

SYSTEMS AND METHODS FOR USING PAIRED-END DATA IN DIRECTED ACYCLIC STRUCTURE

Cross-Reference to Related Application

This application claims the benefit of, and priority to, U.S. Patent Application Serial No. 14/157,979, filed January 17, 2014, which application claims the benefit of, and priority to, U.S. Provisional Patent Application Serial No. 61/893,467, filed October 21, 2013, the contents of which are incorporated by reference.

Sequence Listing

This application contains a sequence listing which has been submitted in ASCII format via EFS-Web and is hereby incorporated by reference in its entirety. The ASCII-formatted sequence listing, created on October 14, 2014, is named SBG-009-01WO-2-Sequences_ST25, and is 1,217 bytes in size.

Field of the Invention

The invention generally relates to bioinformatics and particularly to analyses using directed acyclic data structures.

Background

The transcriptome is the complete set of transcripts in a cell, and their quantity, for a specific developmental stage or physiological condition. The transcriptome is central to organismal function, development, and disease. The very nature and vitality of an organism arises from all of the transcripts (including mRNAs, non-coding RNAs and small RNAs), their expression levels, gene splicing patterns, and post-transcriptional modifications. In fact, a much greater fraction of the human genome than was expected is now known to be transcribed. See Bertone, et al., 2004, Global identification of human transcribed sequences with genome tiling arrays, Science 306:2242–2246.

Methods for transcriptome analysis based on next-generation sequencing (NGS) technologies have been reported. See, e.g., Wang, et al., 2009, RNA-Seq: a revolutionary tool for transcriptomics, *Nat Rev Genet* 10(1):57-63 2009. The RNA-Seq (RNA sequencing) method promises to rapidly generate high volumes of transcriptome data.

However, RNA-Seq faces informatics challenges, such as storing and analyzing large amounts of data, which must be overcome to make good use of the reads. Aligning and analyzing RNA-Seq reads presents problems not only in the nature of the information involved but also in the volume. Read mapping with large data sets is computationally challenging and analytical methods for differential expression are only beginning to emerge. See Garber, et al., 2011, Computational methods for transcriptome annotation and quantification using RNA-Seq, *Nat Meth* 8(6):469-477. The computational demands of mapping the large number of reads from RNA-Seq are greatly multiplied by the large amount of reference data that are coming available.

For example, the GENCODE Consortium seeks to identify all gene features in the human genome that have been reported. Harrow, et al., 2012, GENCODE: The reference human genome annotation for The ENCODE Project, *Genome Res* 22:1760-1774. The volume of material encompassed by the GENCODE project is formidable. Not only do new protein-coding loci continue to be added, the number of alternative splicing transcripts annotated steadily increases. The GENCODE 7 release includes more than 20,000 protein-coding and almost 10,000 long noncoding RNA loci (lncRNA), as well as more than 33,000 coding transcripts not represented in other sources. GENCODE also includes other features such as untranslated regions (UTRs), long intergenic noncoding RNA (lincRNA) genes, short noncoding RNAs, and alternative splice patterns. Even if an RNA-Seq study started with only a limited amount of new data, the volume of potential reference data from a source such as GENCODE is so large that making sense of the new data is computationally challenging.

Summary

The disclosure generally provides systems and methods for analysis of paired-end reads using a reference organized as a directed acyclic graph (DAG) or similar data structure. Features in the reference, such as exons and introns, provide nodes in the DAG and those features are linked to one another in their canonical genomic order by edges. The DAG can scale to any size and can, in fact, be populated in the first instance by imported sequences from an extrinsic

annotated reference. Where the new data includes paired end reads, each read of a pair may potentially align to any location along a genome of length L and in a complex DAG, there may be many more locations (i.e., if a DAG is, on average, n -plex across any loci, aligning a pair of paired-ends reads amounts to choosing from $(nL)^2$ options). However, insert length information provides a tool for constraining the potential alignments within a DAG—not just in distance away within a path, but also by excluding paths, limiting the number of options to be linear with respect to L . Paired-end reads may provide other information, as well. For example, paired-end reads can provide a basis for modifying a DAG, can be used to infer isoform frequency, or others. In fact, sets of paired-end reads with insert length information carry qualities of information that are well-suited to focusing DAG-based alignments and to updating DAGs to represent actually-occurring natural isoforms that may not otherwise have been discovered or noticed. Thus paired-end reads are particularly well suited for RNA-Seq projects, allowing very large sets of RNA-Seq data to be analyzed against very large reference sets. Using paired-end RNA-Seq reads to discover all isoforms of a transcriptome is not computationally intractable and transcriptomics can move at a pace and throughput not previously possible. One or any number of transcriptomes can be studied from one or any number of organisms to increase understanding of organismal function, development, and disease.

The disclosure relates to creation of a DAG that includes features (e.g., exons, introns, etc.) of a reference connected by their proximal edges in genomic order. This DAG data structure has applications in transcriptomics including, for example, identifying isoforms. Such a DAG reference is used in transcriptomics involving RNA-Seq data that includes paired-reads. RNA-Seq data may be analyzed using the DAG reference that includes features as nodes (e.g., created from an annotated source genome or transcriptome). Systems and methods of the invention provide for improved alignments using paired-end reads, building the DAG using paired-end read information, inferring isoform probabilities or frequencies, or combinations thereof.

In one example, there is disclosed a method of analyzing a transcriptome, the method comprising obtaining a pair of paired-end reads from a transcriptome; finding an alignment with an optimal score between a first read of the pair and a node in a directed acyclic data structure, the data structure comprising nodes representing RNA sequences and edges connecting pairs of the nodes; identifying candidate paths that include the node connected to a

downstream node by a path having a length substantially similar to an insert length of the pair of paired-end reads; and aligning the paired-end reads to the candidate paths to determine an optimal-scoring alignment. In one aspect, the present invention refers to a method of analyzing a transcriptome, the method comprising obtaining a pair of paired-end reads from a transcriptome; finding an alignment with an optimal score between a first read of the pair and a node representing an exon in a directed acyclic graph (DAG), the DAG comprising nodes representing a plurality of exons and introns from features in an annotated transcriptome database, and edges connecting pairs of the nodes in their canonical genomic order; identifying candidate paths within the DAG that include the node connected to a downstream node by a path having a length substantially similar to an insert length of the pair of paired-end reads; aligning the paired-end reads to the candidate paths to determine an optimal-scoring alignment by excluding any paths that are not candidate paths from any alignment calculations involving the pair of paired-end reads; and identifying an isoform from the transcriptome using the optimal-scoring alignment of the paired-end reads. In another aspect, the present invention refers to a system for transcriptomics, the system comprising a computer system comprising a processor coupled to a non-transitory memory having stored therein a directed acyclic graph (DAG), the DAG comprising nodes representing a plurality of exons and introns from features in an annotated transcriptome database, and edges connecting pairs of the nodes in their canonical genomic order, wherein the system is operable to obtain a pair of paired-end reads from a cDNA fragment; find an alignment with an optimal score between a first read of the pair and a node in the DAG; identify candidate paths within the DAG that include the node connected to a downstream node by a path having a length substantially similar to an insert length of the paired end reads; align the paired-end reads to the candidate paths to determine an optimal-scoring alignment by excluding any paths that are not candidate paths from any alignment calculations involving the pair of paired-end nodes; and identify an isoform from the transcriptome using the optimal scoring-alignment of the paired-end reads. In yet another aspect, the present invention refers to a method of inferring isoform frequency in a transcriptome, the method comprising aligning a plurality of pairs of paired-end reads from cDNA to a reference directed acyclic graph (DAG) comprising a plurality of nodes representing exons and introns from annotated features in a genetic feature database, wherein pairs of the plurality of nodes are connected in their canonical genomic order at their proximal ends by an edge, wherein the aligning for a pair of paired-end reads comprises identifying an exon at which a first member of a pair aligns, determining a set of distances that accounts for an expected insert-size value of the pair, using the determined set of distances to select downstream nodes as potential alignment sites for a second member of the pair, identifying paths that begin with the upstream node and end with one of the downstream nodes as candidate paths, and excluding any paths that are not candidate paths from any alignment calculations involving the pair of paired-end reads; determining distances between the aligned pairs of the paired-end reads and the

frequencies of the distances between the aligned pairs; and determining a set of isoform paths and isoform frequencies such that representing the isoform paths through the DAG at the isoform frequencies results in pairs of features being included in the isoform paths at the frequencies of the distances between the aligned pairs. In a further aspect, the present invention refers to a system of inferring isoform frequency in a transcriptome, the system comprising a computer system comprising a processor coupled to a non-transitory memory having stored therein a reference DAG comprising a plurality of nodes representing exons and introns from features in a genetic feature database, wherein pairs of the plurality of nodes are connected in their canonical genomic order at their proximal ends by an edge, wherein the system is operable to align a plurality of pairs of paired-end reads to the reference DAG, wherein the aligning for a pair of paired-end reads comprises identifying an exon at which a first member of a pair aligns, determining a set of distances that accounts for an expected insert-size value of the pair, using the determined set of distances to select downstream nodes as potential alignment sites for a second member of the pair, identifying paths that begin with the upstream node and end with one of the downstream nodes as candidate paths, and excluding any paths that are not candidate paths from any alignment calculations involving the pair of paired-end reads; determine distances between the plurality of pairs of paired-end reads from cDNA and the frequencies of the distances between the pairs; and determine a set of isoform paths and isoform frequencies such that representing the isoform paths through the DAG at the isoform frequencies results in pairs of features being included in the isoform paths at the frequencies of the distances between the aligned pairs. In certain examples, the disclosure describes methods of analyzing a transcriptome that involves obtaining at least one pair of paired-end reads from a transcriptome from an organism, finding an alignment with an optimal score between a first read of the pair and a node in a directed acyclic data structure (the data structure has nodes representing RNA sequences such as exons or transcripts and edges connecting pairs of nodes), identifying candidate paths that

include the node connected to a downstream node by a path having a length substantially similar to an insert length of the pair of paired-end reads, and aligning the paired-end reads to the candidate paths to determine an optimal-scoring alignment. The method may include excluding any paths that are not candidate paths from any alignment calculations involving the pair of paired-end reads.

