where the confidence in the localization of sequence reads 1048 is low (e.g., regions of the genome that exhibit extensive repeats). To determine whether there is actual no data (perhaps indicating an extensive structural variation) affordance 3304 allows the user to remove (or lower) the stringent threshold value and to permit the display of data from sequence reads 1048 that have been mapped to the reference genome with lower confidence values. In this way, the user can determined whether there is in fact a structural variation at sites that were missing data when the stringent threshold value was

turned on or whether the genomic region simply represents a region where the confidence values for the sequence reads is low.

[00170] In a typical use case scenario associated with affordance 3304, sequence reads 1084 that do not satisfy a quality threshold are discarded and so are not used to in downstream phasing algorithms and structural variation algorithms. The consequence of discarding such sequence reads is that it can introduce what looks like structure in the heat map plot illustrated in Figures 33 and 34. For instance, some regions of the map may lighten up and some lines may be introduced giving rise to the question of whether there something happening in the actual sample that's causing this to change the signal. By selecting affordance 3304, the discarded reads are put back into the phasing and/or structural variation algorithms regardless of their quality score to see if this causes removal of the observed artifacts in the plot. In this way, artifacts of the data can be teased out so that when a region of the plot is missing, before and after applying affordance 3304, confidence that the observed artifact represents an artifact (e.g., structural variation) in the at least one target nucleic acid in a respective sample or an artifact arising from discarding data from sequence reads 1048.

[00171] Referring to Figure 34, the extent of barcode overlap between respective regions of the target nucleic acid is signified on a color scale 3406 by the number of barcodes (from sequence reads localized to the respective regions of the target nucleic acid) that overlap. Thus, in some embodiments, a color scheme is used, with each particular color in the color scheme uniquely representing a certain number of overlapping barcodes. For instance, if a first and second section of the target nucleic acid have in common a first number of barcodes, the color associated with the first number in the color scheme is used to represent the combination of the first and second section of the target nucleic acid. As illustrated in Figure 34, the X axis 3308 and Y axis 3310 each represent the target nucleic acid and thus the coordinates of the first and second section of the target nucleic acid within the target nucleic acid define an X,Y position in the two dimensional grid, and the color associated with the value of the first number of barcodes is used to color this X,Y position in the two dimensional grid in accordance with the color scheme. In some embodiments, when a first and second section of the target nucleic acid have no barcodes in common, the color scheme dictates that the color used for the X,Y position that represents the combination of the first and second section of the target nucleic acid be white. In some embodiments, when a first and second section of the target nucleic acid have only a few barcodes in common (e.g, in various embodiments, only one barcode in common, only two barcodes in common, only three barcodes in common, only four barcodes in common

or only five barcodes in common), the color scheme dictates that the color used for the X,Y position that represents the combination of the first and second section of the target nucleic acid be grey. That is, in such embodiments, the first position in the color scheme is white, meaning no shared barcodes and the second position in the color scheme is grey, meaning a minimal set of barcodes in common. In some embodiments, there are 10 different values in the color scheme corresponding to 10 different values of shared sequence reads. In some embodiments, there are 11 different values in the color scheme corresponding to 11 different values of shared sequence reads. In some embodiments, there are 12 different values in the color scheme corresponding to 12 different values of shared sequence reads. In some embodiments, there are 13 different values in the color scheme corresponding to 13 different values of shared sequence reads. In some embodiments, there are 14 different values in the color scheme corresponding to 14 different values of shared sequence reads. In some embodiments, there are 15 different values in the color scheme corresponding to 15 different values of shared sequence reads. In some embodiments, there are between five and one hundred different values in the color scheme corresponding to between five and one hundred different values of shared sequence reads.

[00172] Referring to Figure 34, affordance 3308 can be used to pan (translational movement of) the view initially selected by search field 1250 so that different regions of the reference genome can be viewed. Referring to Figure 34, affordance 3310 can be used to zoom the view initially selected by search field 1250 so that different amounts the reference genome can be viewed.

[00173] In some embodiments, the different views offered (e.g., haplotype / phase 152, structural variants 154, and reads 156) by the haplotype visualization tool 148 are all linked. For instance, a user may navigate from one view to another to see the same data using an alternate visualization without reentering information using affordances 1252, 1254, and 1256. For instance, the user may toggle between the matrix view of the structural variants module 154 and the haplotype view of the phase visualization module 152.

[00174] A “smart” search affordance 1250 is employed in the various views. Referring to Figure 17, as a user types in the search affordance 1250, the program will attempt to auto-complete the partial query with real gene names or other forms of chromosomal locations in real time. In some embodiments, each time the user enters another character in the search affordance 1250, the partial query in the search affordance 1250 is queried against a lookup table in the subject nucleic acid sequencing dataset 126. In some embodiments, this lookup

table is the gene track 320 and/or the exon track 322. Advantageously, in some embodiments, the haplotype visualization tool 148 maintains a history of past user queries. Thus, when a user starts to enter a new query, matches (or partial matches) against former queries are also displayed to the user for selection. This is particularly useful given the complex query syntax that is supported by the search bar 1250 in some embodiments. For example, as discussed above a user may query for multiple regions at once by separating queries with a variety of punctuators. A user may also enter a genomic coordinate directly in a number of formats.

[00175] In some embodiments, system 100 stores genomic data to be displayed in a custom file format (e.g., the format of nucleic acid sequencing dataset 126). The file is generated by a “preprocessor” which takes reference data, the VCF file, the BAM, file and the structural variant file as inputs and produces a single output nucleic acid sequencing dataset 126. The nucleic acid sequencing dataset 126 contains all of the information that is required to display a given dataset. The file is organized into several sections. A small synopsis section 308 that is roughly 25MB and a much larger data section 340 (100MB to 20GB). These sections are further subdivided as described above. When the nucleic acid sequencing dataset 126 is loaded, it loads just the index section into memory. System 100 uses that data to find appropriate ranges of the data section to load into memory on-demand. Variant calls and read information is stored in the data section, the rest of the data loupe needs is small enough to store in the index section.

[00176] The data section is organized to chunks which are about ~250KB in some embodiments. When system 100 requires information stored in the data section it consults the relevant index in the synopsis section (e.g., gene track, exon track, etc.) to find the chunk that should have the data and loads the entire chunk into memory. In some embodiments, the chunks for variant data are JSON-encoded structures containing the variant data as well as the supporting barcode information. In some embodiments, the chunks for read data have an array of small (8-byte) data structures in which each structure contains the position, length, and barcode of a single read. In some embodiments, both variant and read data is sorted by genomic position so that in general, system 100 will make only a small number of on-disk reads to acquire all of the data it needs to display a given subset of the data. In some embodiments, the rest of the data that system 100 needs for visualization (such as the location of genes, structural variant breakpoints, etc) is stored in the index (synopsis) section of the nucleic acid sequencing dataset 126 file as an “itree”. An itree is an implementation of an interval tree. It is a reusable data structure (usually encoded in JSON) for annotating ranges of the genome. Thus

exons, genes, phase blocks, and structural variant breakpoints are all encoded with the same mechanism even though they are displayed differently.

