

Trans-ABySS v1.2.0: User Manual

07 January 2011

Prepared by:

Readman Chiu, Rong She, Hisanaga Mark Okada, Gordon Robertson, Shaun Jackman, Jenny Qian, Lucas Swanson

On behalf of:

Gordon Robertson, Jacqueline Schein, Readman Chiu, Richard Corbett, Matthew Field, Shaun D Jackman, Karen Mungall, Sam Lee, Hisanaga Mark Okada, Jenny Q Qian, Malachi Griffith, Anthony Raymond, Nina Thiessen, Timothee Cezard, Yaron S Butterfield, Richard Newsome, Simon K Chan, Rong She, Richard Varhol, Baljit Kamoh, Anna-Liisa Prabhu, Angela Tam, YongJun Zhao, Richard A Moore, Martin Hirst, Marco A Marra, Steven J M Jones, Pamela A Hoodless & Inanc Birol

Genome Sciences Centre, BC Cancer Agency
Vancouver, BC, Canada V5Z 4S6

Contact: Readman Chiu (rchiu@bcgsc.ca)

User forum: <http://groups.google.com/group/trans-abyss?hl=en>

Table of contents

ABySS and Trans-ABySS

Licenses

Getting ABySS

Getting Trans-ABySS

 Download

 Unpacking

Installation

 Trans-ABySS Software

 External software

Assembling and analyzing transcriptome data

 Trans-ABySS pipeline overview

 ABySS assemblies and folder structure

 Trans-ABySS folder structure

 Run Trans-ABySS pipeline

 Setup configuration files

 Setup transcript annotations and genome sequence

 Setup input file

Run trans-ABYSS

- Setting up contigs for analysis

 - Process ABYSS contigs for each k-mer assembly

 - Create the merged assembly

 - Using the wrapper

- Contig and read alignments

 - Read alignments to contigs

 - Contig alignments to a reference genome

 - Aligning reads to a reference genome

- Transcriptome assembly analysis

 - Identify candidate novel transcript structures

 - Estimate gene-level expression

 - Identify candidate gene fusion events

 - Identify putative chimeric transcripts

- Additional Trans-ABYSS functions

 - Identify candidate SNVs and INDELS

 - Identify candidate polyadenylation sites

Datasets

- Large

- Small

 - Insr_UTR

 - Polyadenylation site analysis

References

ABYSS and Trans-ABYSS

ABYSS is a de Bruijn graph-based short-read assembler that can process genome or transcriptome sequence data (Simpson et al. 2009, Birol et al. 2009).

Trans-ABYSS is an analysis pipeline for post-processing ABYSS assemblies of transcriptome sequencing data. It addresses varying transcript expression levels by processing multiple assemblies across a range of k values (Robertson et al. 2010).

The current pipeline can map assembled contigs to annotated transcripts (e.g. RefSeq, Ensembl,...), and can identify candidate novel splicing events such as exon-skipping, novel exons, retained introns, novel introns, and alternative splicing. It can also extract candidate SNVs, INDELS, and gene fusion events from contig alignment data. It also finds putative chimeric transcript events and candidate polyA sites.

The Trans-ABYSS pipeline consists of a) Perl wrapper scripts; b) Python, Perl and bash scripts; and c) command line applications. The pipeline can be run on any POSIX-compliant platform. Processing large datasets will require a computer

cluster.

Licenses

ABySS and Trans-ABySS are released under the terms of the BC Cancer Agency software license agreement. http://www.bcgsc.ca/platform/bioinfo/license/bcca_2010

Getting ABySS

The current Trans-ABySS pipeline will process outputs from ABySS v1.1.2+. ABySS v1.1.1 was used for the Nature Methods publication. Source code for v1.1.1 and for the most current release of ABySS is available at:

www.bcgsc.ca/platform/bioinfo/software/abyss

The ABySS-users discussion group is available at:

<http://groups.google.com/group/abyss-users>

ABySS can be compiled to run on any POSIX-compliant system. Use the following commands to read ABySS man pages:

```
man doc/abyss-pe.1
man doc/ABYSS.1
```

Getting Trans-ABySS

1. Download

The pipeline software can be downloaded from:
<http://www.bcgsc.ca/platform/bioinfo/software/trans-abyss>

2. Unpacking

After unpacking, files will be automatically organized into five folders:

analysis	Contains Python modules and Perl scripts that are used for analyzing ABySS-assembled transcriptome assemblies.
----------	--

annotations	Contains transcript and repeat annotation files used in analysis. It is organized by reference genome assembly (e.g. hg18, mm9, etc).
configs	Contains configuration files (.cfg) that are used for running the trans-ABYSS pipeline.
utilities	Contains Python modules (.py) and ABySS-related binaries that support the analysis modules.
wrappers	Contains Perl scripts (.pl) that are wrappers for running the Trans-ABYSS pipeline.
sample_data	Contains a small sample dataset that can be used for testing

Installation

1. Trans-ABYSS Software

Most of the software is written in Python. Because Trans-ABYSS uses Pysam (<http://code.google.com/p/pysam/>) to parse .sam files, Python 2.6 or later is required.

The wrapper scripts for running the pipeline are written in Perl. All Perl5 versions should work. To use the wrappers, you must add to the Perl path a simple custom configuration module for parsing config files. This module is supplied in the “wrappers” folder.

By setting the following environmental variables you should be ready to run the Trans-ABYSS software:

```
export TRANSABYSS_PATH=/home/user/trans-ABYSS
export PYTHONPATH=.:$PYTHONPATH:$TRANSABYSS_PATH
export PERL5LIB=.:$PERL5LIB:$TRANSABYSS_PATH/wrappers
```

For convenience, a “setup” file is included in the trans-ABYSS root folder, which includes the setup of the environmental variables. Change “TRANSABYSS_PATH” to your own trans-ABYSS directory. Then type “source setup” at the command line. Note that these environment variables need to be set up in each shell where you want to run trans-ABYSS codes. For convenience, you can set it up systematically, e.g. in your .bashrc file.

In addition, the reference genomes and their annotations that are used in analysis should be present in the “annotations” folder. For each reference genome, there should be a “genome.fa” file in its corresponding folder. To run chimeric transcript event finder code, a “2bit” format genome file is used. Please

refer to Section 4.2 for details of required annotation files. The current trans-ABYSS package provides annotation files for two reference genomes: “hg18” (human) and “mm9” (mouse), which can be downloaded separately from the software download page. For analysis on other genomes, please set up their annotation folders in the same fashion.

For convenience, the trans-ABYSS software package provides two scripts in the root directory: “setup_hg18” and “setup_mm9”, to allow automatic download and setup of the “hg18” and “mm9” annotation folders. Simply type “source setup_hg18” and/or “source setup_mm9” to download and set up annotations for hg18 and mm9.

The ‘External software’ section (below) lists other required software.

2. External Software

In addition to Python and Perl, Trans-ABYSS requires the following:

1. Blat (<http://users.soe.ucsc.edu/~kent/src/>)

Blat is used for:

1. Merging: pairwise alignment of contigs to remove redundant contigs.
2. Aligning contigs to a reference genome.

2. Pysam (<http://code.google.com/p/pysam/>)

Pysam is used for parsing .bam files for parsing read-to-contig alignments.

3. BioPython (<http://www.biopython.org/wiki/Download>)

Biopython is used in two parts of Trans-ABYSS analysis:

1. Translating DNA sequence into peptide sequence for identifying potential open reading frames.
2. The “NCBIStandalone.py” module is used for parsing Blast-format output from Blat to extract candidate single nucleotide variants (SNVs) and insertion-deletions (INDELs). After downloading the module, edit the following line in so that HSPs of all scores will be parsed:

```
r"Score =\s*([0-9.e+]+) bits \(([0-9]+\))", line,
```

should be changed to:

```
r"Score =\s*([0-9.e+-]+) bits \(([0-9-]+\))", line,
```

4. Samtools (<http://samtools.sourceforge.net/>)

Samtools is used for merging and indexing read alignment files.