The method is suited for any number of paired-end reads. The node and the downstream node may represent a pair of exons from which the pair of paired-end reads were obtained. The method may include identifying an isoform based on the determined optimal-scoring alignment. In some examples, the identified isoform is a novel isoform. The directed acyclic data structure may be updated to represent the novel isoform (e.g., by adding at least one new node or edge).

The disclosure describes a computer system for transcriptomics. The system has a processor coupled to a non-transitory memory and is used to store a directed acyclic data structure having nodes representing RNA sequences and edges connecting pairs of the nodes. The system is operable to obtain a pair of paired-end reads from a transcriptome, find an alignment with an optimal score between a first read of the pair and a node in the data structure, identify candidate paths that include the node connected to a downstream node by a path having a length substantially similar to an insert length of the pair of paired-end reads, and align the paired-end reads to the candidate paths to determine an optimal-scoring alignment. Preferably, the system excludes all non-candidate paths from subsequent alignment attempts. In certain examples, the system is further operable to obtain a plurality of paired-end reads from cDNA fragments. Preferably the system is operable to determine optimal-scoring alignments for each of the plurality of pairs of paired-end reads. The cDNA fragments may represent a single transcriptome. The system may identify an isoform, such as a novel isoform, based on the determined optimal-scoring alignment. The system may be operable to update the directed acyclic data structure to represent the novel isoform. In some examples, updating the directed acyclic data structure to represent the novel isoform comprises adding at least one new node. A node and a downstream node may represent a pair of exons from which the pair of paired-end reads were obtained. The system may be operable to exclude any paths that are not candidate paths from any alignment calculations involving the pair of paired-end reads.

Other examples of the invention provide for expression analysis through, e.g., systems and methods useful for inferring isoform frequency.

In certain examples, disclosed herein are methods of inferring isoform frequency in a transcriptome. The method includes obtaining a plurality of pairs of paired-end reads from a transcriptome of an organism and optionally aligning those reads to a reference data structure. The data structure includes a plurality of features from a genome each represented as a node. Pairs of the nodes, or features, are connected at their proximal ends by an edge (i.e., within pairs, the two pair-members are connected to one another at their proximal ends on the genome by an edge; since any pair-member may also be a member of other pairs, either or both ends of a feature may be connected by one or any number of edges to other features). Distances between the aligned pairs of the paired-end reads and the frequencies of the distances between the aligned pairs are determined. A set of isoform paths and isoform frequencies are determined such that representing the isoform paths through the structure at the isoform frequencies results in pairs of features being included in the isoform paths at the frequencies of the distances the aligned pairs. Each path of the determined set of isoform paths may represent a transcript isoform present in the organism. The distances between the aligned pairs may be path lengths from an upstream end of a first aligned member of each pair to a downstream end of a second aligned member of that pair. In certain examples, determining the set of isoform path includes discovering a novel isoform path. The reference data structure may be updated to include the novel isoform path.

Determining the set of isoform paths can be performed algebraically, e.g., by solving a set of linear equations. In some examples, one of the isoform frequencies is obtained from biological data prior to solving the set of linear equations (e.g., so remaining variables in the linear equations can be determined). Additionally or alternatively, the frequencies of the distances between the aligned pairs may be normalized prior to determining a set of isoform paths and isoform frequencies.

The data structure may be built from an existing genetic feature database. For example, each node may represent a feature from such a database.

In one example, there is disclosed a system for transcriptomics, the system comprising a computer system comprising a processor coupled to a non-transitory memory having stored therein a directed acyclic data structure, the data structure comprising nodes

representing RNA sequences and edges connecting pairs of the nodes, wherein the system is operable to obtain a pair of paired-end reads from a cDNA fragment; find an alignment with an optimal score between a first read of the pair and a node in the data structure; identify candidate paths that include the node connected to a downstream node by a path having a length substantially similar to an insert length of the paired end reads; and align the paired-end reads to the candidate paths to determine an optimal-scoring alignment. In related examples, disclosed herein is a computer system of inferring isoform frequency in a transcriptome. The system includes a processor coupled to a non-transitory memory and stores a reference data structure comprising a plurality of features from a genome each represented as a node wherein pairs of the plurality of features are connected at their

proximal ends by an edge. The system can be used to obtain or determine distances between a plurality of pairs of paired-end reads from a transcriptome and the frequencies of the distances between the pairs. A set of isoform paths and isoform frequencies can then be determined such that representing the isoform paths through the structure at the isoform frequencies results in pairs of features being included in the isoform paths at the frequencies of the distances the aligned pairs. The features may include exons and the system is further operable to align a plurality of pairs of paired-end reads to the reference data structure. Each path of the determined set of isoform paths may represent a transcript isoform present in the organism. Determining the set of isoform path may include discovering a novel isoform path. The system may update the reference data structure to include the novel isoform path. The distances between the aligned pairs may comprise path lengths from an upstream end of a first aligned member of each pair to a downstream end of a second aligned member of that pair. In some examples, determining the set of isoform paths comprises solving a set of linear equations. In certain examples, determining the set of isoform paths comprises obtaining one of the isoform frequencies from biological data prior to solving the set of linear equations. The system may be operable to normalize the frequencies of the distances between the aligned pairs prior to determining the set of isoform paths and isoform frequencies. The data structure may include a node representing each annotated feature in a remote genetic feature database.

Brief Description of the Drawings

FIG. 1 shows a pictorial representation of the read being compared to the DAG.

FIG. 2 shows the actual matrices that correspond to the comparison.

FIG. 3 provides a diagram of methods of the invention.

FIG. 4 diagrams a hypothetical DAG and pair of paired-end reads.

FIG. 5 shows candidate paths through the hypothetical DAG.

FIG. 6 shows the candidate paths from the hypothetical DAG.

FIG. 7 presents three alternative histories involving sequences.

FIG. 8 depicts an alignment according to the invention.

FIG. 9 shows an isoform frequency distribution.

FIG. 10 illustrates a system for implementing methods of the invention.

Detailed Description

It is understood that genes do not in general code for single, specific proteins. Instead, different subsets of a given gene's exons are transcribed in different proportions at different times. The different proteins created by these different transcripts are called isoforms, and by extension the different transcripts themselves are also called isoforms. As discussed above, contemporary bioinformatics has provided at least partially annotated transcriptomes, potentially useful as references. For example, the Gencode project has catalogued not only genes but exons (along with other regions) in the human genome.

Systems and methods described herein are used to efficiently exploit the information in those annotated transcriptomes. Indeed, the invention includes the insight that comparison of paired-end reads to an annotated transcriptome is ideally suited to analytical approaches in which sequence data is stored as a DAG-based data structure.

Embodiments of the invention provide methods of creating a directed acyclic data structure with the data in an annotated transcriptome. In such a data structure, features (e.g., exons and introns) from the annotated transcriptome are represented as nodes, which are connected by edges. An isoform from the real-world transcriptome is thus represented by a corresponding path through the DAG-based data structure.

Aspects of the invention relate to creation of a DAG that includes features such as introns and exons from one or more known references. A DAG is understood in the art to refer to data that can be presented as a graph as well as to a graph that presents that data. The invention provides methods for storing a DAG as data that can be read by a computer system for bioinformatic processing or for presentation as a graph. A DAG can be saved in any suitable format including, for example, a list of nodes and edges, a matrix or a table representing a matrix, an array of arrays or similar variable structure representing a matrix, in a language built with syntax for graphs, in a general markup language purposed for a graph, or others.

In some embodiments, a DAG is stored as a list of nodes and edges. One such way is to create a text file that includes all nodes, with an ID assigned to each node, and all edges, each with the node ID of starting and ending node. Thus, for example, were a DAG to be created for two sentences, "See Jane run," and "Run, Jane run," a case-insensitive text file could be created. Any suitable format could be used. For example, the text file could include comma-separated values. Naming this DAG "Jane" for future reference, in this format, the DAG "Jane" may read

as follows: 1 see, 2 run, 3 jane, 4 run, 1-3, 2-3, 3-4. One of skill in the art will appreciate that this structure is applicable to the introns and exons represented in FIGS. 7 and 8, and the accompanying discussion below.

In certain embodiments, a DAG is stored as a table representing a matrix (or an array of arrays or similar variable structure representing a matrix) in which the (i,j) entry in the N×N matrix denotes that node i and node j are connected (where N is a vector containing the nodes in genomic order). For the DAG to be acyclic simply requires that all non-zero entries be above the diagonal (assuming nodes are represented in genomic order). In a binary case, a 0 entry represents that no edge exists from node i to node j, and a 1 entry represents an edge from i to j. One of skill in the art will appreciate that a matrix structure allows values other than 0 to 1 to be associated with an edge. For example, any entry may be a numerical value indicating a weight, or a number of times used, reflecting some natural quality of observed data in the world. A matrix can be written to a text file as a table or a linear series of rows (e.g., row 1 first, followed by a separator, etc.), thus providing a simple serialization structure.

One useful way to serialize a matrix DAG would be to use comma-separated values for the entries, after defining the nodes. Using this format, the DAG “Jane” would include the same node definitions as for above, followed by matrix entries. This format could read as:

```
1 see, 2 run, 3 jane, 4 run
```

```
.,1,\,,1,\,,1\,,
```

where entries of zero (0) are simply omitted and ‘\’ is a newline character.

Embodiments of the invention include storing a DAG in a language built with syntax for graphs. For example, The DOT Language provided with the graph visualization software packages known as Graphviz provides a data structure that can be used to store a DAG with auxiliary information and that can be converted into graphic file formats using a number of tools available from the Graphviz web site. Graphviz is open source graph visualization software. Graph visualization is a way of representing structural information as diagrams of abstract graphs and networks. It has applications in networking, bioinformatics, software engineering, database and web design, machine learning, and in visual interfaces for other technical domains. The Graphviz layout programs take descriptions of graphs in a simple text language, and make diagrams in useful formats, such as images and SVG for web pages; PDF or Postscript for inclusion in other documents; or display in an interactive graph browser.

In related embodiments, a DAG is stored in a general markup language purposed for a graph, or others. Following the descriptions of a linear text file, or a comma-separated matrix, above, one of skill in the art will recognize that a language such as XML can be used (extended) to create labels (markup) defining nodes and their headers or IDs, edges, weights, etc. However a DAG is structured and stored, embodiments of the invention involve using nodes to represent features such as exons and introns. This provides a useful tool for analyzing RNA-Seq reads and discovering, identifying, and representing isoforms.