[00177] Plural instances may be provided for components, operations or structures described herein as a single instance. Finally, boundaries between various components, operations, and data stores are somewhat arbitrary, and particular operations are illustrated in the context of specific illustrative configurations. Other allocations of functionality are envisioned and may fall within the scope of the implementation(s). In general, structures and functionality presented as separate components in the example configurations may be implemented as a combined structure or component. Similarly, structures and functionality presented as a single component may be implemented as separate components. These and other variations, modifications, additions, and improvements fall within the scope of the implementation(s).

[00178] It will also be understood that, although the terms “first,” “second,” etc. may be used herein to describe various elements, these elements should not be limited by these terms. These terms are only used to distinguish one element from another. For example, a first object could be termed a second object, and, similarly, a second object could be termed a first object, without changing the meaning of the description, so long as all occurrences of the “first object” are renamed consistently and all occurrences of the “second object” are renamed consistently. The first object and the second object are both objects, but they are not the same object.

[00179] The terminology used herein is for the purpose of describing particular implementations only and is not intended to be limiting of the claims. As used in the description of the implementations and the appended claims, the singular forms “a,” “an” and “the” are intended to include the plural forms as well, unless the context clearly indicates otherwise. It will also be understood that the term “and/or” as used herein refers to and encompasses any and all possible combinations of one or more of the associated listed items. It will be further understood that the terms “comprises” and/or “comprising,” when used in this specification, specify the presence of stated features, integers, steps, operations, elements, and/or components, but do not preclude the presence or addition of one or more other features, integers, steps, operations, elements, components, and/or groups thereof.

[00180] As used herein, the term “if” may be construed to mean “when” or “upon” or “in response to determining” or “in accordance with a determination” or “in response to detecting,” that a stated condition precedent is true, depending on the context. Similarly, the phrase “if it is

determined (that a stated condition precedent is true)” or “if (a stated condition precedent is true)” or “when (a stated condition precedent is true)” may be construed to mean “upon determining” or “in response to determining” or “in accordance with a determination” or “upon detecting” or “in response to detecting” that the stated condition precedent is true, depending on the context.

[00181] The foregoing description included example systems, methods, techniques, instruction sequences, and computing machine program products that embody illustrative implementations. For purposes of explanation, numerous specific details were set forth in order to provide an understanding of various implementations of the inventive subject matter. It will be evident, however, to those skilled in the art that implementations of the inventive subject matter may be practiced without these specific details. In general, well-known instruction instances, protocols, structures and techniques have not been shown in detail.

[00182] The foregoing description, for purpose of explanation, has been described with reference to specific implementations. However, the illustrative discussions above are not intended to be exhaustive or to limit the implementations to the precise forms disclosed. Many modifications and variations are possible in view of the above teachings. The implementations were chosen and described in order to best explain the principles and their practical applications, to thereby enable others skilled in the art to best utilize the implementations and various implementations with various modifications as are suited to the particular use contemplated.

What is claimed:

1. A system for providing structural variation or phasing information over a network connection to a remote client computer, the system comprising one or more microprocessors, a persistent memory and a non-persistent memory, the persistent memory and the non-persistent memory collectively storing one or more nucleic acid sequence datasets, wherein

each respective nucleic acid sequencing dataset in the one or more nucleic acid sequence datasets corresponds to at least one target nucleic acid in a respective sample in a plurality of samples, wherein

the respective sample is associated with a genome of at least one species,

the respective nucleic acid sequencing dataset comprises (i) a header, (ii) a synopsis, and (iii) a data section,

the data section comprises a plurality of sequencing reads,

each respective sequencing read in the plurality of sequencing reads comprises a first portion that corresponds to a subset of at least one target nucleic acid in the respective sample and a second portion that encodes a respective identifier for the respective sequencing read in a plurality of identifiers,

each respective identifier is independent of the sequence of the at least one target nucleic acid, and

the plurality of sequencing reads collectively include the plurality of identifiers, and wherein

the persistent memory and the non-persistent memory further collectively store one or more programs that use the one or more microprocessors to:

provide a visualization tool for installation on the remote client computer;

obtain a request, sent from the remote client computer from a user, over a network connection, for structural variation or phasing information using a first dataset in the one or more datasets; and

responsive to obtaining the request, automatically parse the request by:

(i) loading the header and the synopsis of the first dataset into the non-persistent memory if not already loaded into the non-persistent memory while retaining the data section in persistent memory,

(ii) comparing the request to the synopsis of the first dataset thereby identifying one or more portions of the data section of the first dataset,

(iii) loading the one or more identified portions of the data section into non-persistent memory, wherein the loading loads less than the entirety of the data section,

(iv) formatting structural variation or phasing information for display on the client computer using the first dataset, and

(v) transmitting the formatted structural variation or phasing information over the network connection to the remote client computer for display on the remote client computer.

2. A system for providing structural variation or phasing information, the system comprising one or more microprocessors, a persistent memory and a non-persistent memory, the persistent memory and the non-persistent memory collectively storing one or more nucleic acid sequence datasets, wherein

each respective nucleic acid sequencing dataset in the one or more nucleic acid sequence datasets corresponds to at least one target nucleic acid in a respective sample in a plurality of samples, wherein

the respective sample is associated with a genome of at least one species,

the respective nucleic acid sequencing dataset comprises (i) a header, (ii) a synopsis, and (iii) a data section,

the data section comprises a plurality of sequencing reads,

each respective sequencing read in the plurality of sequencing reads comprises a first portion that corresponds to a subset of at least one target nucleic acid in the respective sample and a second portion that encodes a respective identifier for the respective sequencing read in a plurality of identifiers,

each respective identifier is independent of the sequence of the at least one target nucleic acid, and

the plurality of sequencing reads collectively include the plurality of identifiers, and wherein

the persistent memory and the non-persistent memory further collectively store one or more programs that use the one or more microprocessors to:

provide a visualization tool;

obtain a request from a user, through the visualization tool, for structural variation or phasing information using a first dataset in the one or more datasets,

responsive to obtaining the request, automatically parse the request by:

(i) loading the header and the synopsis of the first dataset into the non-persistent memory if not already loaded into the non-persistent memory while retaining the data section in persistent memory,

(ii) comparing the request for sequence information to the synopsis of the first dataset thereby identifying one or more portions of the data section of the first dataset,

(iii) loading the one or more identified portions of the data section into non-persistent memory, wherein the loading loads less than the entirety of the data section,

(iv) formatting structural variation or phasing information for display in the visualization tool using the first dataset, and

(v) displaying the formatted structural variation or phasing information in the visualization tool.

