

2013 Workshop on Genomics, Cesky Krumlov

Part 1: Short read genomics: Remapping

Instructors:

- Konrad Paszkiewicz k.h.paszkiewicz at exeter ac uk

Objectives:

By the end of the workshop you will be expected to:

- Understand how short reads are generated
- Interpret FASTQ quality metrics
- Remove poor quality data
- Trim adaptor/contaminant sequences from FASTQ data
- Count the number of reads before and after trimming and quality control
- Align reads to a reference sequence to form a SAM file (Sequence Alignment/Map file) using BWA
- Convert the SAM file to BAM format (Binary Alignment/Map format)
- Identify and select high quality SNPs and Indels using SAMtools
- Identify missing or truncated genes with respect to the reference genome
- Identify SNPs which overlap with known coding regions

Background:

- This course assumes that you have successfully completed a Unix and Perl workshop or have equivalent knowledge.

1.1 Introduction

Welcome to the Short Read Genomics workshop. Generating reams of data in Biology is easy these days. In little more than a fortnight we can generate more data than the entire human genome project generated in over a decade of work. Making biological sense out of that data, understanding its limitations and how the analysis algorithms work is now the major challenge for researchers. The aim of this workshop is to take you through an example project. On the way you will learn how to evaluate the quality of data as provided by a sequencing facility, how to align the data against a known and annotated reference genome and how to perform a *de-novo* assembly. In addition you will also learn how to compare results between different samples.

This workshop is broken into 6 parts. You should feel free to take as long as you like on each part. It is much more important that you have a thorough understanding of each part, rather than try to race through the entire workshop. We can always schedule additional sessions if necessary.

The six parts are:

1. Introduction to Illumina sequencing-by-synthesis
2. Setting up your Amazon EC2 instance
3. Remapping a strain of *E.coli* to a reference sequence
4. Assembly of unmapped reads
5. Complete *de-novo* assembly of all reads
6. Repeating parts 3-5 on two other strains and comparing them

For this first workshop we will assume little background knowledge, save a basic familiarity with the Linux operating system and the Amazon cloud. We will cover the basics of how genomic DNA libraries are generated and sequenced, and the principles behind short read paired-end sequencing. We will look at why data can vary in quality, why adaptor sequences need to be filtered out and how to quality control data.

In the second part we will set up and start your Amazon EC2 instance. This will be familiar to those of you who have completed the Unix and Perl workshop.

In the third part we will take the plunge and align the filtered reads to a reference genome, call variants and compare them against the published genome to identify missing, truncated or altered genes. This will involve the use of a publicly available set of bacterial *E.coli* Illumina reads and reference genome.

In parts 4 and 5 we will look at how one can identify novel sequences which are not present in the reference genome. In part 6, you will be asked to repeat the steps in workshops 1, 2 and 3 on 2 other data sets and to compare the results.

A word on notation. If you see something like this:

```
cd ~/genomics_tutorial/reference_sequence
```

It means, type the highlighted text into your terminal.


1.2 Principles of Illumina-based sequencing:

There are several second generation (i.e. non-Sanger) sequencers currently on the market. These include the Life 5500 (formerly known as the ABI SOLiD), the Roche 454 GS FLX and 454 Junior, Ion Torrent, and the Illumina HiSeq and MiSeq systems. All of these systems rely on making hundreds of thousands of clonal copies of a fragment of DNA and sequencing the ensemble of fragments using DNA polymerase or in the case of the SOLiD via ligation. This is simply because the detectors (basically souped-up digital cameras), cannot detect fluorescence (Illumina, SOLiD, 454) or pH changes (Ion Torrent) from a single molecule.

The 'third-generation' Pacific Biosciences SMRT (Single Molecule Real Time) sequencer, is able to detect fluorescence from a single molecule of DNA. However, the machine weighs 2 tons, produces 1/30000th of the data of an Illumina run and has a 15-20% error rate. Nonetheless the sequencer does have its uses and doubtless the error rates will improve with time, but currently the second generation sequencers are dominating the sequencing world.


We will only look at the Illumina sequencing pipeline here, but the basic principles apply to all other sequencers. If you would like further details on other platforms then I recommend reading *Mardis ER. Next-generation DNA sequencing methods. Annual Reviews Genomics Hum Genet 2008; 9 :387–402* Other papers are listed at <http://biosciences.ex.ac.uk/facilities/sequencing/usefulresources/>

A typical sequencing run would begin with the user supplying 1-10 ug of genomic DNA to a facility along with quality control information in the form of an Agilent Bioanalyser trace or gel image and quantification information. The following flowchart illustrates the basic workflow.