If nodes represent features or fragments of features, then isoforms can be represented by paths through those fragments. An exon's being "skipped" is represented by an edge connecting some previous exon to some later one. Presented herein are techniques for constructing a DAG to represent alternative splicing or isoforms of genes. Alternative splicing is discussed in Lee and Wang, 2005, Bioinformatics analysis of alternative splicing, *Brief Bioinf* 6(1):23-33; Heber, et al., 2002, Splicing graphs and EST assembly problems, *Bioinformatics* 18Suppl:s181-188; Leipzig, et al., 2004, The alternative splicing gallery (ASG): Bridging the gap between genome and transcriptome, *Nucl Ac Res* 32(13):3977-2983; and LeGault and Dewey, 2013, Inference of alternative splicing from RNA-Seq data with probabilistic splice graphs, *Bioinformatics* 29(18):2300-2310, the contents of each of which are incorporated by reference. Additional discussion may be found in Florea, et al., 2005, Gene and alternative splicing annotation with AIR, *Genome Research* 15:54–66; Kim, et al., 2005, ECgene: Genome-based EST clustering and gene modeling for alternative splicing, *Genome Research* 15:566–576; and Xing, et al., 2006, An expectation-maximization algorithm for probabilistic reconstructions of full-length isoforms from splice graphs, *Nucleic Acids Research*, 34, 3150–3160, the contents of each of which are incorporated by reference.

1. Using paired-end reads

The invention relates to transcriptomics involving RNA-Seq data that includes paired-reads. Systems and methods of the invention provide for improved alignments using paired-end reads. In certain embodiments, paired-end reads are used to constrain alignments, thereby greatly improving computational cost of an RNA-Seq project.

Much sequencing data is paired-end data. There exist ways in which features of such data ought to manifest themselves in alignments: the distance between the locations to which paired-

end mates are assigned in alignment ought approximately to correspond to the insert size given by the sequencing run. In general, information given by insert sizes is related to the information that alignment algorithms attempt to uncover.

Several methods are known in the art for producing paired nucleic acid reads. Typically, the methods involve some form of selectively fragmenting a nucleic acid sample, resulting in sequences of a known size distribution. The ends of the fragmented samples are typically protected with a functional group or by self-cyclizing, whereupon the remaining nucleic acid sample material is removed. The ends are then amplified and sequenced, resulting in a population of nucleic acid reads that are known to have originated some distance from each other.

Mate pair sample preparation kits are commercially available, for example from Nextera Illumina (San Diego, CA). To prepare paired nucleic acids for sequencing on such a platform, a nucleic acid sample is fragmented into segments between 2 and 5 kb in length. The fragment ends are then biotinylated to facilitate later recovery using affinity column separation. The biotinylated ends are joined to create a circular DNA segment. At this point, non-circularized nucleic acids from the sample are digested and subsequently removed. The remaining circularized DNA is then re-fragmented to yield fragments suitable for clustering and sequencing. The original ends, which are biotinylated, can be removed from the remaining fragments of the circularized DNA with an affinity column or streptavidin coated beads. The recovered ends are then end repaired, poly-A tailed, and sequencing adapters are added prior to amplification and sequencing.

2. Alignment.

The invention provides methods and systems for aligning sequence reads to a directed acyclic data structure that represents an annotated reference. Using alignment algorithms of the invention, reads can be rapidly mapped despite the large numbers associated with RNA-Seq results and substantially comprehensive references. As in other situations, DAGs are ideal objects against which to align. Performing alignments against DAGs constructed out of annotated transcriptomes is much more effective than using the annotated transcriptomes later in one's workflow.

Numerous benefits obtain by using a DAG as a reference. For example, aligning against a DAG is more accurate than aligning against one reference and then attempting to adjust one's results in light of an annotated transcriptome. This is primarily because the latter approach enforces an unnatural asymmetry between the sequence used in the initial alignment and the other sequences represented in the transcriptome. Aligning against an object that potentially represents all the relevant physical possibilities is much more computationally efficient than attempting to align against a linear sequence for each physical possibility (the number of such possibilities will generally be exponential in the number of junctions).

Embodiments of the invention include aligning one or more reads against a reference data structure. That alignment can include a DAG-based alignment that finds the edges and nodes of the data structure that best represent the read. That alignment can also include instances of aligning a read to an exon from a node of the DAG in a pairwise fashion.

Pairwise alignment generally involves placing one sequence along part of target, introducing gaps according to an algorithm, scoring how well the two sequences match, and preferably repeating for various position along the reference. The best-scoring match is deemed to be the alignment and represents an inference about what the sequence data represents. In some embodiments, scoring an alignment of a pair of nucleic acid sequences involves setting values for the probabilities of substitutions and indels. When individual bases are aligned, a match or mismatch contributes to the alignment score by a substitution probability, which could be, for example, 1 for a match and -0.33 for a mismatch. An indel deducts from an alignment score by a gap penalty, which could be, for example, -1. Gap penalties and substitution probabilities can be based on empirical knowledge or *a priori* assumptions about how sequences evolve. Their values affects the resulting alignment. Particularly, the relationship between the gap penalties and substitution probabilities influences whether substitutions or indels will be favored in the resulting alignment.

Stated formally, an alignment represents an inferred relationship between two sequences, x and y . For example, in some embodiments, an alignment A of sequences x and y maps x and y respectively to another two strings x' and y' that may contain spaces such that: (i) $|x'|=|y'|$; (ii) removing spaces from x' and y' should get back x and y , respectively; and (iii) for any i , $x'[i]$ and $y'[i]$ cannot be both spaces.

A gap is a maximal substring of contiguous spaces in either x' or y' . An alignment A can include the following three kinds of regions: (i) matched pair (e.g., $x'[i]=y'[i]$); (ii) mismatched pair, (e.g., $x'[i]\neq y'[i]$ and both are not spaces); or (iii) gap (e.g., either $x'[i..j]$ or $y'[i..j]$ is a gap). In certain embodiments, only a matched pair has a high positive score a . In some embodiments, a mismatched pair generally has a negative score b and a gap of length r also has a negative score $g+rs$ where $g, s < 0$. For DNA, one common scoring scheme (e.g. used by BLAST) makes score $a=1$, score $b=-3$, $g=-5$ and $s=-2$. The score of the alignment A is the sum of the scores for all matched pairs, mismatched pairs and gaps. The alignment score of x and y can be defined as the maximum score among all possible alignments of x and y .

In some embodiments, any pair has a score a defined by a 4×4 matrix B of substitution probabilities. For example, $B(i,i)=1$ and $0 < B(i,j)_{i \neq j} < 1$ is one possible scoring system. For instance, where a transition is thought to be more biologically probable than a transversion, matrix B could include $B(C,T)=.7$ and $B(A,T)=.3$, or any other set of values desired or determined by methods known in the art.

Alignment according to some embodiments of the invention includes pairwise alignment. A pairwise alignment, generally, involves—for sequence Q (query) having m characters and a reference genome T (target) of n characters—finding and evaluating possible local alignments between Q and T . For any $1 \leq i \leq n$ and $1 \leq j \leq m$, the largest possible alignment score of $T[h..i]$ and $Q[k..j]$, where $h \leq i$ and $k \leq j$, is computed (i.e. the best alignment score of any substring of T ending at position i and any substring of Q ending at position j). This can include examining all substrings with cm characters, where c is a constant depending on a similarity model, and aligning each substring separately with Q . Each alignment is scored, and the alignment with the preferred score is accepted as the alignment. One of skill in the art will appreciate that there are exact and approximate algorithms for sequence alignment. Exact algorithms will find the highest scoring alignment, but can be computationally expensive. Two well-known exact algorithms are Needleman-Wunsch (J Mol Biol, 48(3):443-453, 1970) and Smith-Waterman (J Mol Biol, 147(1):195-197, 1981; Adv. in Math. 20(3), 367-387, 1976). A further improvement to Smith-Waterman by Gotoh (J Mol Biol, 162(3), 705-708, 1982) reduces the calculation time from $O(m^2n)$ to $O(mn)$ where m and n are the sequence sizes being compared and is more amenable to parallel processing. In the field of bioinformatics, it is Gotoh's modified algorithm that is often referred to as the Smith-Waterman algorithm. Smith-Waterman approaches are being used

to align larger sequence sets against larger reference sequences as parallel computing resources become more widely and cheaply available. See, e.g., Amazon's cloud computing resources. All of the journal articles referenced herein are incorporated by reference in their entirety.

The Smith-Waterman (SW) algorithm aligns linear sequences by rewarding overlap between bases in the sequences, and penalizing gaps between the sequences. Smith-Waterman also differs from Needleman-Wunsch, in that SW does not require the shorter sequence to span the string of letters describing the longer sequence. That is, SW does not assume that one sequence is a read of the entirety of the other sequence. Furthermore, because SW is not obligated to find an alignment that stretches across the entire length of the strings, a local alignment can begin and end anywhere within the two sequences.

In some embodiments, pairwise alignment proceeds according to dot-matrix methods, dynamic programming methods, or word methods. Dynamic programming methods generally implement the Smith-Waterman (SW) algorithm or the Needleman-Wunsch (NW) algorithm. Alignment according to the NW algorithm generally scores aligned characters according to a similarity matrix $S(a,b)$ (e.g., such as the aforementioned matrix B) with a linear gap penalty d . Matrix $S(a,b)$ generally supplies substitution probabilities. The SW algorithm is similar to the NW algorithm, but any negative scoring matrix cells are set to zero. The SW and NW algorithms, and implementations thereof, are described in more detail in U.S. Pat. 5,701,256 and U.S. Pub. 2009/0119313, both herein incorporated by reference in their entirety.

An alignment program that implements a version of the Smith-Waterman algorithm is MUMmer, available from the SourceForge web site maintained by Geeknet (Fairfax, VA). MUMmer is a system for rapidly aligning genome-scale sequences (Kurtz, S., et al., *Genome Biology*, 5:R12 (2004); Delcher, A.L., et al., *Nucl. Acids Res.*, 27:11 (1999)). For example, MUMmer 3.0 can find all 20-basepair or longer exact matches between a pair of 5-megabase genomes in 13.7 seconds, using 78 MB of memory, on a 2.4 GHz Linux desktop computer. MUMmer can handle the 100s or 1000s of contigs from a shotgun sequencing project, and will align them to another set of contigs or a reference using the NUCmer program included with the system. If the species are too divergent for a DNA sequence alignment to detect similarity, then the PROmer program can generate alignments based upon the six-frame translations of both input sequences.