3. A system for obtaining structural variation or phasing information over a network connection from a remote computer, wherein the system comprises one or more microprocessors, and a memory that stores one or more programs, wherein the one or more programs use the one or more microprocessors to execute a method comprising:

(A) invoking a visualization tool;

(B) obtaining, through the visualization tool from a user, a request for structural variation or phasing information in a first nucleic acid sequencing dataset from among one or more nucleic acid sequencing datasets stored on the remote computer, wherein each respective nucleic acid sequencing dataset in the one or more nucleic acid sequence datasets corresponds to at least one target nucleic acid in a respective sample in a plurality of samples, wherein

the respective sample is associated with a genome of at least one species,

the respective nucleic acid sequencing dataset comprises (i) a header, (ii) a synopsis, and (iii) a data section,

the data section comprises a plurality of sequencing reads,

each respective sequencing read in the plurality of sequencing reads comprises a first portion that corresponds to a subset of at least one target nucleic acid in the respective sample and a second portion that encodes a respective identifier for the respective sequencing read in a plurality of identifiers,

each respective identifier is independent of the sequence of the at least one target nucleic acid, and

the plurality of sequencing reads collectively include the plurality of identifiers;

(C) sending the request to the remote computer over the network connection, wherein the remote computer has persistent memory and non-persistent memory, thereby causing the remote computer to execute a method comprising:

(i) loading the header and the synopsis of the first dataset into the non-persistent memory if not already loaded into the non-persistent memory of the remote computer while retaining the data section in persistent memory,

(ii) comparing the request for sequence information to the synopsis of the first dataset thereby identifying one or more portions of the data section of the first dataset,

(iii) loading the one or more identified portions of the data section into non-persistent memory, wherein the loading loads less than the entirety of the data section, and

(iv) formatting structural variation or phasing information; and

(D) receiving the formatted structural variation or phasing information over the network connection from the remote computer for display in the visualization tool.

4. A system for providing structural variation or phasing information, the system comprising one or more microprocessors, and a memory, wherein the system has access to one or more nucleic acid sequence datasets, wherein

each respective nucleic acid sequencing dataset in the one or more nucleic acid sequence datasets corresponds to at least one target nucleic acid in a respective sample in a plurality of samples, wherein

the respective sample is associated with a genome of at least one species,

the respective nucleic acid sequencing dataset comprises (i) a header, (ii) a synopsis, and (iii) a data section,

the data section comprises a plurality of sequencing reads,

each respective sequencing read in the plurality of sequencing reads comprises a first portion that corresponds to a subset of at least one target nucleic acid in the respective sample and a second portion that encodes a respective identifier for the respective sequencing read in a plurality of identifiers,

each respective identifier is independent of the sequence of the at least one target nucleic acid, and

the plurality of sequencing reads collectively include the plurality of identifiers, and wherein

the memory further collectively stores one or more programs that use the one or more microprocessors to:

provide a visualization tool;

obtain a request from a user, through the visualization tool, for structural variation or phasing information using a first dataset in the one or more datasets,

responsive to the obtained request, automatically parse the request by:

(i) comparing the request for sequence information to the synopsis of the first dataset thereby identifying one or more portions of the data section of the first dataset,

(ii) formatting structural variation or phasing information for display in the visualization tool using the identified one or more portions of the data section of the first dataset, wherein the one or more portions of the data section of the first dataset is less than the entirety of the data section of the first dataset, and

(iii) displaying the formatted structural variation or phasing information in the visualization tool.

5. The system of any one of claims 1-4, wherein the header delineates a plurality of components in the respective nucleic acid sequencing dataset.

6. The system of claim 5, wherein the plurality of components comprises two or more components selected from the group consisting of a summary, an index to variant call data, a phase block track, a refseq index track, a gene track, an exon track, an index to read data, a structural variant dataset track, an index to a target dataset, and an index to a fragment dataset.

7. The system of claim 6, wherein the plurality of components comprises the summary and wherein the summary comprises two or more items in the group consisting of:

a percentage of known SNPs phased in the respective nucleic acid sequencing dataset,

a longest phase block in the respective nucleic acid sequencing dataset,

a number of unique barcodes used in the respective nucleic acid sequencing dataset,

an average fragment length in the respective nucleic acid sequencing dataset,

a mean of the average fragment length in the respective nucleic acid sequencing dataset,

a percentage of fragments greater than a lower threshold in the respective nucleic acid sequencing dataset,

a fragment length histogram in the respective nucleic acid sequencing dataset,

an N50 phase block size in the respective nucleic acid sequencing dataset,

a phase block histogram in the respective nucleic acid sequencing dataset,

a number of sequence reads represented by respective the nucleic acid sequencing dataset,
 a median insert size in the respective nucleic acid sequencing dataset,
 a median depth in the respective nucleic acid sequencing dataset,
 a percent of the target genome with zero coverage in the respective nucleic acid sequencing dataset,
 a mapped reads percentage for the respective nucleic acid sequencing dataset,
 a PCR duplication percentage for the respective nucleic acid sequencing dataset,
 a coverage histogram for the in the respective nucleic acid sequencing dataset,
 an identity of a test nucleic acid that forms the basis for the respective nucleic acid sequencing dataset,
 a genome source for the respective nucleic acid sequencing dataset,
 a sex of an organism that originated the at least one test nucleic acid of the respective nucleic acid sequencing dataset,
 a sex of the organism that originate the respective sample of the in the respective nucleic acid sequencing dataset,
 a dataset file format version of the in the respective nucleic acid sequencing dataset, and
 a pointer to a plurality of structural variant calls made for the respective nucleic acid sequencing dataset.

8. The system of claim 6, wherein the plurality of components comprises the index to variant call data that provides a correspondence between respective ranges of the genome of the species to offsets in the data section where variant call data for the respective ranges is found.

9. The system of claim 6, wherein the plurality of components comprises the phase block track and wherein the phase block track comprises (i) a dictionary and (ii) a track data section comprising phase information for one or more chromosomes in the genome of the at least one species.

10. The system of claim 9, wherein the dictionary comprises a plurality of names, and for each respective name in the plurality of names, an offset into the track data where records for the corresponding name are found.