1.3 DNA Library preparation

For most sequencing applications, paired-end libraries are generated. Genomic DNA is sheared into 300-500bp fragments (usually via sonication) and size-selected accordingly. Ends are repaired and an overhanging adenine base is added, after which oligonucleotide adaptors are ligated. In many cases the adaptors contain unique DNA sequences of 6-8bp which can be used to identify the sample if they are 'multiplexed' together for sequencing. This type of sequencing is used extensively when sequencing small genomes such as those of bacteria because it lowers the overall per-genome cost.



A) Steps a through e explain the main steps in Illumina sample preparation: a) the initial genomic DNA, b) fragmentation of genomic DNA into 500bp fragments, c) end repair, d) addition of A bases to the fragment ends and e) ligation of the adaptors to the fragments.

B) Overview of the automated the size selection protocol: The first precipitation discards fragments larger than the desired interval. The second precipitation selects all fragments larger than the lower boundary of the desired interval.

Borgström E, Lundin S, Lundeberg J, 2011 Large Scale Library Generation for High Throughput Sequencing. *PLoS ONE* 6: e19119. doi:10.1371/journal.pone.0019119



1.4 Sequencing

(adapted from Margulis, E.R., reference below)

Once sufficient libraries have been prepared, the task is to amplify single strands of DNA to form monoclonal clusters. The single molecule amplification step for the Illumina HiSeq 2000 starts with an Illumina-specific adapter library and takes place on the oligo-derivatized surface of a flow cell, and is performed by an automated device called a cBot Cluster Station. The flow cell is either a 2 or 8-channel sealed glass microfabricated device that allows bridge amplification of fragments on its surface, and uses DNA polymerase to produce multiple DNA copies, or clusters, that each represent the single molecule that initiated the cluster amplification.

Separate or multiple libraries can be added to each of the eight channels, or the same library can be used in all eight, or combinations thereof. Each cluster contains approximately one million copies of the original fragment, which is sufficient for reporting incorporated bases at the required signal intensity for detection during sequencing. The Illumina system utilizes a sequencing- by-synthesis approach in which all four nucleotides are added simultaneously to the flow cell channels, along with DNA polymerase, for incorporation into the oligo-primed cluster fragments (see figure below for details). Specifically, the nucleotides carry a base-unique fluorescent label and the 3'-OH group is chemically blocked such that each incorporation is a unique event. An imaging step follows each base incorporation step, during which each flow cell lane is imaged in three 100-tile segments by the instrument optics at a cluster density of 600,000-800,000 per mm². After each imaging step, the 3' blocking group is chemically removed to prepare each strand for the next incorporation by DNA polymerase. This series of steps continues for a specific number of cycles, as determined by user-defined instrument settings, which permits discrete read lengths of 40–100 bases. A base-calling algorithm assigns sequences and associated quality values to each read and a quality checking pipeline evaluates the Illumina data from each run.

The figure on the following page summarises the process:




The Illumina sequencing-by-synthesis approach: Cluster strands created by bridge amplification are primed and all four fluorescently labeled, 3'-OH blocked nucleotides are added to the flow cell with DNA polymerase. The cluster strands are extended by one nucleotide. Following the incorporation step, the unused nucleotides and DNA polymerase molecules are washed away, a scan buffer is added to the flow cell, and the optics system scans each lane of the flow cell by imaging units called tiles. Once imaging is completed, chemicals that effect cleavage of the fluorescent labels and the 3'-OH blocking groups are added to the flow cell, which prepares the cluster strands for another round of fluorescent nucleotide incorporation.

Base-calling:

Base-calling involves evaluating the raw intensity values for each fluorophore and comparing them to determine which base is actually present at a given position during a cycle. To call bases on the Illumina or SOLiD platform, the positions of clusters need to be identified during the first few cycles. This is because they are formed in random positions on the flowcell as the annealing process is stochastic. This is in contrast to the 454 system where the position of each cluster is defined by steel plate with picolitre sized holes in which the reaction takes place.

If there are too many clusters the edges of the clusters will begin to merge and the image analysis algorithms will not be able to distinguish one cluster from another (remember, the software is dealing with upwards of half a million clusters per square millimeter – that's a lot of dots!).