Other exemplary alignment programs include: Efficient Large-Scale Alignment of Nucleotide Databases (ELAND) or the ELANDv2 component of the Consensus Assessment of Sequence and Variation (CASAVA) software (Illumina, San Diego, CA); RTG Investigator from Real Time Genomics, Inc. (San Francisco, CA); Novoalign from Novocraft (Selangor, Malaysia); Exonerate, European Bioinformatics Institute (Hinxton, UK) (Slater, G., and Birney, E., BMC Bioinformatics 6:31 (2005)), Clustal Omega, from University College Dublin (Dublin, Ireland) (Sievers F., et al., Mol Syst Biol 7, article 539 (2011)); ClustalW or ClustalX from University College Dublin (Dublin, Ireland) (Larkin M.A., et al., Bioinformatics, 23, 2947-2948 (2007)); and FASTA, European Bioinformatics Institute (Hinxton, UK) (Pearson W.R., et al., PNAS 85(8):2444-8 (1988); Lipman, D.J., Science 227(4693):1435-41 (1985)).

As discussed above, it may be preferable or desirable to implement the SW alignment algorithm or a modified version of (discussed in greater detail below) when aligning RNA-Seq reads to a direct acyclic annotated reference genome.

The SW algorithm is easily expressed for an $n \times m$ matrix H , representing the two strings of length n and m , in terms of equation (1) below:

$$\begin{aligned}
 H_{k0} &= H_{0l} = 0 \quad (\text{for } 0 \leq k \leq n \text{ and } 0 \leq l \leq m) \\
 H_{ij} &= \max\{H_{i-1,j-1} + s(a_i, b_j), H_{i-1,j} - W_{in}, H_{i,j-1} - W_{del}, 0\} \\
 &\quad (\text{for } 1 \leq i \leq n \text{ and } 1 \leq j \leq m)
 \end{aligned} \tag{1}$$

In the equations above, $s(a_i, b_j)$ represents either a match bonus (when $a_i = b_j$) or a mismatch penalty (when $a_i \neq b_j$), and insertions and deletions are given the penalties W_{in} and W_{del} , respectively. In most instance, the resulting matrix has many elements that are zero. This representation makes it easier to backtrack from high-to-low, right-to-left in the matrix, thus identifying the alignment.

Once the matrix has been fully populated with scores, the SW algorithm performs a backtrack to determine the alignment. Starting with the maximum value in the matrix, the algorithm will backtrack based on which of the three values ($H_{i-1,j-1}$, $H_{i-1,j}$, or $H_{i,j-1}$) was used to compute the final maximum value for each cell. The backtracking stops when a zero is reached. The optimal-scoring alignment may contain greater than the minimum possible number of

insertions and deletions, while containing far fewer than the maximum possible number of substitutions.

When applied as SW or SW-Gotoh, the techniques use a dynamic programming algorithm to perform local sequence alignment of the two strings, S and A , of sizes m and n , respectively. This dynamic programming technique employs tables or matrices to preserve match scores and avoid re-computation for successive cells. Each element of the string can be indexed with respect to a letter of the sequence, that is, if S is the string ATCGAA, $S[1] = A$.

Instead of representing the optimum alignment as $H_{i,j}$ (above), the optimum alignment can be represented as $B[j,k]$ in equation (2) below:

$$B[j, k] = \max(p[j, k], i[j, k], d[j, k], 0) \quad (\text{for } 0 < j \leq m, 0 < k \leq n) \quad (2)$$

The arguments of the maximum function, $B[j,k]$, are outlined in equations (3)-(5) below, wherein MISMATCH_PENALTY, MATCH_BONUS, INSERTION_PENALTY, DELETION_PENALTY, and OPENING_PENALTY are all constants, and all negative except for MATCH_BONUS. The match argument, $p[j,k]$, is given by equation (3), below:

$$\begin{aligned} p[j,k] &= \max(p[j-1,k-1], i[j-1,k-1], d[j-1,k-1]) + \text{MISMATCH_PENALTY, if } S[j] \neq A[k] \\ &= \max(p[j-1,k-1], i[j-1,k-1], d[j-1,k-1]) + \text{MATCH_BONUS, if } S[j] = A[k] \end{aligned} \quad (3)$$

the insertion argument $i[j,k]$, is given by equation (4), below:

$$\begin{aligned} i[j,k] &= \max(p[j-1,k] + \text{OPENING_PENALTY}, i[j-1,k], d[j-1,k] + \\ &\quad \text{OPENING_PENALTY}) + \text{INSERTION_PENALTY} \end{aligned} \quad (4)$$

and the deletion argument $d[j,k]$, is given by equation (5), below:

$$\begin{aligned} d[j,k] &= \max(p[j,k-1] + \text{OPENING_PENALTY}, i[j,k-1] + \\ &\quad \text{OPENING_PENALTY}, d[j,k-1]) + \text{DELETION_PENALTY} \end{aligned} \quad (5)$$

For all three arguments, the $[0,0]$ element is set to zero to assure that the backtrack goes to completion, i.e., $p[0,0] = i[0,0] = d[0,0] = 0$.

The scoring parameters are somewhat arbitrary, and can be adjusted to achieve the

behavior of the computations. One example of the scoring parameter settings (Huang, Chapter 3: *Bio-Sequence Comparison and Alignment*, ser. *Curr Top Comp Mol Biol*. Cambridge, Mass.: The MIT Press, 2002) for DNA would be:

MATCH_BONUS: 10

MISMATCH_PENALTY: -20

INSERTION_PENALTY: -40

OPENING_PENALTY: -10

DELETION_PENALTY: -5

The relationship between the gap penalties (INSERTION_PENALTY, OPENING_PENALTY) above help limit the number of gap openings, i.e., favor grouping gaps together, by setting the gap insertion penalty higher than the gap opening cost. Of course, alternative relationships between MISMATCH_PENALTY, MATCH_BONUS, INSERTION_PENALTY, OPENING_PENALTY and DELETION_PENALTY are possible.

In some embodiments, the methods and systems of the invention incorporate multi-dimensional alignment algorithms. Multi-dimensional algorithms of the invention provide for a “look-back” type analysis of sequence information (as in Smith-Waterman), wherein the look back is conducted through a multi-dimensional space that includes multiple pathways and multiple nodes. The multi-dimensional algorithm can be used to align RNA-Seq reads against the DAG-type reference. That alignment algorithm identifies the maximum value for $C_{i,j}$ by identifying the maximum score with respect to each sequence contained at a position on the DAG (e.g., the reference sequence construct). In fact, by looking “backwards” at the preceding positions, it is possible to identify the optimum alignment across a plurality of possible paths.

The algorithm of the invention is carried out on a read (a.k.a. “string”) and a directed acyclic graph (DAG), discussed above. For the purpose of defining the algorithm, let S be the string being aligned, and let D be the directed acyclic graph to which S is being aligned. The elements of the string, S , are bracketed with indices beginning at 1. Thus, if S is the string ATCGAA, $S[1] = A$, $S[4] = G$, etc.

In certain embodiments, for the DAG, each letter of the sequence of a node will be represented as a separate element, d . A *predecessor* of d is defined as:

- (i) If d is not the first letter of the sequence of its node, the letter preceding d in its node is its (only) predecessor;

(ii) If d is the first letter of the sequence of its node, the last letter of the sequence of any node (e.g., all exons upstream in the genome) that is a parent of d 's node is a predecessor of d .

The set of all predecessors is, in turn, represented as $P[d]$.

In order to find the “best” alignment, the algorithm seeks the value of $M[j, d]$, the score of the optimal alignment of the first j elements of S with the portion of the DAG preceding (and including) d . This step is similar to finding H_{ij} in equation 1 above. Specifically, determining $M[j, d]$ involves finding the maximum of a , i , e , and 0, as defined below:

$$M[j, d] = \max\{a, i, e, 0\} \quad (6)$$

where

$$e = \max\{M[j, p^*] + \text{DELETE_PENALTY}\} \text{ for } p^* \text{ in } P[d]$$

$$i = M[j-1, d] + \text{INSERT_PENALTY}$$

$$a = \max\{M[j-1, p^*] + \text{MATCH_SCORE}\} \text{ for } p^* \text{ in } P[d], \text{ if } S[j] = d;$$

$$\max\{M[j-1, p^*] + \text{MISMATCH_PENALTY}\} \text{ for } p^* \text{ in } P[d], \text{ if } S[j] \neq d$$

As described above, e is the highest of the alignments of the first j characters of S with the portions of the DAG up to, but not including, d , plus an additional DELETE_PENALTY. Accordingly, if d is not the first letter of the sequence of the node, then there is only one predecessor, p , and the alignment score of the first j characters of S with the DAG (up-to-and-including p) is equivalent to $M[j, p] + \text{DELETE_PENALTY}$. In the instance where d is the first letter of the sequence of its node, there can be multiple possible predecessors, and because the DELETE_PENALTY is constant, maximizing $[M[j, p^*] + \text{DELETE_PENALTY}]$ is the same as choosing the predecessor with the highest alignment score with the first j characters of S .

In equation (6), i is the alignment of the first $j-1$ characters of the string S with the DAG up-to-and-including d , plus an INSERT_PENALTY, which is similar to the definition of the insertion argument in SW (see equation 1).

Additionally, a is the highest of the alignments of the first j characters of S with the portions of the DAG up to, but not including d , plus either a MATCH_SCORE (if the j th character of S is the same as the character d) or a MISMATCH_PENALTY (if the j th character of S is not the same as the character d). As with e , this means that if d is not the first letter of the sequence of its node, then there is only one predecessor, i.e., p . That means a is the alignment score of the first $j-1$ characters of S with the DAG (up-to-and-including p), i.e., $M[j-1, p]$, with

either a MISMATCH_PENALTY or MATCH_SCORE added, depending upon whether d and the j th character of S match. In the instance where d is the first letter of the sequence of its node, there can be multiple possible predecessors. In this case, maximizing $\{M[j, p^*] + \text{MISMATCH_PENALTY or MATCH_SCORE}\}$ is the same as choosing the predecessor with the highest alignment score with the first $j-1$ characters of S (i.e., the highest of the candidate $M[j-1, p^*]$ arguments) and adding either a MISMATCH_PENALTY or a MATCH_SCORE depending on whether d and the j th character of S match.