5. Bowtie (<http://bowtie-bio.sourceforge.net/index.shtml>)

Bowtie is used in single-end alignment for aligning reads to contigs.

6. The CPAN Perl module Config::General, IO::Compress, Set::IntSpan

<http://search.cpan.org/~tlinden/Config-General-2.49/General.pm>

<http://search.cpan.org/~pmqs/IO-Compress-2.030/lib/IO/Uncompress/Gunzip.pm>

<http://search.cpan.org/dist/Set-IntSpan/>

IO::Compress is only required if the input reads are gzipped or bziped.

The Config::General and IO::Compress modules are used by the polyadenylation site scripts. The Config::General and Set::IntSpan module is used by the chimeric transcript event finder scripts.

7. BWA (<http://bio-bwa.sourceforge.net/bwa.shtml>)

BWA is used to align PAM and EJ reads to known transcript sequences in the polyadenylation site analysis.

Assembling and analyzing transcriptome data

1. Trans-ABYSS Pipeline Overview

Because transcriptome samples typically contain transcripts with a wide range of expression levels, and assemblies generated with different k-mer lengths perform differently in capturing transcripts expressed at different levels, we recommend using a wide range of k-mer values to assemble read data from an RNA-seq library (Robertson *et al.* 2010). Currently, for a read length L , we typically use a range from $L/2$ to $L-1$ for libraries with $L \leq 50$ bp, and a range from $L/2$ to $L-1$, using every other k , for libraries with $L > 50$ bp.

Trans-ABYSS starts with a set of ABySS assemblies for a range of k values. It processes them into a merged assembly, which is then used to generate alignments, identify novel events and perform other analyses. Figure 1 shows an overview of the pipeline.

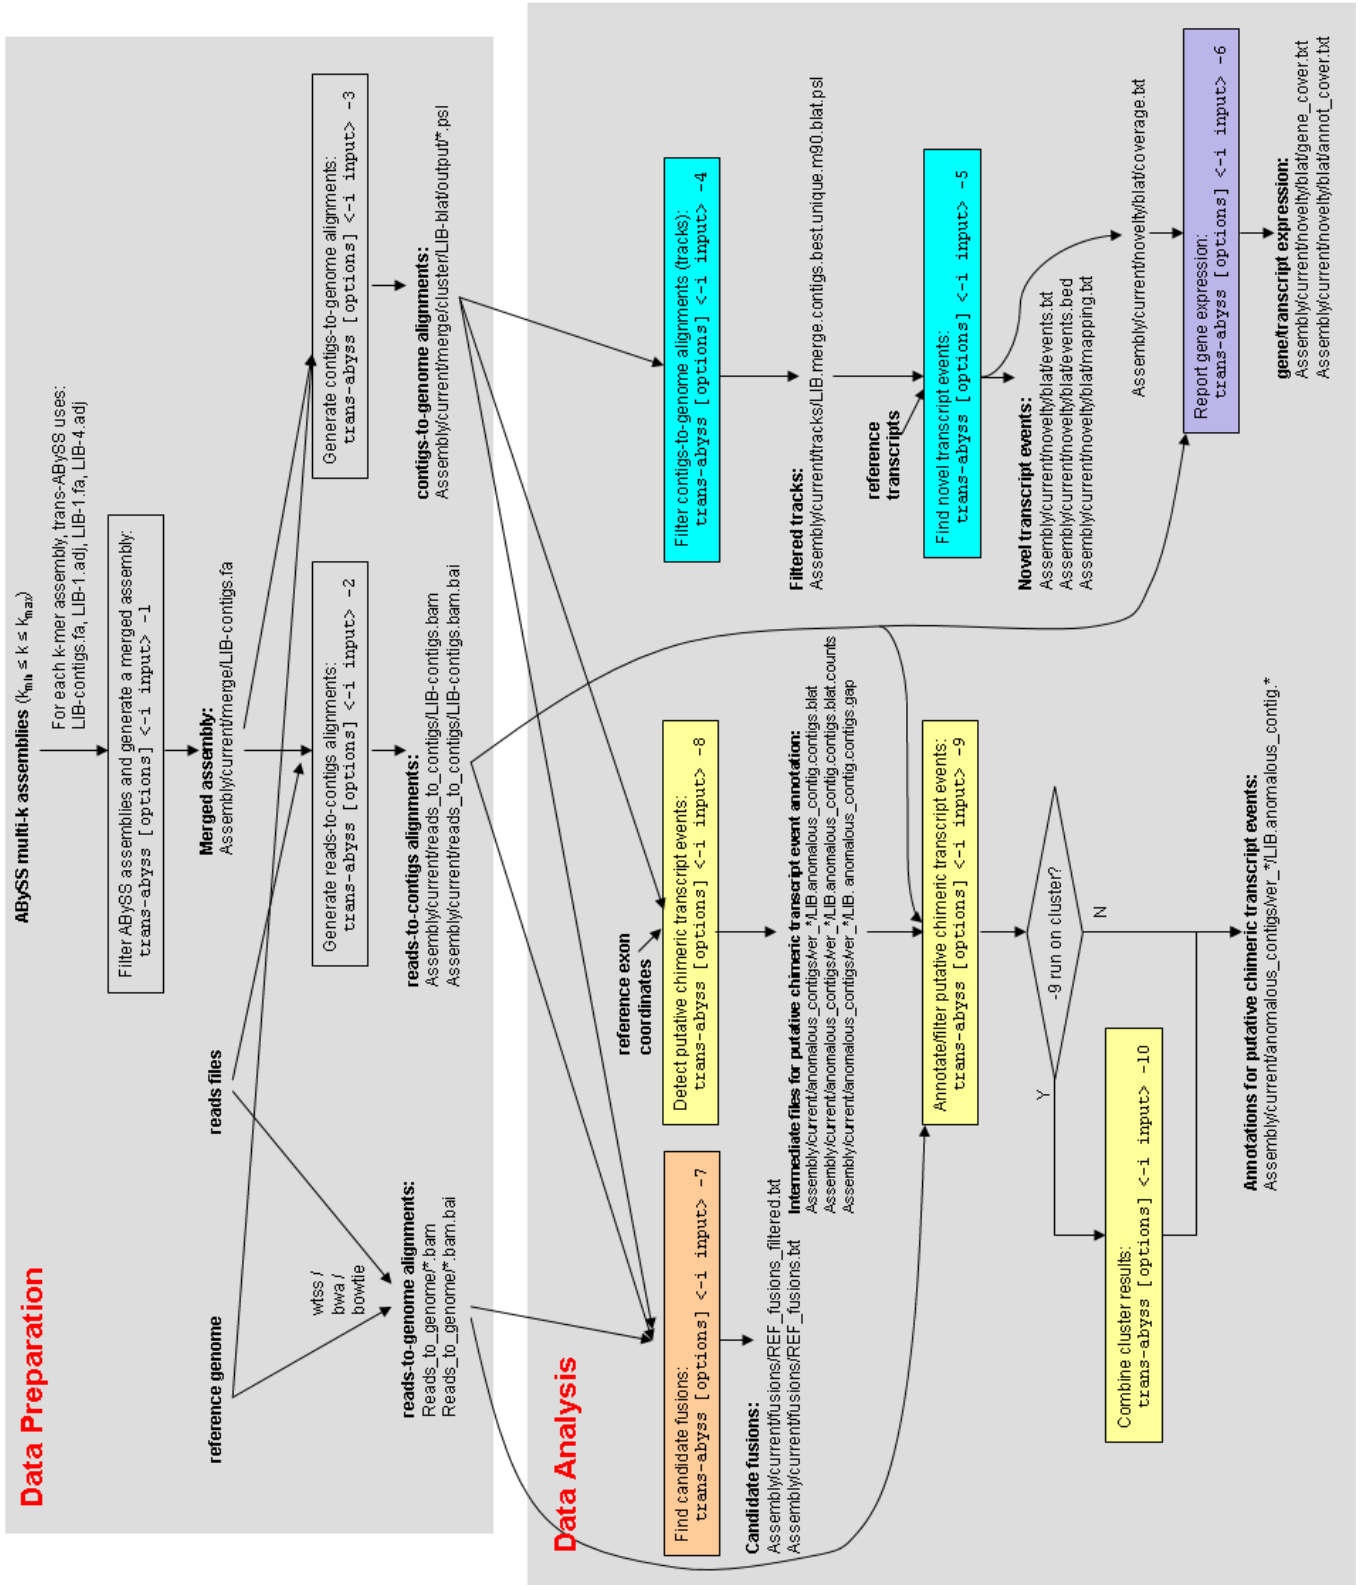


Figure 1. Trans-ABYSS pipeline overview.