11. The system of claim 9, wherein the track data section comprises a plurality of records and wherein each record in the plurality of records represents a phase block in the target nucleic acid.

12. The system of claim 11, wherein the tract data section is in the JSON file format.

13. The system of claim 11, wherein each respective record in the plurality of records specifies:

- (i) a chromosome number corresponding to the respective record,
- (ii) a position where the phase block starts on the chromosome,
- (iii) a position where the phase block ends,
- (iv) a unique name for the record, and
- (v) phasing information about the phase block.

14. The system of claim 11, wherein each respective record in the plurality of records is represented by a node in a plurality of nodes in a respective interval tree in a plurality of interval trees, wherein each interval tree in the plurality of interval trees represents a chromosome in a plurality of chromosomes for the at least one species.

15. The system of claim 14, wherein

a node in the plurality of nodes of a first interval tree in the plurality of interval trees stores a midpoint of the node,

the midpoint of the node is a position of the midpoint, on the corresponding chromosome, of the phase block corresponding to the node,

each respective node in the plurality of nodes of the first interval tree has a link to a left child node, which corresponds to the phase block immediately to the left of the phase block represented by the respective node in the genome of the at least one species,

each respective node in the plurality of nodes of the first interval tree has a link to a right child node, which corresponds to the phase block immediately to the right of the phase block represented by the respective node in the genome of the at least one species,

each respective node in the plurality of nodes of the first interval tree has a sorted set of nodes that represent phase blocks that overlap the midpoint of the respective node sorted by left hand position of such phase block, and

each respective node in the plurality of nodes of the first interval tree has a sorted set of nodes that represent phase blocks that overlap the midpoint of the respective node sorted by right hand position of such phase blocks.

16. The system of claim 15, wherein each respective node in the plurality of nodes of the first interval tree further includes a name, which is an offset in the track data section to the record in the plurality of records that contains phase information for the phase block corresponding to the respective node.

17. The system of claim 5, wherein the header further comprises the version of the dataset structure used by the nucleic acid sequencing dataset.

18. The system of claim 6, wherein the plurality of components comprises the refseq index, wherein the refseq index comprises an index of a plurality of molecular variation identifiers that are called in the sample.

19. The system of claim 18, wherein each respective molecular variation identifier in the plurality of molecular variation identifiers is dbSNP identifier.

20. The system of claim 6, wherein the plurality of components comprises the gene track, wherein the gene track comprises a plurality of genes and, for each respective gene in the plurality of genes, a number of single nucleotide polymorphisms in the respective gene.

21. The system of claim 6, wherein the plurality of components comprises the gene track and wherein the gene track comprises (i) a gene track dictionary and (ii) a gene track data section.

22. The system of claim 21, wherein the gene track dictionary comprises a plurality of names, and for each respective name in the plurality of names, an offset into the gene track data section where records for the corresponding name are found.

23. The system of claim 21, wherein the gene track data section comprises a plurality of records and wherein each record in the plurality of records represents a gene in the target nucleic acid.

24. The system of claim 23, wherein the gene tract data section is in the JSON file format.
25. The system of claim 23, wherein each respective record in the plurality of records specifies:
- (i) a chromosome number corresponding to the respective record,
 - (ii) a position on the chromosome where the gene represented by the respective record begins,
 - (iii) a position on the chromosome where the gene represented by the respective record ends,
 - (iv) a unique name for the respective record, and
 - (v) genetic information about the gene.
26. The system of claim 23 wherein the genetic information about the gene comprises an alternative name for the gene, a count of single nucleotide polymorphisms on the gene, or direction of the gene on the respective chromosome.
27. The system of claim 23, wherein each respective gene record in the plurality of gene records is represented by a node in a plurality of nodes in a respective interval tree in a plurality of interval trees, wherein each interval tree in the plurality of interval trees represents a chromosome in a plurality of chromosomes for the at least one species.
28. The system of claim 27, wherein
- a node in the plurality of nodes of a first interval tree in the plurality of interval trees stores a midpoint of the node,
 - the midpoint of the node is a position of the midpoint, on the corresponding chromosome, of the gene corresponding to the node,
 - each respective node in the plurality of nodes of the first interval tree has a link to a left child node, which corresponds to the gene immediately to the left of the gene represented by the respective node in the genome of the at least one species,
 - each respective node in the plurality of nodes of the first interval tree has a link to a right child node, which corresponds to the gene immediately to the right of the gene represented by the respective node in the genome of the at least one species,

each respective node in the plurality of nodes of the first interval tree has a sorted set of nodes that represent genes that overlap the midpoint of the respective node sorted by left hand position of the respective gene, and

each respective node in the plurality of nodes of the first interval tree has a sorted set of nodes that represent genes that overlap the midpoint of the respective node sorted by right hand position of such genes.

29. The system of claim 28, wherein each respective node in the plurality of nodes of the first interval tree further includes a name, which is an offset in the track data section to the gene record in the plurality of gene records that contains genetic information for the gene corresponding to the respective node.

30. The system of claim 6, wherein the plurality of components comprises the index to read data wherein the index to read data comprises a lookup table between a respective identifier in the plurality of identifiers and a shortened version of the respective identifier.

31. The system of claim 30, wherein

the data section is stored as a blocked index of records organized into a plurality of chunks,

each respective chunk in the plurality of chunks comprises a subset of the plurality of sequencing reads, and

the index to read data further comprises a per chromosome array of chromosome-offset to file-offset associations between the plurality of chunks and absolute positions in the genome of the at least one species that thereby assign each chunk in the plurality of chunks a different absolute position in genome of a species in the at least one species.

32. The system of claim 31, wherein

a respective chunk in the plurality of chunks is an array of structures, each respective structure in the array representing a corresponding sequencing read in the plurality of sequence reads in the respective chunk, and wherein each structure includes:

a first bit reserved to indicate read quality of the corresponding sequencing read,

a second bit to represent origin from a first parental haplotype for the corresponding sequencing read,

a third bit to represent origin from a second parental haplotype for the corresponding sequencing read,

a first plurality of bits to represent the shortened version of an identifier in the plurality of identifiers for the corresponding sequencing read,

a second plurality of bits to represent a length of the corresponding sequencing read, and

a third plurality of bits to represent a start position of the corresponding sequencing read relative to the absolute position in genome of the species assigned to the respective chunk.

33. The system of claim 32, wherein the species is human and the respective chunk represents about 1 million base pairs or less.

34. The system of claim 32, wherein the first plurality of bit is 20 bits, the species is human, the respective chunk represents about 1 million base pairs or less, and each structure in the array of structures is a predetermined bit size.

35. The system of claim 32, wherein the identifier, stored in the index to read data, that corresponds to the shortened version of the identifier requires 24 bits.