Again, as in the SW algorithm, the penalties, e.g., DELETE_PENALTY, INSERT_PENALTY, MATCH_SCORE and MISMATCH_PENALTY, can be adjusted to encourage alignment with fewer gaps, etc.

As described in the equations above, the algorithm finds the optimal (i.e., maximum) value for each read by calculating not only the insertion, deletion, and match scores for that element, but looking backward (against the direction of the DAG) to any prior nodes on the DAG to find a maximum score. Thus, the algorithm is able to traverse the different paths through the DAG, which contain the known mutations. Because the graphs are directed, the backtracks, which move against the direction of the graph, follow the preferred isoform toward the origin of the graph, and the optimal alignment score identifies the most likely alignment within a high degree of certainty.

FIG. 1 shows a pictorial representation of the read being compared to the DAG. The top area of FIG. 1 presents a sequence read “ATCGAA” as well as a reference sequence TTGGATATGGG (SEQ ID NO: 1) and a known insertion event TTGGATTCGAATTTATGGG (SEQ ID NO: 2), where the insertion is underlined. The middle of FIG. 1 shows a relationship among the read, SEQ ID NO:1, and SEQ ID NO: 2, keeping each sequence linear to aid in visualizing the relationship. The bottom area of FIG. 1 shows the alignment constructed using a DAG. In the depicted DAG, SEQ ID NOS. 1 and 2 can both be read by reading from the 5’ end of the DAG to the 3’ end of the DAG, albeit along different paths. The sequence read is shown as aligning to the upper path as depicted.

FIG. 2 shows the actual matrices that correspond to the comparison. Like the Smith-Waterman technique, the illustrated algorithm of the invention identifies the highest score and performs a backtrack to identify the proper location of the read. FIGS. 1 and 2 also highlight that the invention produces an actual match for the string against the construct. In the instances where

the sequence reads include variants that were not included in the DAG, the aligned sequence will be reported out with a gap, insertion, etc.

Aligning RNA-Seq reads to a DAG of an annotated reference allows one to discover isoforms in a transcriptome. Additionally, the possibility that different forms of alternative splicing have occurred (e.g., the possibility that one codon has been skipped in one of the relevant exons) can be anticipated “for free.”

Exons can be automatically separated into two nodes, one corresponding to the first codon and one corresponding to the rest of the exon. This allows both kinds of splicing to be represented in the DAG even if they have not yet been observed. Anticipating these situations can be useful, given that one codon’s worth of differences to an alignment will often not punish a candidate alignment enough to induce us to look for alternatives. (The largest such penalty would be that for three insertions or three deletions, and often the penalty will be smaller due to chance.)

Representing the transcriptome as a DAG permits a natural interpretation of an isoform as a path through the DAG. As is common when one represents a physical item in a mathematical structure, the benefits are numerous. For example, systems and methods of the invention allow sequences to be visualized in an intuitive way. Sequences can be compared, for example, by showing two paths through a DAG.

Additionally, methods of the invention are conducive to exploiting graph theorems and algorithms. The analysis of paths through DAGs is a vibrant research area in modern mathematics and computer science. The ability to exploit this research to improve the study of the transcriptome is therefore valuable. For example, if edges were to be given weights corresponding to the probabilities of junctions’ being spanned by reads, then the total weight of a path could carry information about the *a priori* possibilities that different isoforms are realized. Maximizing (or minimizing) weights of paths through DAGs is known, and a shortest-path algorithm could be used to find those shortest paths quickly.

FIG. 3 provides a diagram of methods of the invention. Once a DAG is built with exons connected in canonical order, as discussed above, paired-end reads are obtained. The paired-end reads are then aligned to the DAG and an alignment is found (having an alignment score that meets a predetermined criteria) between the reads and the DAG comprising nodes representing RNA sequences and edges connecting pairs of nodes. The predetermined criteria may be, for

example, the highest-scoring alignment. The alignment may proceed by the modified Smith-Waterman algorithm described above. The alignment may include, for a pair of paired-end reads, identifying transcripts in the transcriptome based on paths through the directed acyclic data structure that include the alignments. The alignment may be performed efficiently by using insert sizes as a constraint on alignment. The paired-end reads may be used to build the DAG (e.g., add edges).

The method is suited for analysis of reads obtained by RNA-Seq. In certain embodiments, the directed acyclic data structure represents substantially all known features for at least one chromosome. The directed acyclic data structure may be created and stored prior to the finding the alignments. In some embodiments, the alignments are found by comparing each sequence read to at least a majority of the possible paths through the directed acyclic data structure. The method may include assembling the plurality of sequence reads into a contig based on the found alignments.

3. Using insert sizes as a constraint on alignment

One way in which insert size information can be combined with a DAG is to use insert sizes as a constraint on alignment. There are possible gains both in speed and in accuracy from using this method. The method may be more accurate than aligning the paired-end mates separately because it keeps one from discarding important information relevant to the alignment. It is noted that aligning two paired-end mates independently amounts, approximately, to choosing two locations among the L possible locations in the genome (if the genome has length L). Thus the aligning involves choosing among roughly L^2 possibilities. However, the first alignment is used to strongly constrain the other alignment to some region of constant size k around the first alignment, then there are only roughly kL possibilities--on the order of L , not L^2 . This will be much faster.

FIGS. 4-6 illustrate use of insert size information to constrain alignment.

FIG. 4 illustrates a hypothetical DAG of arbitrary structure and a hypothetical pair of paired-end reads defining an insert size k . In the DAG as drawn, nodes are squares and edges are lines. In the depicted embodiment, insert size k is drawn from an approximately normal distribution with a mean of k bp. The first member of the pair is aligned to the DAG (e.g., using Smith-Waterman or modified Smith-Waterman described above). In general, paired end

sequencing involves a fragment with adapters on each end, with paired-end reads that extend inward, towards each other, from the fragment-adapter boundaries. There may exist a gap between the paired end reads. Insert length may be taken to mean the length of the fragment—i.e., the distance from one fragment-adapter boundary to the other fragment-adapter boundary. See, e.g., Lindgreen, 2012, AdapterRemoval: easy cleaning of next-generation sequence reads, BMC Res Notes 5:337 (e.g., Lindgreen, Figure 1, panel C), the contents of which are incorporated by reference.

In FIG. 5, the left-most black square represents the exon at which the first member of the pair aligns. A set of distances is determined that accounts for most (e.g., 99.95%) of the expected insert-size values. For example, where the insert-sizes are normally distributed, the set of distances may be those distances within 3 standard deviations of k . The determined set of distances is used to select downstream nodes and thus candidate paths. The selected downstream nodes represent potential alignment sites for the second member of the pair of paired-end reads. In FIG. 5, the set of downstream nodes are also filled in black squares. Paths that begin with the upstream node and end with one of the downstream nodes are candidate paths to which to align the pair of paired-end reads.

FIG. 6 illustrates the candidate paths of the DAG from FIG. 4. It is noted that FIG. 6 also presents a DAG, a DAG that is contained within the DAG of FIG. 4. All of the candidate paths and only the candidate paths are illustrated in FIG. 6. The hypothetical pair of paired-end reads that define insert size k (from FIG. 4) can be aligned to this DAG presented in FIG. 6 for computational savings. Without recognizing this constraint on alignment, each read of the pair may potentially align to any location along a genome of length L and in a complex DAG (e.g., the DAG in FIG. 4 is approximately 4-plex across its length), there may be many more locations (i.e., if a DAG is, on average, n -plex across any loci, aligning a pair of paired-ends reads amounts to choosing from $(nL)^2$ options). Performing the alignment with a DAG according to the illustrated methods can make a very large project computationally much less costly, yielding a much greater throughput with high volumes of paired-end read data from RNA-Seq projects.

4. *Sensitivity to varieties of alternative splicing.*

Current techniques make it difficult or impossible to distinguish between several kinds of situations when a read spans the junction of exons. One such situation is where mutations have

occurred. Another such situation arises when, even in the absence of mutations, a read may appear not to match on either side of a junction because a form of alternative splicing has occurred in which part of a codon has been included or skipped. Also, sequencing errors present another such situation.

Consider, as an example, the following sequence:

AUGCAUUUCCGUGGAUAGCGAUGCAUACUCACGAUG AUGAAA
AUGCAUCAGAAAUAG (SEQ ID NO: 3)

where plain text regions represent exons, the first underlined portion represents an intron, and the second underlined portion is a region that, due to alternative splicing, might or not be included in the second exon. Thus the mature RNA might take either of two forms:

AUGCAUUUCCGUGGAUAGAUGAAAAUGCAUCAGAAAUAG (SEQ ID NO: 4)
AUGCAUUUCCGUGGAUAGAUGCAUCAGAAAUAG (SEQ ID NO: 5)

Now suppose that we obtain the following read: TGGATAGATGAA (SEQ ID NO: 6). Such a read could have several causal histories that differ in important ways.

FIG. 7 presents three alternative histories involving these sequences. In the first situation depicted across the top third of FIG. 7, the read is taken from SEQ ID NO: 4, and no mutation or sequencing error has occurred. In the second situation, the read is taken from SEQ ID NO: 5, and a sequencing error has occurred. In the third situation, the read is taken from transcript SEQ ID NO: 5, and the read accurately reflects a mutation.

If a linear reference sequence consists only of the plain text region (dashed underlines in FIG. 7), it is very likely that the alignment score will be so high—boosted by many matches and having suffered only one mismatch penalty—that the algorithm in use will not flag the alignment as possibly erroneous. Rather than examine some other file or database for splicing information, it will simply posit a mutation or single sequencing error. Aligning against a DAG reference does better.

FIG. 8 depicts an alignment of SEQ ID NO: 3 against a DAG reference (appearing now in FIG. 8 as SEQ ID NO: 4 since the intron is gone). Here the accurate alignment is immediately recovered, as it will have the highest score of any possible alignment to this DAG. Due to the nature of a DAG, FIG. 8 also depicts:

AUGCAUUUCCGUGGAUAGAUGCAUCAGAAAUAG (SEQ ID NO: 5)

This is a very general phenomenon and will occur often, given how often small parts of exons are subject to inclusion or exclusion by alternative splicing. A system involving linear references will usually fail to detect mischaracterizations when these situations arise, and the only way to prevent these mistakes is to align against so many linear references as to make the process prohibitively expensive (in both time and money).