2. ABySS assemblies and folder structure

Trans-ABYSS expects the output from ABySS multi-k assemblies for each library to be organized as follows: a single parent folder is used to hold all k-assemblies, where each subfolder is named “*kn*” (*n* is the value of *k*, e.g. k35) and stores the ABySS assembly output files for that particular *k* value (Fig. 2). In addition, in order to generate reads-to-contigs alignment, a simple text file named “in” should be present in the ABySS assembly folder, which lists all the paths of all input read files.

```
LIB0001/  
  k1/  
    LIB-contigs.fa  
    LIB-1.adj  
    LIB-1.fa  
    LIB-4.adj  
    [other ABySS output files]  
  k2/  
    LIB-contigs.fa  
    LIB-1.adj  
    LIB-1.fa  
    LIB-4.adj  
    [other ABySS output files]  
  ...  
  in
```

Figure 2. The ABySS assembly folder structure that Trans-ABYSS expects. Each ‘k’ folder holds the output of an ABySS assembly that was generated using that *k* value. Here, schematic folder ‘k’ names are shown; typical names might be: k26, k27,

The read files specified by the “in” file can be in any of the following formats: bam, qseq, export, or fastq. They can be compressed using gzip or bzip2 (with “.gz” or “.bz2” extensions). Fig. 3 shows an example “in” file.

```
/archive/solexa1_4/analysis2/HS1136/3153YAAXX_2/  
3153YAAXX_2_1_export.txt.gz  
/archive/solexa1_4/analysis2/HS1136/3153YAAXX_2/  
3153YAAXX_2_2_export.txt.gz  
/archive/solexa1_4/analysis3/HS1136/42HVVAAXX_1/  
42HVVAAXX_1_1_export.txt.gz  
/archive/solexa1_4/analysis3/HS1136/42HVVAAXX_1/  
42HVVAAXX_1_2_export.txt.gz  
/archive/solexa1_4/analysis3/HS1136/42HVVAAXX_2/  
42HVVAAXX_2_1_export.txt.gz
```



```
/archive/solexa1_4/analysis3/HS1136/42HVVAAXX_2/  
42HVVAAXX_2_2_export.txt.gz
```

Figure 3. An example “in” file that specifies paths to all input read files.

Once the “in” file is created, ABySS can be run with multiple *k* values on the same set of read files. For convenience, an example shell script “run-abyss” is included in Trans-ABySS “utilities” folder, which demonstrates how to generate ABySS multi-*k* assemblies in the required folder structure:

```
for k in {26..49}; do mkdir k$k; cd k$k; abyss-pe  
in=`paste -sd' ' in` OVERLAP_OPTIONS=--no-scaffold  
SIMPLEGRAPH_OPTIONS=--no-scaffold E=0 n=10 v=-v; cd ..;  
done
```

The following are example scripts to run ABySS on a computer cluster and generate multi-*k* assemblies in the required folder structures using qsub:

Script: utilities/qsub-150-64

```
#!/bin/sh  
set -eu  
qsub -N `basename $PWD` -t 33-49 ../qsub-150-k64
```

Script: utilities/qsub-150-k64

```
#!/bin/env qsub  
#$ -q mpi.q  
#$ -pe openmpi-1.3.1 16  
#$ -l hostname=qn*  
##$ -l mem_used=300M  
setenv PATH [PATH-TO-ABYSS-BINARIES]:$PATH  
setenv in `paste -sd' ' in`  
mkdir k$SGE_TASK_ID && cd k$SGE_TASK_ID && \  
    abyss-pe OVERLAP_OPTIONS=--no-scaffold  
SIMPLEGRAPH_OPTIONS=--no-scaffold E=0 n=10 v=-v
```

Note that ABySS scaffold option produces sequences with Ns and ambiguity codes, which may cause difficulty in some process of the current trans-ABySS pipeline. Thus we recommend either running ABySS with scaffold option off or breaking scaffolds after ABySS assembly is done. For more help on how to generate ABySS multi-*k* assemblies or other ABySS-related problems, please refer to the ABySS help group at <http://groups.google.com/group/abyss-users>.

3. Trans-ABySS folder structure

Trans-ABySS and ABySS have similar working directory structures (Fig. 4). The Trans-ABySS folder structure can be set up by running the “trans-abyss” script in the “wrappers” folder (see Section 3.3 below). Each k-mer sub-folder should initially be empty, and will be populated by the script to hold the processed assembly file from the corresponding ABySS k-mer assembly. The other directories are used to hold various Trans-ABySS output files, which will be discussed in detail in the following sections. The “log” file is used to keep track of trans-ABySS pipeline progress and is generated automatically when running various trans-ABySS scripts.

```
Project/
  Library/
    log
    Reads_to_genome/
      Reads_to_genome.bam

  Assembly/
    Abyss-1.2.1/
      source -> ABySS assembly path
      k1/
        Library-contigs.fa
      k2/
        Library-contigs.fa
      ...
      merge/
        Library-contigs.fa
      reads_to_contigs/
      tracks/
      novelty/
      fusions/
      anomalous_contigs/
      snv/
```

Figure 4. A typical Trans-ABySS working folder.

4. Running the Trans-ABySS Pipeline

4.1 Set up configuration files

Trans-ABySS analyses are performed on individual libraries, i.e. short-read sequencing datasets. However, to support work on a project that involves multiple related libraries (e.g. tens of patients for a disease), sets of libraries can be organized under a common project directory, and can share common run configuration settings. Settings are specified in configuration files, as follows.

The wrapper scripts use the following configuration files (in “configs” folder) to run the pipeline:

- **projects.cfg**
For each project, the user specifies the reference genome and the top-level ‘project’ directory (Fig. 5). A ‘project’ directory will contain a subdirectory for each of its libraries. In “projects.cfg”, default parameters for each script are specified in the “default” section. Defaults can be overridden by values set with each project. Uppercase words (e.g. MERGINGDIR) are used by calling scripts as templates that will be automatically replaced with appropriate values during a pipeline run.

```
[default]
merge.pl: VERDIR LIB contigs MERGINGDIR
align_parser.py: BLAT_DIR blat -n 1 -u -m 90 -d -k
TRACK_NAME -o PSL -f CONTIGS
...

[projectA]
topdir: /projects/projectA
reference: hg18

[projectB]
...
```

Figure 5. Example organization of “projects.cfg”

- **binaries.cfg**
Paths to external software are specified in “software: path” format, with one line for each executable (Fig. 6). An example file is provided in the distribution. Note in the example file, two versions of “python” are specified: “python” points to the executable of the correct version of python that runs on GSC’s cluster, and “python_xhost” points to the executable of the version of python that runs locally. Please replace all paths to point to proper binaries in your own computing environment. Do not change the name of software (the part before “:”).

```
[binaries]
python: /gsc/software/linux-x86_64/python-builder-2.6.4/bin/
python
python_xhost: /home/rshe/bin/bin/python
perl: /usr/local/bin/perl5.8.3
blat: /home/pubseq/BioSw/blat/blat34/blat
exonerate: /home/pubseq/BioSw/exonerate/exonerate-2.2.0-x86_64/
bin/exonerate
bwa: /home/pubseq/BioSw/bwa/bwa-0.5.6/bwa
bowtie: /home/rchiu/bin/bowtie-0.12.5/bowtie
```

```
bowtie_build: /home/rchiu/bin/bowtie-0.12.5/bowtie-build
samtools: /home/pubseq/BioSw/samtools/0.1.6/samtools
export2fq: /home/pubseq/BioSw/Maq/maq-0.7.1_x86_64-linux/
scripts/fq_all2std.pl export2std
biopython: /home/rchiu/python/biopython-1.52
mqsub: /opt/mqtools/bin/mqsub
```