The above figure illustrates the principles of base-calling from cycles 1 to 9. If we focus on the highlighted cluster, one can observe that the colour (wavelength) of light observed at each cycle changes along with the brightness (intensity). This is due to the incorporation of complementary ddNTPs containing fluorophores. So at cycle 1 we have a T base, at 2 a G base and so on. If the colour or intensity is ambiguous the sequencer will mark it as an N. Other clusters are also visible in the images; these will represent different monoclonal clusters with different sequences.

The base calling algorithms turn the raw intensity values into T,G,C,A or N base calls. There are a variety of methods to do this and the one mentioned here is by no means the only one available, but it is often used as the default method on the Illumina systems. Known as the 'Chastity filter' it will only call a base if the intensity divided by the sum of the highest and second highest intensity is less than a given threshold (usually 0.6). Otherwise the base is marked with an N. In addition the standard Illumina pipeline will reject an entire read if two or more of these failures occur in the first 4 bases of a read (it uses these cycles to determine the boundary of a cluster).

Note that these processes are carried out at the sequencing facility and you will not need to perform any of these tasks under normal circumstances. They are explained here as useful background information.

CHASTITY formula:


$$C = \frac{I_A}{I_A + I_B}$$



1.5 What are paired-end reads and why are they necessary?

Paired-end sequencing is a remarkably simple and powerful modification to the standard sequencing protocol. It is nearly always worth obtaining paired-end reads if performing genomic sequencing. Typically sequencers of any type are only able to sequence a portion of DNA (e.g. 100bp in the case of Illumina) before the fidelity of the enzyme and de-phasing of clusters (see later) increase the error rate beyond tolerable levels. As a result, on the Illumina system, a fragment which is 500bp long will only have the first 100bp sequenced.

If the size selection is tight enough and you know that nearly all the fragments are close to 500bp long, you can repeat the sequencing reaction from the other end of the fragment. This will yield two reads for each DNA fragment separated by a known distance I.e:



The added information gained by knowing the distance between the two reads can be invaluable for spanning repetitive regions. In the figure below, the light coloured regions indicate repetitive sections of DNA. If a read contains only repetitive DNA, an alignment algorithm will be able to map the read to many locations in a reference genome. However, with paired-end reads, there is a greater chance that at least one of the two reads will map to a unique region of DNA. In this way one of the reads can be used to anchor the other read in the pair and help resolve the repetitive region. Paired-end reads are often used when performing de-novo genome sequencing (i.e. when a reference is not available to align against) because they enable contiguous regions of DNA to be ordered, or when characterizing variants such as large insertions or deletions.


Other forms of paired-end sequencing with much larger distances (e.g. 10kb) are possible with so called 'mate-pair' libraries. These are usually used in specific projects to help order contigs in de-novo sequencing projects. We will not cover them here, but the principles behind them are similar.

Repetitive DNA

Unique DNA

Paired read maps uniquely

Single read maps to
multiple positions




1.6 Inherent sources of error

No measurement is without a certain degree of error. This is true in sequencing. As such there is a finite probability that a base will not be called correctly. There are several possible sources:

Frequency cross-talk and normalisation errors:


When reading an A base, a small amount of C will also be measured due to frequency overlap and vice-versa. Similarly with G and T bases. Additionally, from the figure below, it should be clear that the extent to which the dyes fluoresce differs. As such it is necessary to normalize the intensities. This normalisation process can also introduce errors.



Frequency response curve for A and C dyes
(Intensity y-axis and frequency on the x-axis)

Phasing/Pre-phasing:

This occurs when a strand of DNA lags or leads the other DNA strands within a cluster. This introduces additional background noise into the signal and reduces the intensity of the true base. E.g. below we have a cluster with 7 strands of DNA (very small, but this is just an example). Five strands are on a C-base, whilst 1 is lagging behind (called phasing) on a G base and the remaining strand is running ahead of the pack (confusingly called pre-phasing) on an A base. As such the C signal will be reduced and A and G boosted for the rest of the sequencing run. Too much phasing or pre-phasing (i.e. > 15-20%) usually causes problems for the base calling algorithm and result in clusters being filtered out.