5. *Using paired-end data properly: modifying the DAG.*

The invention provides additional exploits of the relationship between the insert size and the correct alignments. One can view the current paradigm as treating the insert size as a constraint on a given alignment or pair of alignments. This is an incomplete description of the situation. It may be more accurate to view insert sizes as a constraint on the combination of the alignments and the transcript.

Methods of the invention involve using insert sizes as a quality control not only on the alignments but also on the transcript. If each of two paired-end mates has a high-scoring alignment in the transcript DAG, and if this pair of alignments corresponds to a grossly implausible insert size, it will often be more reasonable to adjust beliefs about the transcript (e.g., by adding edges to the graph) than to adjust our beliefs about the correctness of each of the alignments.

This can be implemented by—whenever two paired-end mates align to a DAG—calculate the insert size or sizes implied by that alignment. (More than one size might be involved because the read might be consistent with several features.) If the quotient of the smallest such insert size and the expected value of the insert size is greater than some constant (perhaps 2), examine the DAG to find a pair of unconnected nodes which, if connected, would entail the existence of an isoform for which the quotient of the implied insert size for that isoform and the expected value of the insert size is between 0.5 and 1.5. If such a pair of nodes exists, connect them with an edge. If several such pairs exist, choose the one that yields the ratio closest to 1 and connect it.

In certain aspects, the invention provides a method of analyzing a transcriptome that involves obtaining at least one pair of paired-end reads from a transcriptome, finding an

alignment with an optimal score between a first read of the pair and a node in a directed acyclic data structure (the data structure has nodes representing RNA sequences such as exons or transcripts and edges connecting each pair of nodes), identifying candidate paths that include the node connected to a downstream node by a path having a length substantially similar to an insert length of the pair of paired-end reads, and aligning the paired-end reads to the candidate paths to determine an optimal-scoring alignment and modifying the DAG to include a new edge if necessary or preferred to create a candidate path indicated by a pair of paired end reads. One of skill in the art will recognize that a path having a length substantially similar to an insert length of the pair of paired-end reads can be defined by any suitable criteria. For example, substantially similar may mean the same number of nucleotides, or may preferably mean ± 10 or ± 100 nucleotides. In certain embodiments, substantially similar can be defined by a user to mean within 10% of, or preferably within 25% of, or within 50% of, or within 100% of each other in length. In a preferred embodiment, substantially similar in length means that one is not more than twice the length of the other. The method is suited for a number of paired-end reads such as those from cDNA fragments (e.g., cDNA fragments representing a single transcriptome). The node and the downstream node may represent a pair of exons from which the pair of paired-end reads were obtained. The method may include identifying an isoform based on the determined optimal-scoring alignment. In some embodiments, the identified isoform is a novel isoform. The directed acyclic data structure may be updated to represent the novel isoform (e.g., by adding a new edge edge). The method may include excluding any paths that are not candidate paths from any alignment calculations involving the pair of paired-end reads. Other aspects of the invention provide for expression analysis through, e.g., systems and methods useful for inferring isoform frequency.

6. *Using paired-end data: inferring isoform frequency.*

A powerful technique arises not from using paired-end structure to align properly, but from using alignment to infer isoform frequency with the intermediate step of inferring insert sizes from alignment. If the frequency with which pairs of exons appear in a sample is known (see, e.g., Harrow, 2012, Genome Research 22:1760), those frequencies can be used to construct the frequencies with which isoforms appear in a sample. Moreover, different pairs of exons will

usually be separated by different distances. Thus, the distances observed between paired-end mates can be used (slightly indirectly) to keep track of exon frequencies.

An illustrative example is as follows. For each pair of paired end mates, the pair may first optionally be aligned to the DAG reference. The distances between a reference point (e.g., the left-most position) on each pair member is noted and tracked. This can be obtained from the alignment to the DAG, although it is noted that methods of inferring isoform frequency may not require aligning the pairs to the DAG.

One suitable method for noting and tracking the distances between the reference point for each aligned pair is through the use of a hash table. In the distance hash, keys may be all distances between 1 and 1000 and the values of each key are initialized to zero. Incrementing can then be done by looking up the value of the key corresponding to the distance and incrementing that value. Note that this sort of manipulation of a hash table can be done extremely fast--indeed, in constant time. Additionally it is noted that hash tables are well-supported by existing bioinformatics packages such as BioPerl. Each pair of exons is then associated with the sum of the values corresponding to the distances nearest the true distance between the exons. (That is: add up all the “hits” observed for distances close to the distance between those exons. The quotient of that number and the total number of pairs will be very close to the frequency with which that pair of exons occurs, relative to all pairs of exons, in the sample.)

From this data, isoform frequency can be determined algebraically. The preferred form of equation may depend on the structure of relevant isoforms, but a general technique obtains, as follows:

Suppose that there are three isoforms, one of which contains exons A, B, and C, another of which contains exons A, B, and D, and the third of which contains exons B, C, and D. The frequencies of the exon pairs—where $[AB]$ is the frequency of the exon pair AB among all of the aligned reads—is obtained from the aligned reads (i.e., each time exon i is connected to exon j in an alignment, the corresponding frequency is incremented by one).

Then, if the frequencies of the three isoforms are (in order) $[A]=f$, $[B]=g$, and $[C]=h$, the frequencies of the exon-pairs should be as follows:

$$[AB] = f + g$$

$$[BC] = f + h$$

$$[BD] = g + h$$

As the left-hand values are known from the alignment, the isoform frequencies can be solved. In general, methods of the invention include writing equations relating exon-pair-frequencies to isoform-frequencies and use the former to obtain the latter. One of skill in the art will recognize algebraic methods for solving the depicted equations. To give an illustrative example, it is noted that algebraic manipulation of the above yields:

$$f = [AB] - g,$$

and:

$$h = [BD] - g$$

Thus:

$$[BC] = f + [BD] - g$$

and:

$$[BC] = [AB] - g + [BD] - g$$

Rearrangement provides:

$$[BC]/2([AB]+[BD]) = g$$

Since $[BC]$, $[AB]$, and $[BD]$ are all known from the alignment, g is thus calculated.

Solving for f and h is then trivial.

In some circumstances, there may appear to be a “degree of freedom” problem in that the number of variables to be solved is not less than the number of input variables. Any suitable technique for addressing this apparent problem could be employed and may depend on the particular circumstance. Suitable techniques may be found to include normalizing the variables; solving for more than one set of isoform frequencies and choosing one after-the-fact; obtaining at least one input value extrinsically (e.g., from the biological sample), or providing input information (e.g., from the literature or an assumption).

It is important to note a graphical interpretation of the key step: suppose a graph is produced having distances on the x-axis and frequencies on the y-axis. Exon-pairs will then appear as peaks on that graph, and frequently some of those peaks will be readily interpretable as isoform-frequencies also.

Thus, the invention provides a method of inferring isoform frequency in a transcriptome by obtaining a plurality of pairs of paired-end reads from a transcriptome and optionally aligning those reads to a reference data structure. The data structure includes a plurality of exons from a

genome each represented as a node. Pairs of nodes, or exons, are connected at their proximal ends by an edge. Distances between the aligned pairs of the paired-end reads and the frequencies of the distances between the aligned pairs are determined. A set of isoform paths and isoform frequencies are determined such that representing the isoform paths through the structure at the isoform frequencies results in pairs of exons being included in the isoform paths at the frequencies of the distances the aligned pairs. Each path of the determined set of isoform paths may represent a transcript isoform present in the organism. The distances between the aligned pairs may be path lengths from an upstream end of a first aligned member of each pair to a downstream end of a second aligned member of that pair. In certain embodiments, determining the set of isoform path includes discovering a novel isoform path. The reference data structure may be updated to include the novel isoform path.

Determining the set of isoform paths can be performed algebraically, e.g., by solving a set of linear equations. In some embodiments, the one of the isoform frequencies is obtained from biological data prior to solving the set of linear equations (e.g., so remaining variables in the linear equations can be determined). Additionally or alternatively, the frequencies of the distances between the aligned pairs may be normalized prior to determining a set of isoform paths and isoform frequencies.

7. *Incorporating a priori exon- and isoform-count distributions.*

In some embodiments, the invention exploits situations in which there is *a priori* data for the relative frequency of isoform counts. Isoform distributions for transcripts may be inferred to correlate to distributions of protein isoforms. Questions about the distribution of exon- and isoform-counts for a given gene, and the way in which these distributions are sensitive to other facts about a sequence (e.g., whether it is a coding or non-coding sequence), have been studied thoroughly, if incompletely.

Isoform frequencies may be obtained from the literature or by performing an analysis. For example, isoform frequencies may be inferred from paired-end read RNA-Seq reads, as described in Section 6, ‘Using paired-end data: inferring isoform frequency’, above. In some embodiments, insert sizes from paired-end read data used to infer isoform frequencies, as described in Section 8, ‘Using insert sizes more effectively to infer isoform probabilities’, below.

In certain embodiments, isoform frequencies are obtained from a source such as a database or the literature.

FIG. 9 represents protein isoform distributions obtained by a study of SwissProt amino acid sequence, and is a reproduction of the third figure from Nakao, et al., 2005, Large-scale analysis of human alternative protein isoforms: pattern classification and correlation with subcellular localization signals, Nucl Ac Res 33(8):2355-2363. Additionally, Chang, et al., have used programming methods to deduce quantitative distributions of various alternative splicing forms. Chang, et al., 2005, The application of alternative splicing graphs in quantitative analysis of alternative splicing form from EST database, Int J. Comp. Appl. Tech 22(1):14.

The information presented by FIG. 9 represents an isoform count distribution that can be used to aid in mapping reads to a DAG. Similarly, exon count distribution information may be obtained.

For example, there is a demonstrated distribution of number of exons per transcript. See, e.g., Harrow, et al., (2012, GENCODE: The reference human genome annotation for The ENCODE Project, Genome Res 22:1760-1774. Transcripts at protein-coding loci demonstrate a peak at four exons per transcript, while lncRNAs show a distinct peak at two exons (see Figure 2 of Harrow, et al., 2012). The exon count distribution from Harrow, et al., or the isoform distribution from Nakao, et al., could be interpreted probabilistically, with the y-axis giving relative instead of absolute frequency. Such information may be used to inform judgments about the likelihood that candidate alignments correspond to actual nucleotide sequences. Prior art techniques rely primarily on the scores of the alignments. The invention provides ways of including *a priori* knowledge of isoform number.