Figure 6. Example of “binaries.cfg”

- cluster.cfg
This file is required when running jobs on a cluster. It specifies cluster job settings including the memory requirement for running different scripts on cluster, and the file used for each reference genome (Fig. 7). There is also an optional section “[email]” that specifies an email address. This is used to receive automatic notifications about the cluster job status when a job is completed or failed. If this section is missing or no email address is entered, no email notifications will be sent.

```
[memory]
merge.pl: 1G
fusion.py: 1G
model_matcher.py: 10G
reads_to_contigs.py: 1G
align_parser.py: 1G
cluster_align.py: 5G
gene_coverage.py: 1G

[genomes]
hg18: /var/tmp/genome/lymphoma/ucsc-hg18.fa
mm9: /var/tmp/genome/mouse/mm9_build37_mouse.fasta

[email]
email: someone@somewhere.com
```

Figure 7. Example of “cluster.cfg”

- align.cfg
Specify parameters used by each aligner when contigs are aligned to the reference genome.
- model_matcher.cfg
Specify settings of annotations for “model_matcher.py”, for finding novel transcripts and transcript events, relative to reference transcript annotations (Fig. 8). Each section specifies the annotation files and their order for a reference genome. See section 4.4.3A for more details.

```
[hg18]
k: knownGene_ref.txt
```

```
e: ensGene_ref.txt
r: refGene.txt
a: acembly_ref.txt
x: ensg.txt
order: k,e,r,a

[mm9]
k: knownGene_ref.txt
e: ensGene_ref.txt
r: refGene.txt
a: acembly_ref.txt
order: k,e,r,a
```

Figure 8. Example of “model_matcher.cfg”

- submitjobs.sh

The “submitjobs.sh” script in the “utilities” folder is used to run jobs on a computer cluster. The GSC cluster is currently a ~2000+ core (CPUs) Beowulf-style cluster running Red Hat Enterprise Linux 4. The infrastructure consists of a headnode to which users submit jobs, and the rest of the cluster consists of compute nodes that are involved only in computation. The headnode, called “apollo”, runs OSCAR 5.0pre with Sun Grid Engine 6.1u3. We submit jobs to apollo with this command:

```
submitjobs.sh apollo /opt/mqtools/bin/mqsub <job_dir>
<job_file> <job_name> <job_memory_requirement>
[email_address]
```

This script is used by several wrapper scripts when jobs need to be submitted to the computer cluster. Please adjust “submitjobs.sh” to be appropriate to your cluster.

4.2 Set up transcript annotations and genome sequence

Trans-ABySS compares genome alignments of assembled contigs to known annotations to discover transcript variants that are novel relative to reference transcript annotations.

Transcript annotation files are downloaded from the UCSC genome browser. Annotation files are organized by reference genome (Fig. 9). Currently trans-ABySS comes with “hg18” and “mm9” annotations (under “annotations” folder) that include Ensembl, UCSC, Aceview and Refseq transcripts.

```
annotations/
  hg18/
```

```
genome.fa -> ucsc-hg18.fa
knownGene.txt
knownGene_ref.txt
knownGene_ref.idx
knownGene_exons.txt
hg18.2bit
hg18_all_rmsk.coord
RNA.repeat.txt
splice_motives.txt
...
mm9/
...
shared/
    splice_motives.txt
README
```

Figure 9. Organization of the “annotations” folder.

4.2.1 Reference Genome

The reference genome fasta file is expected to be either copied or linked to genome folder as “genome.fa”. In addition, a “<GENOME>.2bit” genome file is needed when running scripts that find chimeric transcript events (transabyss stage 8 and 9), where <GENOME> is the reference name such as “hg18” or “mm9”.

4.2.2 Transcript Annotations

Transcript annotation files

(“knownGene.txt”, “ensGene.txt”, “acembly.txt”) (downloaded from UCSC) need to be modified slightly to include the common gene names at the end of each record

(“knownGene_ref.txt”, “ensGene_ref.txt”, “acebml_y_ref.txt”).

The “refGene.txt” file (also downloaded from UCSC) does not require such processing. The “README” file in the “annotations” folder describes how to do this.

4.2.3 Indexes (“.idx” files)

The transcript annotation files are indexed by genomic locations to expedite the searching and matching of contigs. Indexing is achieved by running the scripts in the “analysis/annotations” folder, one for each transcript model. For example, to index Ensembl transcripts, run the following command:

```
python ~mapper/trans-ABYSS/analysis/annotations/
ensembl.py ~mapper/trans-ABYSS/annotations/hg18/
ensGene.txt -i ~mapper/trans-ABYSS/annotations/hg18/
```

ensGene.idx

Currently, the following four scripts are supplied: `ensembl.py`, `knownGene.py` (for ucsc known genes), `aceview.py` (aceview genes), `refGene.py` (for refseq genes).

Transcript model files with other formats can be used if you create a custom parser. This can be easily done by modifying any of the existing parsers in the “analysis/annotations/” folder.

4.2.4 Exon coordinates

To find chimeric transcript events (trans-abyss stage 8, 9 and 10, see Section 4.4), each genome should have the exon coordinate annotation file (*-exons.txt). There can be multiple exon coordinate files, with one file being used as the primary annotation (required) and multiple other files can be used as secondary annotations (optional) that provide support for putative events that are detected based on the primary annotation.

These files are generated based on transcript annotation files downloaded from UCSC. The “README” file in the “annotations” folder describes how to generate them from UCSC files.

4.2.5 Repeat annotations

The chimeric transcript finder script (trans-abyss stage 8, 9 and 10) is also able to take into consideration of repeats in annotation. To enable this, each genome should have a coordinate file that contains coordinates of all repeats (<GENOME>_all_rmsk.coord) and/or coordinate file for RNA repeats (RNA.repeat.txt).

These files are generated using repeatmasker files downloaded from UCSC. The “README” file in the “annotations” folder also contains description on how to generate them.

An example config file can be seen at “sample_data/Insr_UTR/Insr_UTR/Assembly/current/anomalous_contigs/ver_15.5.0/config”, which lists the annotation files that are used by the chimeric transcript event finder script.

4.2.6 Splice motifs

The “splice_motives.txt” file specifies motives of known splice sites for each genome. It is used by “align_parser.py” and “model_matcher.py” (see Section 4.4.3A) for determining whether a splice site is novel or known.

There is also a shared “splice_motives.txt” file (under “shared” folder) that

can be used as common splice motifs, and the “splice_motifs.txt” file in each reference genome can be a symbolic link to this file.

4.3 Set up the input file

The wrapper scripts take an input file as the argument (Fig. 10). The input file specifies the libraries that need to be processed. Each library is specified by 5 fields in one line:

```
<library> <ABYSS-version> <ABYSS-assembly-location> <project> <min_read_length>
```

The fields are separated by spaces. The <ABYSS-assembly-location> is the name of the parent folder that holds all ABySS k-assemblies for that library, where each k-assembly is in a subfolder k_n (see Fig. 2). <min_read_length> is the smallest read length in the library (note that it is possible to have different-length reads in one library). If <min_read_length> is not specified, it is defaulted to 50.

```
LIB0001 1.2.1 /projects/ABYSS/assemblies/LIB0001 projectA 50
LIB0002 1.2.1 /projects/ABYSS/assemblies/LIB0002 projectA 50
LIB0003 1.2.1 /projects/ABYSS/assemblies/LIB0003 projectA 75
```

Figure 10. Example input file.

Note that library names should be unique in the same input file. To process a library with different parameters (e.g. with different ABySS-versions), each run should be put in a different input file.