Other issues:

- **Biases introduced by sample preparation** – your sequencing is only as good as your experimental design and DNA extraction. Also, remember that your sample will be put through several cycles of PCR before sequencing. This also introduces a potential source of bias.
- **High AT or GC content sequences** – this reduces the complexity of the sequence and can result in higher error rates
- **Homopolymeric sequences** – long stretches of a single base can make it difficult to determine phasing and pre-phasing rates. This can introduce errors in determining the precise length of a homopolymeric stretch of sequence. This is much more of a problem on the 454 and Ion Torrent than Illumina platforms but still worth bearing in mind. Especially if you encounter indels which have been called in homopolymeric tracts.
- Some motifs can cause loops and other steric clashes

See Nakamura et al, Sequence-specific error profile of Illumina sequencers Nuc. Acid Res. first published online May 16, 2011 doi:10.1093/nar/gkr344

1.7 Quality scores

To account for the possible errors and provide an estimate of confidence in a given base-call, the Illumina sequencing pipeline assigns a quality score to each base called. Most quality scores are calculated using the Phred scale. Each base call has an associated base call quality which estimates chance that the base call is incorrect.

Q10 = 1 in 10 chance of incorrect base call

Q20 = 1 in 100 chance of incorrect base call

Q30 = 1 in 1000 chance of incorrect base call

Q40 = 1 in 10,000 chance of incorrect base call

For most 454, SOLiD and Illumina runs you should see quality scores between Q20 and Q40. Note that these are only estimates of base-quality based on calibration runs performed by the manufacturer against a sample of known sequence with (typically) a GC content of 50%. Extreme GC bias and/or particular motifs or homopolymers can cause the quality scores to become unreliable.

Accurate base qualities are an essential part in ensuring variant calls are correct. As a rough and ready rule we generally assume that with Illumina data anything less than Q20 is not useful data and should be excluded from the analysis.

1.8 Reads containing adaptors

Some reads will contain adaptor sequences after sequencing, usually at the end of the read. This is usually because of short sample DNA fragments, which result in the polymerase reading into the adaptor region. Occasionally this can also happen because of mis-priming. It is important to remove or trim sequences containing these reads as the adaptor sequences can prevent reads mapping to a reference sequence and will adversely affect de-novo assembly.

2013 Workshop on Genomics, Cesky Krumlov

Part 2: Short read genomics: Starting your Amazon virtual machine

Instructors:

- Konrad Paszkiewicz k.h.paszkiewicz at exeter ac uk

Objectives:

By the end of this section you will be expected to:

- Log into the Amazon Console and start your instance of the Genomics workshop AWS Image
- Log in to an Amazon EC2 Linux instance using your own computer
- Continue with the Short read genomics tutorials at your own pace

Background:

- This course assumes that you have successfully completed the Unix and Perl workshop or equivalent knowledge

Introduction

For this workshop we will provide an overview of the cloud as described by Amazon and how, as researchers, we can use this flexible resource to get work done quickly and relatively inexpensively. We will dive right into starting up the pre-prepared Virtual Machine (VM) and logging into the Amazon management console. We will give you a whirlwind tour of the features of Amazon's cloud and then get you to start your Genomics AMI and log-in via the NX-client.

For this tutorial I borrowed documentation from the following sites:

- <http://ged.msu.edu/angus/tutorials/unix-and-ssh-and-scp.html>
- <http://aws.amazon.com/documentation>

Task 1 – Tour of Amazon's Cloud

In this section of the workshop we will log into Amazon's cloud (referred to as Amazon Web Services or AWS) and take a look at the various services offered by Amazon. These include:

- Elastic Cloud Compute (EC2): the service AWS is known for. It enables you to rent hourly Linux and Windows machines by the hour. Amazon now has also special High Performance Computing nodes (HPC) and Graphical Computing nodes (GPU nodes)
- Simple Storage Service (S3): a storage service, not particularly fast but great for storing large “buckets” of data for long term storage, sharing, or temporary storage for use between instances
- Elastic Block Storage (EBS): similar to S3 but limited in size (max 1TB), these are virtual hard drives that you can attach and detach very quickly to and from your running instances. Think of these as the USB flash drive of the cloud computing world
- A ton of other services that are geared towards building highly scalable and fault-tolerant web-based services. Many can be co-opted for use in research!

We have set up user accounts for each student in the class, take a look at the form given to you in your welcome packet. You need to know the console URL, your username, and your password. You will also find the access key and secret access key, both are used by various tools to log in and use the Amazon cloud programmatically.