36. The system of claim 32, wherein the identifier, stored in the index to read data, that corresponds to the shortened version of the identifier requires 30 bits, 32 bits, 34 bits, or 36 bits.

37. The system of claim 6, wherein

the plurality of components comprises the structural variant dataset track, and

the structural variant dataset track comprises (i) a dictionary and (ii) a track data section comprising structural variant call information identified in the plurality of sequencing reads.

38. The system of claim 37, wherein the dictionary comprises a plurality of names, and for each respective name in the plurality of names, an offset into the track data where records for the corresponding name are found.

39. The system of claim 38, wherein a name in the plurality of names corresponds to a chromosome.

40. The system of claim 38, wherein

the track data section comprises a plurality of structural variant records, and

each structural variant record in the plurality of structural variant records represents a structural variant call made in the at least one target nucleic acid in the sample.

41. The system of claim 40, wherein the tract data section is in the JSON file format.

42. The system of claim 40, wherein a respective structural variant record in the plurality of structural variant records specifies a start point on a first chromosome and an end point on a second chromosome for the structural variant represented by the respective structural variant record and at least one on of:

(i) a name of the structural variant represented by the respective structural variant record, and

(ii) an expression of the quality in the confidence of the identification of the structural variant represented by the respective structural variant record.

43. The system of claim 40, wherein

each respective structural variant record in the plurality of structural variant records is represented by a node in a plurality of nodes in a respective interval tree in a plurality of interval trees, and

each interval tree in the plurality of interval trees represents a chromosome in a plurality of chromosomes for the species.

44. The system of claim 43, wherein

a node in the plurality of nodes of a first interval tree in the plurality of interval trees stores a midpoint of the node,

the midpoint of the node is a position of the midpoint, on the corresponding chromosome, of the structural variant corresponding to the node,

each respective node in the plurality of nodes of the first interval tree has a link to a left child node, which corresponds to the structural variant immediately to the left of the structural variant represented by the respective node in the genome of the species,

each respective node in the plurality of nodes of the first interval tree has a link to a right child node, which corresponds to the structural variant immediately to the right of the structural variant represented by the respective node in the genome of the species,

each respective node in the plurality of nodes of the first interval tree has a sorted set of nodes that represent structural variant that overlap the midpoint of the respective node sorted by left hand position of such structural variant, and

each respective node in the plurality of nodes of the first interval tree has a sorted set of nodes that represent structural variants that overlap the midpoint of the respective node sorted by right hand position of such structural variants.

45. The system of claim 6, wherein

the plurality of components comprises the index to the target dataset,

the target dataset comprises the regions of the at least one target nucleic acid in the sample that were selected for sequencing in the respective nucleic acid sequencing dataset,

the target dataset is indexed by a target dataset index stored in the synopsis, and

the target dataset is stored in the data section.

46. The system of claim 45, further comprising instructions, executed by the one or more physical processors to use the target dataset index to calculate which portions of the target dataset must be read from the target dataset to fulfill the request.

47. The system of claim 45, wherein

the target dataset index is split up by chromosome and, for each respective chromosome, and

the target dataset index stores a corresponding array associating ranges on the respective chromosome with the offset at which specific data for that range may be found in the target dataset.

48. The system of claim 6, wherein

the plurality of components comprises the index to the fragment dataset,

the fragment dataset comprises a length, chromosomal position, identifier, and phase of each fragment of the at least one target nucleic acid in the sample,

the fragment dataset is indexed by a fragment dataset index stored in the synopsis, and

the fragment dataset is stored in the data section.

49. The system of claim 48, further comprising instructions, executed by the one or more physical processors to use the fragment dataset index to calculate which portion of the fragment dataset must be read from the fragment dataset to fulfill the request.

50. The system of claim 49, wherein the fragment dataset index is split up by chromosome and, for each respective chromosome, the fragment dataset index stores a corresponding array associating ranges on the respective chromosome with the offset at which specific data for that range may be found in the fragment dataset.

51. The system of any one of claims 1-4, wherein the request is for phasing information in a region of the genome and the formatted phasing information includes a graphic representation comprising:

- a first haplotype track corresponding to a first parental haplotype of a first species in the at least one species in the region of the genome for the first dataset,

- a second haplotype track, corresponding to a second parental haplotype of the first species in the region of the genome for the first dataset,

- an indeterminate track corresponding to regions of the at least one nucleic acid sample that have not been assigned a parental haplotype in the region of the genome for the first dataset.

52. The system of claim 51, wherein each phase block in the region in the first dataset is encompassed by a rectangular box.

53. The system of claim 52, wherein the region in the first dataset comprises two or more phase blocks, five or more phase blocks, or ten or more phase blocks.

54. The system of claim 51, wherein the first haplotype track, the second haplotype track and the indeterminate track includes a plurality of vertical bars, wherein each vertical bar represents a single nucleotide polymorphism, insertion, or deletion, in the at least one target nucleic acid.

55. The system of claim 54, wherein each vertical bar in the plurality of vertical bars is color coded to indicate one of a reference genotype and an alternative genotype.

56. The system of claim 51, wherein the graphic representation further comprises a chromosome map and a location on the chromosome map that is encompassed by the region of the genome.

57. The system of claim 51, wherein the graphic representation further comprises a graphic representation of each gene that is in the region of the genome.

58. The system of claim 51, wherein the graphic representation further comprises a graphic representation of each exon that is in the region of the genome.

59. The system of claim 51, wherein the graphic representation further comprises a coverage track for the region of the genome, wherein the coverage track comprises a plurality of vertical bars, and wherein each respective vertical bar in the plurality of vertical bars indicates an average coverage-per-base in the first dataset for a corresponding portion of the genome under the bar.

60. The system of claim 51, wherein the graphic representation further comprises a breakpoints track for one or more structural variants arising in the region of the genome.

61. The system of claim 60, wherein the breakpoints track wherein the one or more structural variants comprise an inter-chromosomal translocation, a gene fusion, an inversion, or a deletion.

62. The system of claim 60, wherein a first structural variant in the one or more structural variants is formatted as a graphic that, when selected by the user, provides a structural variant zoom affordance that, when selected, provides an additional haplotype track zoomed to the breakpoints of the first structural variant.

63. The system of any one of claims 1-4, wherein the one or more nucleic acid sequencing datasets is a single nucleic acid sequencing dataset.

64. The system of any one of claims 1-4, wherein the one or more nucleic acid sequencing datasets is a plurality of nucleic acid sequencing datasets.

65. The system of claim 64, wherein the plurality of nucleic acid sequencing datasets comprises ten nucleic acid sequencing datasets, 100 nucleic acid sequencing datasets, or 1000 nucleic acid sequencing datasets.