4.4 Running trans-ABYSS

The current pipeline can be run with a wrapper script: “trans-abyss” (in the “wrappers” folder). It carries out analysis work in the following stages:

1. generate the transcriptome assembly
 - 1.1 set up the Trans-ABYSS folder structure (as in Section 3 above);
 - 1.2 process each ABySS k-mer assembly (see Section 4.4.1A below);
 - 1.3 merge all k-mer assemblies into one assembly (see Section 4.4.1B below);
2. align reads to contigs (see Section 4.4.2A below);
3. align contigs to the genome (see Section 4.4.2B below);
4. filter contigs-to-genome alignments and generate a track file that can be

- loaded to UCSC browser as a custom track (see Section 4.4.3A below);
5. find candidate novel transcript events (see Section 4.4.3A below);
 6. report gene expression levels (see Section 4.4.3B below).
 7. find candidate fusion genes (see Section 4.4.3C below).
 8. find putative chimeric transcript events (see Section 4.4.3D below).
 9. group and filter putative chimeric transcript events (see Section 4.4.3D).
 10. if stage 9 was run on a computer cluster, this step is needed to combine the cluster results (see Section 4.4.3D).

Each stage depends on the completion of previous stages. Please refer to the pipeline overview for the workflow (Fig. 1).

To run “trans-abyss”, use the following command:

```
trans-abyss [options] <-i input-file> <-1|-2|-3|-4|-5|-6|-7|-8|-9|-10>
```

where “input-file” is the input file described in Section 4.3 above. Specify the stage number to run trans-ABySS in corresponding stage: “-1”, “-2”, “-3”, “-4”, “-5”, “-6”, “-7”, “-8”, “-9”, “-10”.

The options are as follows:

- c <CLUSTER_HEAD>
name of the cluster head node to submit jobs (if applicable). For large datasets, it is necessary to run the jobs on a cluster.
- s <START_LIB>
“start library”, i.e. the name of the first library to be processed in the list of libraries in the input file; can be used in combination with “-num” option (see below)
- n <NUM>
number of libraries to process starting from “start library”. Use this option in combination with “-start” option to specify the libraries that need to be processed.
For example, given an input file as shown in [Figure 10], use “-start LIB0002 -num 2” to start from library “LIB0002” and process 2 libraries, i.e. LIB0002 and LIB0003.
If only “-start” option is specified, and “-num” is not specified, then all libraries from the start library onwards will be processed.
- l <LIB>
library that needs to be processed (for processing single library).
If only a single library needs to be processed, use this option instead of “-start” and “-num” combination.
Running “trans-abyss” without “-start”, “-num”, or “-lib” options will process all libraries in the input file.
- h | --help

Print help message.
--version
Print version message.

Note that most scripts in trans-ABYSS package can be run with “-h”, “--help” or “--man” option to get help on the description and usage of the script.

4.4.1 Setting up contigs for analysis

A. Processing ABYSS contigs for each k-mer assembly

For each assembly, a working set of contigs will comprise the following:

- all paired-end contigs; “Paired-end contigs” are contigs that were assembled during the pair-end stage of ABYSS (Simpson et al. 2009).
- all junction contigs; “Junction contigs” are single-end contigs that have one and only one neighbour on either side in the ABYSS assembly graph and should represent the lesser-expressed branch of a heterozygous allele (indels, splice variants, etc) in the transcriptome. Junction contigs are fully extended on either side and the extended contigs that are bigger than the read length of the library (or the minimum if multiple read lengths exist) are kept .
- single-end contigs of length greater than 150bp;
- single-end contigs that are not “islands” and are between $(2k-1)$ bp and 150bp in length. “Islands” are contigs that have no neighbours in the ABYSS graph.

The contig set can be generated by running “assembly.py” in the “utilities” folder:

```
assembly.py <library> -d <ABYSS_assembly_path>/k50 -o  
<trans-ABYSS_working_folder>/k50/library-contigs.fa -k  
50 -j <read_length + 1>
```

Running this command for each ABYSS k-mer assembly will generate a single FASTA file (named “LIBRARY-contigs.fa”) in the corresponding trans-ABYSS k-mer directory. This filtered contig set will be used for all downstream analysis.

“assembly.py” makes use of two ABYSS-related programs “MergeContigs” and “SimpleGraph” which are included in trans-ABYSS “utilities” folder. The “MergeContigs” program is the same program as in the ABYSS package and can also be downloaded from ABYSS download page. The “SimpleGraph” program is different from the version that is included in the ABYSS package, thus the source file “SimpleGraph.cpp” is also included in the “utilities” folder if there is need to re-compile it.

Note: “assembly.py” reads from several ABySS files in addition to the final ABySS fasta file, including “library-4.fa”, and “library-4.adj”. To process single-end ABySS assemblies that do not have “library-4.adj”, please create a symbolic link “library-4.adj” in the ABySS assembly folder that points to “library-3.adj”.

B. Creating the merged assembly

After each assembly has been processed, sets of contigs for assemblies across a range of k values are then merged to create a smaller, non-redundant contig set. The merging algorithm (“merge.pl”), which is described in the manuscript (Robertson et al. 2010), uses Blat to perform iterative pairwise alignments between assemblies:

```
merge.pl <trans-ABYSS_working_folder> <library> contigs
<trans-ABYSS_working_folder>/merge/merging
```

The final result is another FASTA file (also named “LIBRARY-contigs.fa”, but under the “merge” sub-folder) that consists of the non-redundant contig set from all k-mer assemblies.

C. Using the wrapper

For convenience, the wrapper script “trans-abyss” is set up to call “assembly.py” and “merge.pl” automatically in its stage “-1” as follows:

```
trans-abyss [options] <-i input-file> -1
```

The options are specified in Section 4.4 above. This command will run through all k-mer assemblies and merge the results into the final FASTA file.

4.4.2 Contig and read alignments

A. Read alignments to contigs

Read alignments to contigs are required for providing evidence ‘support’ for novel transcript events, and for estimating gene-level expression. Trans-ABYSS currently uses Bowtie in single-end mode to perform read-contig alignments. Because contigs can overlap, we allow multi-mapping, but require exact match alignments.

The wrapper script “trans-abyss” can be used to perform reads-to-contigs alignments as follows:

```
trans-abyss [options] <-i input-file> -2
```

B. Contig alignments to a reference genome

All the analyses described below (except for polyadenylation sites) require that assembled contigs be aligned to the reference genome. Trans-ABYSS currently supports Blat and exonerate aligners. However, outputs from other aligners that can generate .psl outputs (e.g. GMAP) can be treated as Blat outputs (by specifying the aligner as “blat” when required) and so can be processed by trans-ABYSS.

As noted in the manuscript (Robertson et al. 2010), to minimize the time required to review candidate novel transcript events, it is important that a contig aligner have a low error rate, and that its error rate be addressed.

Because even after merging there will typically be a large number of contigs, contig alignments are usually performed in parallel on a computer cluster. Computing systems at different laboratories will differ, and the “submitjobs.sh” script in “utilities” should be tailored by users to suit their cluster configuration.

To run BLAT alignments on a cluster, we split the merged assembly file into many smaller files, and run each job independently. The current default is to separate the assembly into 1,000 contigs per file. These files and their corresponding cluster job scripts and output can be found in the “merge/cluster/<LIB-blat-dir>” subdirectory in the trans-ABYSS working directory (Fig. 4):

```
merge/cluster/<LIB-blat-dir>/input
merge/cluster/<LIB-blat-dir>/jobs
merge/cluster/<LIB-blat-dir>/output
```

The “trans-abyss” wrapper script can be used to perform contig-to-genome alignments as follows:

```
trans-abyss [options] <-i input-file> -3
```

Some alignment jobs may not finish successfully on the cluster. To check whether all BLAT alignment jobs for a certain library are finished completely, use the tool “check_complete_blat.pl” (supplied in the “utilities” folder) as follows:

```
check_complete_blat.pl <LIB-blat-dir>
```

<LIB-blat-dir> is the name of the directory that holds the inputs, job scripts, and outputs of the blat jobs.

C. Aligning reads to a reference genome

Mate-pair read alignments to a reference genome are directly used as supporting evidence to rank fusion gene candidates. A fusion candidate that is well supported by mate-pairs will be prioritized for manual review.