For example, for a large set of reads, 75% may unambiguously map to regions in a set of eight isoforms {I1, I2, ..., I8}, while the other 25% of which could be interpreted as mapping either to I5 or to some other combination of exons corresponding to a further isoform I9. In this situation, various kinds of facts are relevant to the decision how to align that last 25% of reads. Such relevant facts include (i) the scores of the alignments to I5 and to I9; (ii) any further knowledge we might have about the regulatory mechanisms; and (iii) whether it is more likely *a priori* for there to be eight or nine isoforms.

The invention incorporates isoform distribution information when mapping reads to a DAG reference. Isoform frequency distribution information may be used in several different

ways. A first exemplary method of including *a priori* isoform number in mapping reads includes: (1) constructing a DAG for the relevant region of the genome (the edges of this DAG can be weighted or unweighted) and (2) associating—with each edge of the DAG—a value that indicates the edge’s usage. Set the usage “off” if the number of reads aligned in a way that spans that edge is smaller than some variable U . This variable intuitively reflects whether the edge has been “used.” One way to estimate it might be by setting it equal to $(\text{the number of reads in the sample}) * (\text{the quotient of a reasonable estimate of the length of an exon and the average read length}) * (.005)$.

Further, the method includes (3) associating with the relevant region of the genome a variable corresponding to a usage threshold T , perhaps setting T equal to the expected value of the number of edges used in the graph plus the standard deviation of the number of edges used in the graph. These numbers could be calculated by estimating the number of edges in a DAG that has the numbers of nodes and maximal paths corresponding to empirical exons frequencies, isoform frequencies, or both that would be expected for the gene in question. The empirical frequencies could be obtained, for example, from Harrow, et al., 2012 for exon count distribution or from Nakao, et al., 2005, for isoform distribution.

The method includes (4) modifying the scoring scheme used in Smith-Waterman or other algorithms, punishing the traversal of an unused edge by some penalty if the total number of “used” edges (i.e., edges for which the Boolean usage value is “on” or “true”) N exceeds T . This penalty may be made to increase as N increases. A reasonable formula for the penalty might be $\max(0, (N - T) * (\text{the insertion penalty}) * 2)$. The zero term appears to avoid negative penalties--i.e., to avoid bonuses for using new edges--but such bonuses might be appropriate in some contexts, in which case there is no need to take the maximum of zero and the latter expression’s value.

To avoid unwelcome path dependencies in the form of the alignments’ being sensitive to the order in which the reads are processed, one might repeat the alignments several times and compare the results. This does not affect the asymptotic running time of the algorithm.

An alternative exemplary method of including *a priori* isoform number in mapping reads includes:

- (i) align each read against the DAG;

(ii) compare an observed number of isoforms to an expected number of isoforms (e.g., obtained from an *a priori* probability distribution of the number of isoforms given background information about the gene); and

(iii) if the difference between the observed number and the expected number exceeds some threshold, re-align the reads that were assigned to the rarest candidate isoforms. In the re-alignment, any alignment to a rare candidate isoform is penalized with a weight that depends on the magnitude of the difference. The penalty may be $\max(0, (N - T) * (\text{the insertion penalty}) * 2)$, where N is a total number of “used” edges (i.e., edges for which the Boolean usage value is “on” or “true”) and T is a usage threshold. In some embodiments, T is an expected value of the number of edges used in the graph plus the standard deviation of the number of edges used in the graph.

By application of methods of using *a priori* probability distribution of the number of isoforms, the reads being mapped to the DAG will be pushed to result in a distribution that approaches the *a priori* distribution.

8. *Using insert sizes more effectively to infer isoform probabilities.*

Suppose that two exons E1 and E2 occur in different isoforms and that they are separated by distances L1 and L2 in those two isoforms. It will occasionally happen that there is obtained a pair of paired-end reads for which one mate maps to E1 and the other to E2. The choice of which isoform to attribute this pair to can then be reduced to the question whether to infer that L1 or L2--or one of those distances shifted by an appropriate amount given the mates' offsets from the edge of the exon--is the true distance between the reads.

Translating the isoform-assignment question to a question about insert sizes is valuable because we usually have good background information about the distribution of insert sizes in a paired-end sequencing run. Given that there will usually be no other reason to believe that the read would more likely to have been taken from one isoform or the other, we can perform a simple Bayesian calculation to assign the read probabilistically to the two isoforms according to the ratio of the probabilities of insert sizes L1 and L2 in the distribution.

This is a marked improvement on current techniques, which generally view reads as carrying equal evidential weight for the two isoforms in this case. Incorporating the insert size in the way described is therefore a way of not ignoring relevant evidence.

Thus, we could count isoforms in these ambiguous cases as follows:

(1) Calculate a probability distribution of insert sizes. (Denote by $P(s)$ the probability of a read's having insert size s .) Note that this needs to be calculated only once and that this could be done in a preprocessing step.

(2) If each of two paired-end mates maps inside an exon:

(2a) Make a list of all the isoforms that include those exons.

(2b) Associate with each isoform the insert size that would be implied by the read's having derived from that isoform.

(2c) Calculate $P(s)$ for each of the insert sizes calculated in (2b).

(2d) Associate with each isoform the relevant value of $P(s)$ divided by the sum of all the $P(s)$ values calculated in (2c).

It is noted that techniques described herein will improve read counts of isoforms in a way that does not affect one's ability to use "downstream" tools--that is, tools that take read counts as inputs and give other useful outputs.

9. *Construction of a "maximal DAG."*

In some embodiments, it may be beneficial to construct a DAG in which every pair of nodes in such a DAG is connected, with the edge going from the node representing the "earlier" or "left-side" exon to the "later" or "right-side" exon.

Using Gencode or a similar annotated transcriptome, nodes are created to represent exons and other features. Gencode provides enough information to associate with these nodes such information as a sequence for the node, the positions it occupies in a reference sequence (and therefore its length), and so on. Thus, using a computer system comprising a processor coupled to a non-transitory memory, each of a plurality of features from a genome is represented as a node. For each pair of the plurality of features, an edge is created connecting the two members of the pair by their proximal ends. The biological interpretation of such a move is to conjecture the existence of, or at least to facilitate the discovery of, isoforms representing every combination of exons that respects the order in which those exons appear in the genome. There are both inductive reasons and biochemical reasons that it is sensible to respect the order in which exons appear in the genome.

For example, reports suggest that exons in multi-exon pre-mRNAs are always maintained in order. See, e.g., Black, 2005, A simple answer for a splicing conundrum, PNAS 102:4927–8. Without being bound by any particular mechanism, it may be theorized that splicing machinery recognizes the start and end of introns, as well as branch-point adenines or other conserved sequences proximal to splice junctions. The splicing machinery, or spliceosome, moves along the pre-mRNA from 5' to 3' removing the introns and splicing the exons. Thus it appears that there is phenomenological as well as literature support for connecting all exons in the 5' to 3' order in which they appear along the genomic DNA strand. See also Shao et al., 2006, Bioinformatic analysis of exon repetition, exon scrambling and trans-splicing in humans, Bioinformatics 22: 692-698. For certain assumptions, methods of the invention follow the assumption that human pre-mRNAs are spliced into linear molecules that retain the exon order defined by the genomic sequence.

Additionally, where it is desirable to include instances of exon scrambling (see, e.g., Carrington et al., 1985, Polypeptide ligation occurs during post-translational modification of concanavalin A, Nature 313:64-67) in a directed acyclic data structure, the assumption that exons are spliced in their original order can be maintained substantially by including duplicate exons as a “fiction” (or, in the alternative, the assumption can be abandoned).

Accordingly where, for example, a genome includes the following exons in the following order: ABCDEFGHI

using dashes as edges and letters as nodes, the complete set of edges can be represented by:

A-C; A-D; A-E; A-F; A-G; A-H; A-I; B-C; B-D; B-E; B-F; B-G; B-H; B-I; C-D; C-E; C-F; C-G; C-H; C-I; D-E; D-F; D-G; D-H; D-I; E-F; E-G; E-H; E-I; F-G; F-H; F-I; G-H; G-I.

If this proves adequate to discover all isoforms but one, and that one isoform is—to give an example: ADEHG

Then, the acyclic data structure can be “spiked” with an artificial, or fictitious, G—call it G'. Then that isoform can be represented by A-D-E-H-G' and the structure need not include any edge from H to G. However, it may be found that exon scrambling can be ignored and genomic exon order maintained by all included edges.

A DAG of the invention may first be prepared by using input data from an annotated transcriptome. As such, *ab initio*, the DAG will include all known isoforms. If a maximal DAG

is made, the DAG will include all possible edges that maintain genomic exon order and will thus then include nodes and edges that represent both known and novel isoforms. The nodes and edges are stored as a directed acyclic data structure in the memory and the method includes aligning a paired-end sequence reads to a path comprising a subset of the features connected by a subset of the edges, thereby identifying an isoform corresponding to the path.

When new data is aligned to the DAG (e.g., paired-end RNA-Seq reads), those data may include exons that had not been identified in the annotated reference. As the exon is novel in this example, so is the isoform. Use of a DAG as the reference is ideally suited to this circumstance as aligning a novel isoform to an existing DAG can also guide the creation of a new node in the DAG to represent the novel exon. Thus, where an identified isoform is novel, aligning the transcript sequence can involve creating a new node in the directed acyclic data structure. A new node can also be created for splice variants, e.g., where one exon differs from another by a single codon.

Preferably, the maximal DAG is initially built—all or partly—prior to and independent of analyzing paired-end reads. Thus, the method may include creating the edge for each pair of the plurality of features prior to aligning the transcript. For reasons related to the nature of directed acyclic data structures, the reference represented in the data structure can be arbitrarily large—e.g., hundreds or thousands of exons, introns, or both.

10. Systems of the invention

FIG. 10 illustrates a computer system 401 useful for implementing methodologies described herein. A system of the invention may include any one or any number of the components shown in FIG. 10. Generally, a system 401 may include a computer 433 and a server computer 409 capable of communication with one another over network 415. Additionally, data may optionally be obtained from a database 405 (e.g., local or remote). In some embodiments, systems include an instrument 455 for obtaining sequencing data, which may be coupled to a sequencer computer 451 for initial processing of sequence reads.