66. The system of any one of claims 1-65, wherein the request is formatted in accordance with an expression syntax.

67. The system of claim 66, wherein the expression syntax is $X_1:N_1-N_2$, wherein

X_1 is an identity of a selected first chromosome or a selected first contig sequence,

N_1 is a selected start position within the first chromosome or the selected first contig sequence, and

N_2 is a selected end position within the first chromosome or the selected first contig sequence.

68. The system of claim 66, wherein the expression syntax is $X_1:N_1-N_2$, wherein

X_1 is an identity within a selected first chromosome or a selected first contig sequence,

N_1 is a selected start position within the first chromosome or the selected first contig sequence, and

N_2 is a selected end position within the first chromosome or the selected first contig sequence.

69. The system of claim 66, wherein the expression syntax is $X_1:N_1$, wherein

X_1 is an identity of a selected first chromosome or a selected first contig sequence, and

N_1 is a number of nucleotides, beginning at the origin of the first chromosome or the selected first contig sequence.

70. The system of claim 66, wherein the expression syntax is $Y_1, Y_2, \dots, Y_N$, wherein

each Y_i in $Y_1, Y_2, \dots, Y_N$ is either an alphanumeric identification of a selected gene, a selection of a chromosomal region, or selection of a region of a contig sequence.

71. The system of claim 70, wherein

a first Y_i in $Y_1, Y_2, \dots, Y_N$ is an identity of a first chromosome or a first contig sequence having the syntax $X_1:N_1-N_2$, wherein X_1 is an identity of the first chromosome or the first contig sequence, N_1 is a selected start position within the first chromosome or the first contig

sequence, and N_2 is a selected end position within the first chromosome or the first contig sequence, and

a second Y_i in $Y_1, Y_2, \dots, Y_N$ is an alphanumeric identification of a selected gene.

72. The system of claim 70, wherein

a first Y_i in $Y_1, Y_2, \dots, Y_N$ is an identity of a first chromosome or a first contig sequence having the syntax $X_1:N_1-N_2$, wherein X_1 is an identity of the first chromosome or the first contig sequence, N_1 is a selected start position within the first chromosome or the first contig sequence, and N_2 is a selected end position within the first chromosome or the first contig sequence, and

a second Y_i in $Y_1, Y_2, \dots, Y_N$ is an alphanumeric identification of a selected gene.

73. The system of any one of claims 1-4, wherein the request is converted, without human intervention, to genomic coordinates by comparison of the request against one or more lookup tables that match alphanumeric entries of genes to genomic coordinates.

74. The system of any one of claims 1-4, wherein the request comprises one or more gene names, one or more genomic coordinates, or a combination thereof.

75. The system of any one of claims 1-74, wherein obtaining the request is facilitated by displaying to the past requests that have been made by the user.

76. The system of any one of claims 1-74, wherein obtaining the request is facilitated by displaying to the user search expressions obtained from a lookup table that match a partial request provided by the user.

77. The system of any one of claims 1-4, wherein the respective sample is associated with a genome of a plurality of species and includes at least a portion of the genome of a first species and a portion of the genome of the second species.

78. The system of claim 77, wherein the portion of the genome of the first species is integrated into the portion of the genome of the second species.

79. The system of claim 78, wherein the first species is a retrovirus.

80. The system of claim 77, wherein the portion of the genome of the first species is not integrated into the portion of the genome of the second species.

81. A system for processing program output over a network connection using a local computer, wherein the local computer comprises one or more microprocessors, and a memory that stores one or more programs, wherein the one or more programs use the one or more microprocessors to execute a method in accordance with a first operating system running on the local computer, the method comprising:

(A) invoking a first instance of a first program;

(B) obtaining, through the first instance of the first program from a user, a login and a password to a user account on a remote computer;

(C) logging the user into the user account on the remote computer automatically, using the login and the password provided by the first instance of the first program, across a network connection between the local computer and the remote computer;

(D) responsive to successful login on the remote computer, automatically sending, without human intervention, a second instance of the first program configured to auto-install on the remote computer upon transmission to the remote computer;

(E) receiving from the remote computer a request to open a panel within the first instance of the first program, wherein the panel is originated by the second instance of the first program running on the remote computer, and wherein the panel solicits input from the user for controlling the second instance of the first program;

(F) responsive to receiving input from the user for controlling the second instance of the first program in the panel on the local computer, sending the input to the second instance of the first program on the remote computer; and

(G) receiving, from the remote computer across the network connection, output from the second instance of the first program responsive to the input and displaying the output at the local computer.

82. The system of claim 81, wherein remote computer is a server computer and the local computer is a desktop computer or laptop computer.

83. The system of claim 81, wherein the first operating system is other than the second operating system.

84. The system of claim 81, wherein the receiving (E) further solicits an encryption key from the user and wherein communication between the first instance on the client computer and the second instance on the remote computer is encrypted.

85. The system of claim 81, wherein the remote computer comprises a persistent memory and a non-persistent memory, the persistent memory and the non-persistent memory collectively storing one or more nucleic acid sequence datasets, wherein

each respective nucleic acid sequencing dataset in the one or more nucleic acid sequence datasets corresponds to at least one target nucleic acid in a respective sample in a plurality of samples, wherein

the respective sample is associated with a genome of a species,

the respective nucleic acid sequencing dataset comprises (i) a header, (ii) a synopsis, and (iii) a data section,

the data section comprises a plurality of sequencing reads,

each respective sequencing read in the plurality of sequencing reads comprises a first portion that corresponds to a subset of at least one target nucleic acid in the respective sample and a second portion that encodes a respective identifier for the respective sequencing read in a plurality of identifiers,

each respective identifier is independent of the sequence of the at least one target nucleic acid, and

the plurality of sequencing reads collectively include the plurality of identifiers, and wherein

the input received from the user for controlling the second instance of the first program is a request for structural variation or phasing information using a first dataset in the one or more datasets, and

responsive to receiving the request, the second instance of the first program filters the request by:

(i) loading the header and the synopsis of the first dataset into the non-persistent memory if not already loaded into the non-persistent memory while retaining the data section in persistent memory,

(ii) comparing the request to the synopsis of the first dataset thereby identifying one or more portions of the data section of the first dataset,

- (iii) loading the one or more identified portions of the data section into non-persistent memory,
- (iv) formatting structural variation or phasing information for display in the panel in the first instance of the first program using the first dataset, and
- (v) transmitting the formatted structural variation or phasing information as the output from the second instance of the first program over the network connection to the local computer for display in the first panel.