The current pipeline does not include code for handling exon-exon junctions for such read alignments. For the results in the publication, we make a BWA-aligned .bam format file, and a .bigWig file derived from it, available for download from the Trans-ABYSS software v1.0 download page. These were generated with an internal GSC pipeline (unpublished).

4.4.3 Transcriptome assembly analysis

Trans-ABYSS currently offers the following functionality.

A. Identify candidate transcript structures that are novel relative to one or more sets of annotated transcript models (e.g. RefSeq, Ensembl, ...).

We recommend filtering contig alignments to retain only the best alignment that is unique (i.e. a contig cannot align to multiple genomic locations with the same score) and covers the majority of the length of the contig (e.g. 90%):

```
python align_parser.py blat_output_dir/blat_output_file
blat -n 1 -u m90 -d -k "track name" -o filtered.psl -f
merged_contigs.fa
```

The wrapper script “trans-abyss” can be used to run this job:

```
trans-abyss [options] <-i input-file> -4
```

After filtering, the resulting PSL-format file can be loaded into the UCSC genome browser for review, and can also be compared to reference transcript model files (e.g. in UCSC gene table format) by the “model_matcher.py” script in order to find novel transcripts and transcript events in the contig alignments:

```
model_matcher.py filtered_track.psl genome -l -d -o
output_dir -f merged_contigs.fa -r
```

The wrapper script “trans-abyss” can be used to do this job by specifying stage “-5”:

```
trans-abyss [options] <-i input-file> -5
```

The output directory will contain:

1. “mapping.txt” – details the mappings of contigs to known transcripts
2. “events.txt” – reports novel transcript variants relative to all transcript models (e.g. skipped exons, novel exons, ...)
3. “events.bed” – novel transcript variants in bed format
4. “coverage.txt” – transcript coverage statistics

The “mapping.txt” file is a text file where each line reports a match between a contig and a transcript, for example:

```
k38:11 matches uc009krd.2(Insr) model:k(wt:4) in 1 blocks total_blocks=1 total_exons=21  
partial_match coord:chr8:3154550-3154852 score:2.0 events:0 coverage:0.032
```

The format of each line is:

```
<contig> matches <transcript>(<gene>) model:<model abbreviation> (wt:<model  
weight> in <number aligned blocks> blocks total blocks=<total alignment blocks>  
total exons=<total number exons> <match> coord:<contig alignment coordinate>  
score:<score> events:<number of events> coverage:<coverage>
```

where:

<contig> = contig id

<transcript> = transcript id

<gene> = gene symbol

<model abbreviation> = gene model abbreviation, as specified in configuration file for model_matcher.py (e.g. ‘k’ = known genes, ‘e’ = Ensembl, ‘r’ = Refseq, ‘a’ = Aceview)

<model weight> = determined by order of models used for matching, which is specified in “configs/model_matcher.cfg” file (the models are specified in the order from highest to lowest weight). This serves as a tie-breaker when contigs are aligned with the same score to different gene models, in which case the gene model with the highest weight will be considered the best match.

<number aligned blocks> = number of alignment blocks matching exons

<total alignment blocks> = total number of alignment blocks in contig alignment

<total number exons> = total number of exons in transcript

<match> = “full_match”: all edges of alignment blocks aligned (outermost edges not included); “partial_match”: a subset of the total number of block edges aligned; “non_match”: none of the block edges aligned

<contig alignment coordinate> = coordinate of contig alignment in UCSC genome browser format (chr:start-end)

<score> = total number of edges perfectly aligned + 0.5 * number of splice site variants

<number of events> = number of novel splicing events

<coverage> = number of bases aligned / transcript length

The “events.txt” file reports an event per contig per line in space-delimited

columns, for example:

```
1.1 novel_exon k40:5-,2+,8- uc009krd.2(Insr) 11,12 12 chr8:3181702-3181737 36
3174889,3181701,--,3181738,3184950 orf:AGV...NPS,1399aa,162-4361,4199nt,0.88,1
```

The columns in each line are as follows:

1. event id: the same event shown by different contigs are grouped using event id, e.g. for event N, N.1, N.2, N.3 indicates that three contigs captured the same event

2. event type:

'AS3' = novel 3' splice site

'AS5' = novel 5' splice site

'AS53' = novel 5' splice site and novel 3' splice site

'skipped_exon' = exon skipping

'retained_intron' = retained intron

'novel_intron' = novel intron

'novel_exon' = novel exon

'novel_utr' = novel UTR

'novel_transcript' = novel transcript

3. transcript id (gene symbol)

4. alignment block number

5. exon number: exons are numbered in ascending order of genome coordinate, regardless of transcript orientation

6. event coordinate: overall event coordinate from start to end

7. splice-site info, may differ depending on the type of event

a) for AS/novel_utr/novel_intron/novel_exon/novel_transcript:

<splice site sequence>(<motif name>)

b) for retained_intron:

3x:False/True

c) for skipped_exon:

not applicable

8. surrounding coordinate: event region masked in "--", surrounded by neighboring coordinates e.g. <upstream neighbour start>,<upstream neighbour end>--,<downstream neighbour start>,<downstream neighbour end>

9. longest open reading frame: <start 3 amino acids>...<end 3 amino acids>,<number amino acids>aa,<start base number of contig>-<end base number of contig>,<total number bases translated>nt,<fraction contig translated>,<orientation>

The "events.bed" file contains several UCSC-format .bed tracks with each track representing one type of novel event (listed in 2, above). Each line represents one event. For reviewing novel event predictions, it is helpful to load this file into the UCSC genome browser, along with the contig alignments.

The "coverage.txt" is a tab-delimited txt file that reports the coverage of individual transcripts by contigs, for example:

InsrandD630014A15Rik.cSep07	InsrandD630014A15Rik	219	547	k26:8
0.043 7	k33:2,k45:18,k42:8,k35:13,k38:11,k26:8,k44:20	0.400	8.2	

The columns are as follows:

1. transcript
2. gene
3. total_coverage - total number of bases of transcript covered by contig
4. transcript_length - transcript length in base pairs
5. best_contig - best contig covering the transcript in terms of bases covered
6. best_contig_coverage - coverage of best contig
7. nbr_contigs - number of contigs covering the transcript
8. contigs - list of contigs covering the transcript
9. coverage - total bases covered (column 3) divided by transcript length (column 4)
10. normalized_k_coverage_best_contig - k-mer coverage divided by contig length

B. Estimate gene-level expression

Trans-ABYSS maps contigs to reference annotated transcripts by default. Gene-level expression is estimated by mapping to the gene the coverage on contigs aligned to a gene's transcripts. For the mouse adult liver data described in the Nature Methods publication (Robertson et al. 2010), trans-ABYSS expression values correlated closely to those from ALEXA-seq (Griffith et al. Nat Methods. 2010 7(10):843-7).

There are three parts to the calculations:

1. Align reads to contigs. This is part of the standard pipeline. Alignments can be generated by running the wrapper script "trans-abyss" with stage "-2" or by directly running the Python script "reads_to_contigs.py".
2. Map contig alignments to annotated transcripts to generate a 'coverage' file. This can be done by directly running "model_matcher.py" or running the wrapper "trans-abyss" with stage "-5".
3. Run `gene_coverage.py` to determine gene coverage:

```
python gene_coverage.py coverage-file-from-model_matcher
reads-to-contigs-bamfile track-file-used-for-model_matcher library-
name output-file
```

This step can also be run with the wrapper "trans-abyss" as follows:

```
trans-abyss [options] <-i input-file> -6
```


The output of "gene_coverage.py" is stored in "gene_cover.txt" as a tab-delimited text file with the following 5 columns:

1. gene name
2. number of reads mapped to gene
3. total read bases mapped to gene
4. union of contig alignment block lengths
5. normalized coverage (column 3 / column 4).