In some embodiments, methods are performed by parallel processing and server 409 includes a plurality of processors with a parallel architecture, i.e., a distributed network of processors and storage capable of comparing a plurality of sequences (e.g., RNA-Seq reads) to a reference sequence construct (e.g., a DAG). The system comprises a plurality of processors that

simultaneously execute a plurality of comparisons between a plurality of reads and the reference sequence construct. While other hybrid configurations are possible, the main memory in a parallel computer is typically either shared between all processing elements in a single address space, or distributed, i.e., each processing element has its own local address space. (Distributed memory refers to the fact that the memory is logically distributed, but often implies that it is physically distributed as well.) Distributed shared memory and memory virtualization combine the two approaches, where the processing element has its own local memory and access to the memory on non-local processors. Accesses to local memory are typically faster than accesses to non-local memory.

Computer architectures in which each element of main memory can be accessed with equal latency and bandwidth are known as Uniform Memory Access (UMA) systems. Typically, that can be achieved only by a shared memory system, in which the memory is not physically distributed. A system that does not have this property is known as a Non-Uniform Memory Access (NUMA) architecture. Distributed memory systems have non-uniform memory access.

Processor–processor and processor–memory communication can be implemented in hardware in several ways, including via shared (either multiported or multiplexed) memory, a crossbar switch, a shared bus or an interconnect network of a myriad of topologies including star, ring, tree, hypercube, fat hypercube (a hypercube with more than one processor at a node), or n-dimensional mesh.

Parallel computers based on interconnected networks must incorporate routing to enable the passing of messages between nodes that are not directly connected. The medium used for communication between the processors is likely to be hierarchical in large multiprocessor machines. Such resources are commercially available for purchase for dedicated use, or these resources can be accessed via “the cloud,” e.g., Amazon Cloud Computing.

A computer generally includes a processor coupled to a memory and an input-output (I/O) mechanism via a bus. Memory can include RAM or ROM and preferably includes at least one tangible, non-transitory medium storing instructions executable to cause the system to perform functions described herein. As one skilled in the art would recognize as necessary or best-suited for performance of the methods of the invention, systems of the invention include one or more processors (e.g., a central processing unit (CPU), a graphics processing unit (GPU),

etc.), computer-readable storage devices (e.g., main memory, static memory, etc.), or combinations thereof which communicate with each other via a bus.

A processor may be any suitable processor known in the art, such as the processor sold under the trademark XEON E7 by Intel (Santa Clara, CA) or the processor sold under the trademark OPTERON 6200 by AMD (Sunnyvale, CA).

Input/output devices according to the invention may include a video display unit (e.g., a liquid crystal display (LCD) or a cathode ray tube (CRT) monitor), an alphanumeric input device (e.g., a keyboard), a cursor control device (e.g., a mouse or trackpad), a disk drive unit, a signal generation device (e.g., a speaker), a touchscreen, an accelerometer, a microphone, a cellular radio frequency antenna, and a network interface device, which can be, for example, a network interface card (NIC), Wi-Fi card, or cellular modem.

Incorporation by Reference

References and citations to other documents, such as patents, patent applications, patent publications, journals, books, papers, web contents, have been made throughout this disclosure. All such documents are hereby incorporated herein by reference in their entirety for all purposes.

Equivalents

Various modifications of the invention and many further embodiments thereof, in addition to those shown and described herein, will become apparent to those skilled in the art from the full contents of this document, including references to the scientific and patent literature cited herein. The subject matter herein contains important information, exemplification and guidance that can be adapted to the practice of this invention in its various embodiments and equivalents thereof.

Claims

1. A method of analyzing a transcriptome, the method comprising:
 - obtaining a pair of paired-end reads from a transcriptome;
 - finding an alignment with an optimal score between a first read of the pair and a node representing an exon in a directed acyclic graph (DAG), the DAG comprising nodes representing a plurality of exons and introns from features in an annotated transcriptome database, and edges connecting pairs of the nodes in their canonical genomic order;
 - identifying candidate paths within the DAG that include the node connected to a downstream node by a path having a length substantially similar to an insert length of the pair of paired-end reads;
 - aligning the paired-end reads to the candidate paths to determine an optimal-scoring alignment by excluding any paths that are not candidate paths from any alignment calculations involving the pair of paired-end reads; and
 - identifying an isoform from the transcriptome using the optimal-scoring alignment of the paired-end reads.
2. The method of claim 1, further comprising obtaining a plurality of paired-end reads from the transcriptome.
3. The method of claim 2, further comprising determining optimal-scoring alignments for each of the plurality of pairs of paired-end reads.
4. The method of claim 3, wherein the plurality of paired-end reads are obtained from cDNA fragments from the transcriptome.
5. The method of claim 1, wherein the identified isoform is a novel isoform.
6. The method of claim 5, further comprising updating the DAG to represent the novel isoform.

7. The method of claim 6, wherein updating the DAG to represent the novel isoform comprises adding at least one new node.
8. The method of claim 7, wherein the node and the downstream node represent a pair of exons from which the pair of paired-end reads were obtained.
9. The method of claim 1, further comprising:
 - aligning a plurality of pairs of paired-end reads to the DAG;
 - determining distances between the aligned pairs of the plurality of pairs of paired-end reads and the frequencies of the distances between the aligned pairs; and
 - determining a set of isoform paths and isoform frequencies such that representing the isoform paths through the DAG at the isoform frequencies results in pairs of features being included in the isoform paths at the frequencies of the distances between the aligned pairs.
10. The method of claim 9, wherein the features include exons.
11. The method of claim 10, wherein each path of the determined set of isoform paths represents a transcript isoform present in the organism.
12. The method of claim 9, wherein the distances between the aligned pairs comprise path lengths from an upstream end of a first aligned member of each pair to a downstream end of a second aligned member of that pair.
13. A system for transcriptomics, the system comprising:
 - a computer system comprising a processor coupled to a non-transitory memory having stored therein a directed acyclic graph (DAG), the DAG comprising nodes representing a plurality of exons and introns from features in an annotated transcriptome database, and edges connecting pairs of the nodes in their canonical genomic order, wherein the system is operable to:
 - obtain a pair of paired-end reads from a cDNA fragment;

find an alignment with an optimal score between a first read of the pair and a node in the DAG;

identify candidate paths within the DAG that include the node connected to a downstream node by a path having a length substantially similar to an insert length of the paired end reads;

align the paired-end reads to the candidate paths to determine an optimal-scoring alignment by excluding any paths that are not candidate paths from any alignment calculations involving the pair of paired-end nodes; and

identify an isoform from the transcriptome using the optimal scoring-alignment of the paired-end reads.

14. The system of claim 13, further operable to obtain a plurality of paired-end reads from cDNA fragments.

15. The system of claim 13, further operable to update the DAG to represent the novel isoform.

16. The system of claim 15, wherein updating the DAG to represent the novel isoform comprises adding at least one new node.

17. A method of inferring isoform frequency in a transcriptome, the method comprising:
aligning a plurality of pairs of paired-end reads from cDNA to a reference directed acyclic graph (DAG) comprising a plurality of nodes representing exons and introns from annotated features in a genetic feature database, wherein pairs of the plurality of nodes are connected in their canonical genomic order at their proximal ends by an edge, wherein the aligning for a pair of paired-end reads comprises identifying an exon at which a first member of a pair aligns, determining a set of distances that accounts for an expected insert-size value of the pair, using the determined set of distances to select downstream nodes as potential alignment sites for a second member of the pair, identifying paths that begin with the upstream node and end with one of the downstream nodes as candidate paths, and excluding any paths that are not candidate paths from any alignment calculations involving the pair of paired-end reads;

determining distances between the aligned pairs of the paired-end reads and the frequencies of the distances between the aligned pairs; and

determining a set of isoform paths and isoform frequencies such that representing the isoform paths through the DAG at the isoform frequencies results in pairs of features being included in the isoform paths at the frequencies of the distances between the aligned pairs.

18. The method of claim 17, wherein the cDNA represents substantially an entire transcriptome of an organism.

19. The method of claim 18, wherein each path of the determined set of isoform paths represents a transcript isoform present in the organism.

20. The method of claim 19, wherein determining the set of isoform paths includes discovering a novel isoform path.

21. The method of claim 20, further comprising updating the reference DAG to include the novel isoform path.

22. The method of claim 17, wherein the distances between the aligned pairs comprise path lengths from an upstream end of a first aligned member of each pair to a downstream end of a second aligned member of that pair.

23. The method of claim 17, further comprising normalizing the frequencies of the distances between the aligned pairs prior to determining a set of isoform paths and isoform frequencies.

24. A system of inferring isoform frequency in a transcriptome, the system comprising:
a computer system comprising a processor coupled to a non-transitory memory having stored therein a reference DAG comprising a plurality of nodes representing exons and introns from features in a genetic feature database, wherein pairs of the plurality of nodes are connected in their canonical genomic order at their proximal ends by an edge, wherein the system is operable to:

align a plurality of pairs of paired-end reads to the reference DAG, wherein the aligning for a pair of paired-end reads comprises identifying an exon at which a first member of a pair aligns, determining a set of distances that accounts for an expected insert-size value of the pair, using the determined set of distances to select downstream nodes as potential alignment sites for a second member of the pair, identifying paths that begin with the upstream node and end with one of the downstream nodes as candidate paths, and excluding any paths that are not candidate paths from any alignment calculations involving the pair of paired-end reads;

determine distances between the plurality of pairs of paired-end reads from cDNA and the frequencies of the distances between the pairs; and

determine a set of isoform paths and isoform frequencies such that representing the isoform paths through the DAG at the isoform frequencies results in pairs of features being included in the isoform paths at the frequencies of the distances between the aligned pairs.

25. The system of claim 24, wherein the cDNA represents substantially an entire transcriptome of an organism.
26. The system of claim 25, wherein each path of the determined set of isoform paths represents a transcript isoform present in the organism.
27. The system of claim 26, wherein determining the set of isoform paths includes discovering a novel isoform path.
28. The system of claim 27, further operable to update the reference DAG to include the novel isoform path.
29. The system of claim 24, wherein the distances between the aligned pairs comprise path lengths from an upstream end of a first aligned member of each pair to a downstream end of a second aligned member of that pair.
30. The system of claim 24, further operable to normalize the frequencies of the distances between the aligned pairs prior to determining the set of isoform paths and isoform frequencies.