86. A system for viewing nucleic acid sequencing data, the system comprising one or more microprocessors and a memory, wherein the memory stores one or more programs that use the one or more microprocessors to:

- obtain a nucleic acid sequencing dataset corresponding to at least one target nucleic acid in a sample, the nucleic acid sequencing dataset comprising a plurality of sequencing reads from the sample, wherein

- each respective sequencing read in the plurality of sequencing reads comprises a first portion that corresponds to a subset of at least one target nucleic acid in the sample and a second portion that encodes a respective identifier for the respective sequencing read in a plurality of identifiers,

- each respective identifier is independent of the sequence of the at least one target nucleic acid, and

- the plurality of sequencing reads collectively include the plurality of identifiers, display a visualization tool;

- obtain a request from a user, through the visualization tool, that specifies a genomic region represented by the nucleic acid sequencing dataset; and

- responsive to the obtained request, parse the request by:

- (i) obtaining a plurality of sequencing reads within the genomic region from the nucleic acid sequencing dataset,

- (ii) running a scan window against the plurality of sequencing reads thereby creating a plurality of windows, each respective window of the plurality of windows corresponding to a different region of the genomic region and including an identity of each identifier of each sequencing read in the different region of the genomic region in the nucleic acid sequencing dataset, and

- (iii) displaying a two dimensional heat map that represents each possible window pair in the plurality of windows, wherein each respective window pair is displayed in

the two dimensional heat map as a color selected from a color scheme based upon the number of identifiers in common in the respective window pair.

87. The system of claim 86, wherein the visualization tool provides a hide expected overlap affordance, wherein

when a user engages the hide expected overlap affordance, identifiers of sequencing reads that are expected to be adjacent to each other in accordance with a reference genome do not contribute to a count of a number of identifiers in common in the respective window pair, and

when a user does not engage the hide expected overlap affordance, identifiers of sequencing reads that are expected to be adjacent to each other in accordance with a reference genome contribute to the count of the number of identifiers in common in the respective window pair.

88. The system of claim 86, wherein the visualization tool provides a quality affordance, wherein

when a user engages the quality affordance, identifiers of respective sequencing reads that do not satisfy a quality threshold quantifying a quality of assignment of the respective sequence reads to locations in a reference genome do not contribute to a count of the number of identifiers in common in the respective window pair,

when a user does not engage the quality affordance, identifiers of respective sequencing reads that do not satisfy a quality threshold quantifying a quality of assignment of the respective sequence reads to locations in a reference genome contribute to the count of the number of identifiers in common in the respective window pair.

89. The system of claim 86, wherein the visualization tool provides a zoom affordance to zoom the heat map.

90. The system of claim 86, wherein the visualization tool provides a pan affordance to pan the heat map.

91. A method for providing structural variation or phasing information over a network connection to a remote client computer, the method comprising:

at a system comprising one or more microprocessors, a persistent memory and a non-persistent memory, the persistent memory and the non-persistent memory collectively storing one or more nucleic acid sequence datasets, wherein

each respective nucleic acid sequencing dataset in the one or more nucleic acid sequence datasets corresponds to at least one target nucleic acid in a respective sample in a plurality of samples, wherein

the respective sample is associated with a genome of at least one species,

the respective nucleic acid sequencing dataset comprises (i) a header, (ii) a synopsis, and (iii) a data section,

the data section comprises a plurality of sequencing reads,

each respective sequencing read in the plurality of sequencing reads comprises a first portion that corresponds to a subset of at least one target nucleic acid in the respective sample and a second portion that encodes a respective identifier for the respective sequencing read in a plurality of identifiers,

each respective identifier is independent of the sequence of the at least one target nucleic acid, and

the plurality of sequencing reads collectively include the plurality of identifiers,

(A) providing a visualization tool for installation on the remote client computer;

(B) obtaining a request, sent from the remote client computer from a user, over a network connection, for structural variation or phasing information using a first dataset in the one or more datasets; and

(C) responsive to obtaining the request, automatically parsing the request by:

(i) loading the header and the synopsis of the first dataset into the non-persistent memory if not already loaded into the non-persistent memory while retaining the data section in persistent memory,

(ii) comparing the request to the synopsis of the first dataset thereby identifying one or more portions of the data section of the first dataset,

(iii) loading the one or more identified portions of the data section into non-persistent memory, wherein the loading loads less than the entirety of the data section,

(iv) formatting structural variation or phasing information for display on the client computer using the first dataset, and

(v) transmitting the formatted structural variation or phasing information over the network connection to the remote client computer for display on the remote client computer.

92. A method for providing structural variation or phasing information, the method comprising:

at a system comprising one or more microprocessors, a persistent memory and a non-persistent memory, the persistent memory and the non-persistent memory collectively storing one or more nucleic acid sequence datasets, wherein

each respective nucleic acid sequencing dataset in the one or more nucleic acid sequence datasets corresponds to at least one target nucleic acid in a respective sample in a plurality of samples, wherein

the respective sample is associated with a genome of at least one species,

the respective nucleic acid sequencing dataset comprises (i) a header, (ii) a synopsis, and (iii) a data section,

the data section comprises a plurality of sequencing reads,

each respective sequencing read in the plurality of sequencing reads comprises a first portion that corresponds to a subset of at least one target nucleic acid in the respective sample and a second portion that encodes a respective identifier for the respective sequencing read in a plurality of identifiers,

each respective identifier is independent of the sequence of the at least one target nucleic acid, and

the plurality of sequencing reads collectively include the plurality of identifiers, and wherein

(A) providing a visualization tool;

(B) obtaining a request from a user, through the visualization tool, for structural variation or phasing information using a first dataset in the one or more datasets,

(C) responsive to obtaining the request, automatically parsing the request by:

(i) loading the header and the synopsis of the first dataset into the non-persistent memory if not already loaded into the non-persistent memory while retaining the data section in persistent memory,

(ii) comparing the request for sequence information to the synopsis of the first dataset thereby identifying one or more portions of the data section of the first dataset,

(iii) loading the one or more identified portions of the data section into non-persistent memory, wherein the loading loads less than the entirety of the data section,

(iv) formatting structural variation or phasing information for display in the visualization tool using the first dataset, and

(v) displaying the formatted structural variation or phasing information in the visualization tool.