C. Identify candidate gene fusion events.

Split genomic alignments of contigs are reported as candidate gene fusion events. To generate a file with such candidate events, run:

```
fusion.py blat_out_dir output -l library -B genomic_bam -
b contig_bam
```

To minimize time spent in manually reviewing fusion candidates, we recommend filtering outputs on: a) the minimum number of read pairs from read-to-genome alignments, b) the minimum number of spanning reads (from read-to-contig alignments), c) the minimum percentage of identity in the contig-to-genome alignments, etc.:

```
fusion.py output filtered_output -X -F
```

The wrapper script "trans-abyss" can be used to do this job, which is equivalent to running the above two "fusion.py" commands:

```
trans-abyss [options] <-i input-file> -7
```

The output of "fusion.py" is a space-delimited txt file that reports one candidate fusion event per line, for example:

```
CTG:k30:253500+,938187+,1140721-(7146bp) TARGET:chr4:13210525-13238093,chr4:13209979-
13210251 CONTIG:1-6881,6874-7146 -,+
TO:0.00,CO:0.00,CC:1.00,I1:100.0,I2:100.0,AF1:0.96,AF2:0.04 READPAIRS:1001 SPAN_READS:2
```

Each line is in the following format:

```
CTG:<ctg id>(<ctg length>bp) TARGET:<region1 target coordinate>,<region
2 target coordinate> CONTIG:<region1 contig coordinate>,<region 2 contig
coordinate> <region1 orientaion>,<region2 orientation> TO:<overlap
target fraction>,<CO:<query overlap fraction>,<CC:<contig coverage
fraction>,<I1:<alignment 1 identity>,<I2:<alignment 2 identity>,<AF1:<alignment
fraction1>,<AF2:<alignment fraction2>
```

where:

<overlap target fraction> = fraction of overlap between target regions over sum of target regions aligned
<query overlap fraction> = fraction of overlap between contig regions over sum of contig regions aligned
<contig coverage fraction> = fraction of contig covered by both alignments
<alignment 1 identity> = identity of alignment 2
<alignment 2 identity> = identity of alignment 1
<alignment fraction1> = fraction of contig aligned to region 1
<alignment fraction2> = fraction of contig aligned to region 2

D. Identify putative chimeric transcript events.

Trans-ABYSS can be used to examine contig-to-reference genome alignments and identify putative chimeric transcript events, then annotate these events and add support from single-end read-to-contig and paired-end read-to-genome alignments.

(1) To identify putative chimeric transcript events from contig-to-reference alignments, run:

```
trans-abyss [options] <-i input-file> -8
```

This will prepare the output directories and essentially calls the “acfinder.pl” script with “-check_aligns” option.

The script will first ask for some user input in order to generate the configuration file needed for running chimeric transcript codes. The configuration file is a text file that lists various reference annotation files. In practice, for large projects that contain many different libraries, usually all libraries could share the same chimeric transcript configure file, since they are usually compared with the same reference. In such cases, the configuration file generated when processing the first library in the project can be stored as the standard configuration for all libraries of that project. Such standard configuration files are stored in “configs” directory with the name “acf_<project>_config”. Then during subsequent chimeric transcript runs for other libraries in the same project, the script will automatically look for such standard configuration file and ask whether the user wants to use that file or generate a new configuration file from scratch.

With the configuration file done, the alignment checking code will start to run. It is preferred to use a computer cluster (with the “-c” option) so that the alignments can be processed in parallel. Then the following script can be used to check whether all cluster jobs are done:

```
../analysis/acf/check_acf_status_4.py <JOB_DIR> [OPTIONS]
```

where <JOB_DIR> is the job directory that contains all cluster job results under trans-ABySS working folder (see Fig. 4), which is usually “<project>/<library>/Assembly/<abyss-version>/anomalous_contigs/current/cluster”.

The “check_acf_status_4.py” script also has the option to resubmit failed cluster jobs while checking the job status. To resubmit jobs to cluster, use:

```
../analysis/acf/check_acf_status_4.py <JOB_DIR> -r --new-jobs=<NEW_JOB_DIR> -c <CLUSTER>
```

where <CLUSTER> is the cluster head node to submit jobs to. <NEW_JOB_DIR> is a new job directory that will contain the newly submitted cluster jobs, such as “<project>/<library>/Assembly/<abyss-version>/anomalous_contigs/current/cluster1”.

(2) Then, to annotate and add support to the putative chimeric transcript events, run:

```
trans-abyss [options] <-i input-file> -9
```

which essentially calls the “acfinder.pl” script with “-add_support” option. It is recommended to run this step on a computer cluster (with the cluster option “-c”), otherwise it may take a very long time to finish. Note that if this step was run on a computer cluster, an extra step (very fast) is needed to combine the results from all cluster jobs. To check whether all “add_support” jobs on the cluster are done, use the “p2g_check_status.py” script in the “analysis/acf/” folder:

```
../analysis/acf/p2g_check_status.py <JOB_DIR>
```

where the <JOB_DIR> is the job directory that contains the cluster “add_support” jobs, usually “<project>/<library>/Assembly/<abyss-version>/anomalous_contigs/current/p2g_cluster”.

For cluster runs, when all the cluster “add_support” jobs are done, run the following command to combine all cluster results (note this is not needed if trans-abyss stage 9 is run locally without the cluster option):

```
trans-abyss [options] <-i input-file> -10
```

The output will be several files containing chimeric transcript event records. The “.data” file will contain all putative events identified. The “.pass” file will contain all putative events that passed the initial filtering values. If the “-split” option was used, there will also be individual files containing the passing events for each event type. These files will all have the same format. Each event record

will be made up of a group data line and one or more member data lines. Event records will be separated by blank lines. Member data lines will be one of two types: split member data lines and gapped member data lines. An example event record is shown below:

```
GRPNUM:413 TYPE:intrachr-opp-strand COORDS:chr2:234054565-234134606;chr2:234147236-234151219 MEMBERS:2 CONTIGS:2 ALIGNERS:1 CLENS:3678bp,846bp LOCAL:Y
PAIR_TO_GENOME_0:13 PAIR_TO_GENOME_10:13 EXON_BOUNDS:2
REPEATS:LINE;Simple_repeat STATUS:PASS VER:15.2.1 A00002
413a) intrachr-opp-strand contigs:k30:1147673(3678bp:2k+3618) ctg:1-3420(3420bp)=chr2:234054565-234134606(80042bp,+);AF:0.93,PID:99.9 ctg:3357-3678(322bp)=chr2:234151219-234147236(3984bp,-);AF:0.09,PID:100.0 ALIGN:blat,neighbour READ_TO_CTG:8 READ_TO_CTG_UNIQUE:8(74bp:0.11) CTGLOCAL:Y NUM_EVENTS:1 OVERLAPPING_GENES:USP40;UCH
CO:64bp(0.199),CR:1.00,JD:16613,TD:12630,QG:1,SP:2,MM:0,GF:N JUNCTION:chr2:234134606(up)-chr2:234151219(up)
413b) intrachr-opp-strand contigs:k29:1253630(846bp:2k+788) ctg:1-588(588bp)=chr2:234127665-234134606(6942bp,+);AF:0.70,PID:100.0 ctg:525-846(322bp)=chr2:234151219-234147236(3984bp,-);AF:0.38,PID:100.0 ALIGN:blat,neighbour READ_TO_CTG:8 READ_TO_CTG_UNIQUE:8(74bp:0.11) CTGLOCAL:Y NUM_EVENTS:1 OVERLAPPING_GENES:USP40;UCH
CO:64bp(0.199),CR:1.00,JD:16613,TD:12630,QG:0,SP:2,MM:0,GF:N JUNCTION:chr2:234134606(up)-chr2:234151219(up)
```

For a detailed specification of the output format, please refer to the file “analysis/acf/README”.

4.4.4 Additional Trans-ABYSS Functions

Trans-ABYSS provides some additional functions that are currently not handled by the wrapper scripts, but can be run separately.

A. Identify candidate SNVs and INDELS.

ABYSS outputs a “bubbles” file (“bubbles.fa”) that contains bubble contigs that represent potential SNVs. Alignments of the bubble contigs to both the genome and the paired-end contig set can be used to report SNVs:

```
python bubble.py k-len bubbles.fa align-genome.psl align-
contigs.psl blat -o logfile -n 1 -u -m 90 -b align-genome-
blast-output
```

Genome alignments of contigs can also be mined to extract potential SNVs and INDELS. Currently, this requires the contigs genome alignment in both psl and blast formats (note that BLAT is able to output in blast format):

```
python align_parser.py blat_output_psl_file blat -n 1 -u m90
-d -k "track name" -o filtered.psl -f merged_contigs.fa -b
```

```
blat_output_blast_file -v -w output
```

The .snv file generated from the above commands reports the following columns:

1. type: snv (single or multiple bases substitution), ins (insertion), or del (deletion)
2. chr: chromosome
3. chr_start: start coordinate of event
4. chr_end: end coordinate of event
5. strand: strand of alignment
6. ctg: contig id
7. ctg_len: contig length
8. ctg_start: contig start base
9. ctg_end: contig end base
10. len: length of event (bases)
11. change: e.g. G->A, or bases inserted or deleted
12. from_end: shortest distance of event from contig end

Work is in progress to rank candidate SNVs and INDELS by using read alignments as evidence. This functionality will be made available in a future version of the pipeline.

B. Identify candidate polyadenylation sites.

Polyadenylation site candidates are detected using a combination of Perl scripts, the BWA short-read aligner and UNIX commands. To perform the basic operations, a configuration file needs to be set up that points to the locations of BWA and the reference transcript models (Refseq, Ensembl, etc. sequence files in FASTA). The wrapper scripts 'polyareads.pl' and 'polyafinder.pl' can run the necessary commands.

Below, we outline the workflow. For more detailed information on each script, please refer to their respective perldocs, and to the README in the "polyascripts" folder under the "analysis" folder.

The "analysis/polyascripts" folder has this structure

```
bin
  Perl scripts
conf
  polyafinder.conf
eg_data
DISCLAIMER.txt
README.txt
```

The wrapper script requires as input: raw Illumina read sequence files (or FASTQ files), reference transcript sequences in FASTA format, and contig FASTA sequences from the assembly pipeline.

As described in the publication's Supplementary Information (Robertson et al. 2010), the method uses two types of reads: paired-end mate (PAM) and end-junction (EJ).

Run the wrapper script to extract PAM reads from the read files:

```
polyareads.pl -p -f <FORWARD_READS> -r <RIGHT_READS> [-F  
<FORWARD_READS2> -l ...] [-r <RIGHT_READS2> -r ...] [-conf  
<CONFIGFILE>]
```

Perform the PAM read alignment to reference transcript sequences:

```
polyafinder.pl -f <FORWARD_PAM> -r <REVERSE_PAM> -t  
<TRANSCRIPT> [-t <TRANSCRIPT>] -a [-cont <CONTIGFILE> -  
capp]
```

Extract EJ reads from the read files:

```
polyareads.pl -e -f <FORWARD_READS> -r <RIGHT_READS> [-F  
<FORWARD_READS2> -l ...] [-r <RIGHT_READS2> -r ...] [-conf  
<CONFIGFILE>]
```

Perform the EJ read alignment to reference transcript sequences:

```
polyafinder.pl -e -a -f <FORWARD_EJ> -r <REVERSE_EJ> -t  
<TRANSCRIPT> [-t <TRANSCRIPT>] [-mf <FORWARD_EJ_MATE> -mr  
<REVERSE_EJ_MATE>]
```

Optionally, it is possible to create visual images of the reads mapped to the genome by creating .bed files with `getremappedBED.pl` or `samse2bed.pl`, .wig files with `bed2Wig.pl`. Both can be used as viewable tracks within the UCSC genome browser. To summarize high-hit alignments and generate genome browser URLs that facilitate reviewing predictions at UCSC, use the script `getpolyTmapcoord_extracols.pl`.

Align reads to the genome by either extracting from the raw Illumina read file:

```
cat <INPUT_READS> | getremappedBED.pl <OUT_UNMAPPED_READS>  
<INPUT_FASTQ > <OUTPUT_BED>
```

or if no raw Illumina reads are available, use FASTQ to align to genome:

```
bwa aln <GENOME_FA> <INPUT_FASTQ> > <OUTPUT_SAI>  
bwa samse -n 20 <GENOME_FA> <OUT_SAI> <INPUT_FASTQ> |  
samse2bed.pl > <OUT_BED>
```

Then, convert the alignments to .wig:

```
cat <OUT_BED> | bed2Wig.pl > <OUT_WIG>
```

To rank the transcripts by number of PAM/EJ reads mapped and to create UCSC linked URLs:

```
getpolyTmapcoord_extracols.pl -t <IN_TSV> [-b <OUT_BED>]
[-z <ZOOM>] [-c CUTOFF1,C2,C3,...,Cn] [-u 'http:/
/genome.ucsc.edu/cgi-bin/hgTracks?
org=<ORGANISM>&db=<DB>&position='] [-o] [-d 500] >
<RANK_OUT>
```

To limit the number of false positives due to high poly-A and poly-T regions, use the script `calcPolyAinSeqs.pl` for the transcript and contig sequences, and use the script `findgenomicpolya.pl` for the genome.

Data sets

1. Large

MM0472 library at the SRA (Short Read Archive). This 147 M read PE dataset can be downloaded from <http://www.ncbi.nlm.nih.gov/sra/SRX017642?report=full>

Once downloaded, run ABySS with multiple k values (we ran every k from 26 to 50) and then use trans-ABYSS to process the assemblies and perform analyses. The reads-to-genome (mm9) alignments and bigWig files can be downloaded from Trans-ABYSS software v1.0 release web page.

2. Small

2.1 Insr_UTR

We include a dataset that consists of all 15,024 reads that aligned with Bowtie to contigs that the pipeline matched to annotated insulin receptor gene, `Insr`, including its UTR regions. Trans-ABYSS identified an exon in this gene that was novel when we discovered it, but was subsequently (but temporarily) included in 'UCSC gene' transcript models for this gene.

This set of reads can be assembled with ABySS, and processed and analyzed with Trans-ABYSS on a single CPU in less than 2 hours. All data/results are stored in "sample_data/Insr_UTR" folder and can be used for testing. You should be able to duplicate the results by running trans-ABYSS on your own machine.

2.2 Polyadenylation site analysis

In the `data_pam` folder, we include a small dataset that contains 25 PAM reads in FASTQ format, 4 genes in FASTA format, and 2 contig sequence in FASTA format. The PAM reads were chosen to illustrate finding novel polyadenylation sites. Specifically, the two 'gene' examples show candidate novel short 3' UTRs

while the 'contig' example shows a candidate novel lengthened 3' UTR. In the data_ej folder, we include 13 EJ reads in FASTQ format, and 1 gene in FASTA format. Similarly, these were chosen to illustrate the detection of novel short polyadenylation site.

The wrapper scripts `polyareads.pl` and `polyafinder.pl` can be used to extract and align the reads to the transcripts. Refer to `cmd.txt` for commands that can be used. `cmd.txt` can also be run as a shell script.

References

Biol I, Jackman SD, Nielsen CB, Qian JQ, Varhol R, Stazyk G, Morin RD, Zhao Y, Hirst M, Schein JE, Horsman DE, Connors JM, Gascoyne RD, Marra MA, Jones SJ. **De novo transcriptome assembly with ABySS**. *Bioinformatics*. 2009 Nov 1;25(21):2872-7.

Robertson G, Schein J, Chiu R, Corbett R, Field M, Jackman SD, Mungall K, Lee S, Okada HM, Qian JQ, Griffith M, Raymond A, Thiessen A, Cezard T, Butterfield Y, Newsome R, Chan SK, She R, Varhol R, Kamoh B, Prabhu A-L, Tam A, Zhao Y-J, Moore R, Hirst M, Marra MA, Jones SJM, Hoodless PA, Birol B. **De novo Assembly and Analysis of RNA-seq data**. *Nature Methods*, *in press*.

Simpson JT, Wong K, Jackman SD, Schein JE, Jones SJ, Birol I. **ABySS: a parallel assembler for short read sequence data**. *Genome Res*. 2009 Jun;19(6):1117-23.