93. A method for obtaining structural variation or phasing information over a network connection from a remote computer, the method comprising:

at a system comprising one or more microprocessors, and a memory that stores one or more programs:

(A) invoking a visualization tool;

(B) obtaining, through the visualization tool from a user, a request for structural variation or phasing information in a first nucleic acid sequencing dataset from among one or more nucleic acid sequencing datasets stored on the remote computer, wherein each respective nucleic acid sequencing dataset in the one or more nucleic acid sequence datasets corresponds to at least one target nucleic acid in a respective sample in a plurality of samples, wherein

the respective sample is associated with a genome of at least one species,

the respective nucleic acid sequencing dataset comprises (i) a header, (ii) a synopsis, and (iii) a data section,

the data section comprises a plurality of sequencing reads,

each respective sequencing read in the plurality of sequencing reads comprises a first portion that corresponds to a subset of at least one target nucleic acid in the respective sample and a second portion that encodes a respective identifier for the respective sequencing read in a plurality of identifiers,

each respective identifier is independent of the sequence of the at least one target nucleic acid, and

the plurality of sequencing reads collectively include the plurality of identifiers;

(C) sending the request to the remote computer over the network connection, wherein the remote computer has persistent memory and non-persistent memory, thereby causing the remote computer to execute a method comprising:

(i) loading the header and the synopsis of the first dataset into the non-persistent memory if not already loaded into the non-persistent memory of the remote computer while retaining the data section in persistent memory,

(ii) comparing the request for sequence information to the synopsis of the first dataset thereby identifying one or more portions of the data section of the first dataset,

(iii) loading the one or more identified portions of the data section into non-persistent memory, wherein the loading loads less than the entirety of the data section, and

(iv) formatting structural variation or phasing information; and
(D) receiving the formatted structural variation or phasing information over the network connection from the remote computer for display in the visualization tool.

94. A method for providing structural variation or phasing information, the method comprising:

at a system comprising one or more microprocessors, and a memory, wherein the system has access to one or more nucleic acid sequence datasets, wherein

each respective nucleic acid sequencing dataset in the one or more nucleic acid sequence datasets corresponds to at least one target nucleic acid in a respective sample in a plurality of samples, wherein

the respective sample is associated with a genome of at least one species,

the respective nucleic acid sequencing dataset comprises (i) a header, (ii) a synopsis, and (iii) a data section,

the data section comprises a plurality of sequencing reads,

each respective sequencing read in the plurality of sequencing reads comprises a first portion that corresponds to a subset of at least one target nucleic acid in the respective sample and a second portion that encodes a respective identifier for the respective sequencing read in a plurality of identifiers,

each respective identifier is independent of the sequence of the at least one target nucleic acid, and

the plurality of sequencing reads collectively include the plurality of identifiers,

(A) providing a visualization tool;

(B) obtaining a request from a user, through the visualization tool, for structural variation or phasing information using a first dataset in the one or more datasets; and

(C) responsive to obtaining the request, automatically parsing the request by a method comprising:

(i) comparing the request for sequence information to the synopsis of the first dataset thereby identifying one or more portions of the data section of the first dataset,

(ii) formatting structural variation or phasing information for display in the visualization tool using the identified one or more portions of the data section of the first

dataset, wherein the one or more portions of the data section of the first dataset is less than the entirety of the data section of the first dataset, and

(iii) displaying the formatted structural variation or phasing information in the visualization tool.

95. A method for processing program output over a network connection, the method comprising:

at a local computer, wherein the local computer comprises one or more microprocessors, and a memory that stores one or more programs to:

(A) invoke a first instance of a first program;

(B) obtain, through the first instance of the first program from a user, a login and a password to a user account on a remote computer;

(C) log the user into the user account on the remote computer automatically, using the login and the password provided by the first instance of the first program, across a network connection between the local computer and the remote computer;

(D) responsive to successful login on the remote computer, automatically send, without human intervention, a second instance of the first program configured to auto-install on the remote computer upon transmission to the remote computer;

(E) receive from the remote computer a request to open a panel within the first instance of the first program, wherein the panel is originated by the second instance of the first program running on the remote computer, and wherein the panel solicits input from the user for controlling the second instance of the first program;

(F) responsive to receiving input from the user for controlling the second instance of the first program in the panel on the local computer, send the input to the second instance of the first program on the remote computer; and

(G) receive, from the remote computer across the network connection, output from the second instance of the first program responsive to the input and displaying the output at the local computer.

96. A method for viewing nucleic acid sequencing data, the method comprising:

at a system comprising one or more microprocessors and a memory, wherein the memory stores one or more programs that use the one or more microprocessors to:

obtain a nucleic acid sequencing dataset corresponding to at least one target nucleic acid in a sample, the nucleic acid sequencing dataset comprising a plurality of sequencing reads from the sample, wherein

each respective sequencing read in the plurality of sequencing reads comprises a first portion that corresponds to a subset of at least one target nucleic acid in the sample and a second portion that encodes a respective identifier for the respective sequencing read in a plurality of identifiers,

each respective identifier is independent of the sequence of the at least one target nucleic acid, and

the plurality of sequencing reads collectively include the plurality of identifiers, display a visualization tool;

obtain a request from a user, through the visualization tool, that specifies a genomic region represented by the nucleic acid sequencing dataset; and

responsive to obtaining the request, parse the request by:

(i) obtaining a plurality of sequencing reads within the genomic region from the nucleic acid sequencing dataset,

(ii) running a scan window against the plurality of sequencing reads thereby creating a plurality of windows, each respective window of the plurality of windows corresponding to a different region of the genomic region and including an identity of each identifier of each sequencing read in the different region of the genomic region in the nucleic acid sequencing dataset, and

(iii) displaying a two dimensional heat map that represents respective window pairs in the plurality of windows, wherein each respective window pair is displayed in the two dimensional heat map as a color selected from a color scheme based upon the number of identifiers in common in the respective window pair.

97. The method of claim 96 wherein a window pair in the two dimensional heat map represents a first window and a second window that are more than 100 kilobases apart from each other in a reference genome corresponding to the sample.

98. The method of claim 96 wherein a window pair in the two dimensional heat map represents a first window and a second window that are more than a megabase apart from each other in a reference genome corresponding to the sample.

99. The method of claim 96 wherein a window pair in the two dimensional heat map represents a first window and a second window that are more than 5 megabases apart from each other in a reference genome corresponding to the sample.

100. The method of claim 96, wherein the number of identifiers in common in the respective window pair is down-weighted to remove bar codes that are expected to be proximate to each other based on the reference genome sequence corresponding to the sample.

101. The method of claim 96, wherein the plurality of sequencing reads is characterized by a mean fragment length and wherein the number of identifiers in common in the respective window pair is down-weighted to remove bar codes that are expected to be within a multiple of the mean fragment length of each other based on the reference genome sequence corresponding to the sample.

102. The method of claim 101, wherein the multiple is a real number between 0 and 1.

103. The method of claim 101, wherein the multiple is a real number of 1 or greater.

104. The method of claim 101, wherein the multiple is a positive integer.