Task 2 - Connecting to Your Personal VM

As part of the course we were extremely fortunate to have received funding from Amazon, enough to provide each student with \$100 of compute and storage resources. This should be plenty to cover the costs of the class. In this section of the tutorial we will log into the Amazon console web application, setup your account and VM instance, and login to kick the tyres.

The Rules

While \$100 is very generous we ask that each student adhere to the following guidelines to ensure we have enough resources for the duration of the workshop:

1. Please only launch a single VM instance of the type specified by the instructor at the beginning of the workshop and “stop” the instance at the end of the workshop. Do not leave the image running overnight or over a whole week – you will quickly run out of funding.

2. Please name your VM instance as I do in my demo. Including your name will make it easy to find your resources in the list of class resources
3. Please do not delete EBS volumes or terminate VM instances that don't belong to you

Step 1 – Logging Into the Console

In addition to being extremely comprehensive, the Amazon cloud has a very easy-to-use interface for interacting with all their cloud offerings. All you have to do is log into a web application and most of the functionality of the Amazon tools are available for you and very easy to use.

Username and passwords were given in your welcome packet along with the URL (which is repeated below).


We have setup a joint account for all the students on this course using something called Identity Access Manager. The nice thing about this is you have pretty much free access to the console and we can have very fine grain control on what your sub-accounts can and cannot do.

To get started go to the following URL and login with your username and password.

<https://genomics2013.signin.aws.amazon.com/console>

After logging in you'll be presented with a wide range of options. Most likely you've logged into the “EC2 Dashboard”. If you haven't, please click the drop down menu at the top and then click on “Amazon Elastic Compute Cloud (EC2)”. On this page you'll get a summary of the EC2 state for your account. You can see mine below.

Step 2 – Launch the “AMI”



From here we can create computers on Amazon's 'cloud'. What this means is that we can create as many computers as we like, start them, log-in to them, do some work, transfer data to/from them or destroy them altogether. Amazon worry about the hardware, power, cooling and maintenance – all we need to do is specify how powerful a computer we want (micro, small, large or extra-large). Amazon charge for each Gb stored every month and for each hour a machine is run. This can vary from a few cents per hour to a few dollars. Whilst it is very convenient if you are only doing analyses occasionally, at the moment it is still cheaper to have your own computing (although you then have the headache of maintaining it!).

The reason we are using the cloud here is that most of you will only be doing occasional analyses. In the case of high-throughput sequencing data (e.g. Illumina) you will find that your desktop PC may not be powerful enough to cope with the data. As such Amazon can offer a good alternative. It also means that you can start and stop your AMI from home and continue to work through the tutorial from there.

Once logged into the console we can select an Amazon Machine Image (AMI) Virtual Machine and launch it. Click on the “Launch Instance” button in the center. Then select ‘Launch Classic Wizard’.

Create a new instanceCancel

Select an option below:

**Launch Classic Wizard**

Continue to the classic wizard which provides you with the full list of AMIs as well as fine-grained control over how you would like your instance to be launched.

**Quicklaunch**

Select from a list of popular configurations to launch your instance into the cloud as quickly as possible.

Launch with the classic wizard

Request Instances WizardCancel

CHOOSE AN AMI

INSTANCES DETAILS

CREATE KEY PAIR

CONFIGURE FIREWALL

REVIEW

Choose an Amazon Machine Image (AMI) from one of the tabbed lists below by clicking its **Select** button.

Quick Start

My AMIs

Community AMIs

**Basic 32-bit Amazon Linux AMI 2011.02.1 Beta** (AMI Id: ami-8c1fecc5)
Amazon Linux AMI Base 2011.02.1, EBS boot, 32-bit architecture with Amazon EC2 AMI Tools.
Root Device Size: 8 GB

**Basic 64-bit Amazon Linux AMI 2011.02.1 Beta** (AMI Id: ami-0e1fece7)
Amazon Linux AMI Base 2011.02.1, EBS boot, 64-bit architecture with Amazon EC2 AMI Tools.
Root Device Size: 8 GB

**Red Hat Enterprise Linux 6.1 32 bit** (AMI Id: ami-0cbb4265)
Red Hat Enterprise Linux version 6.1, EBS-boot, 32-bit architecture.
Root Device Size: 7 GB

**Red Hat Enterprise Linux 6.1 64 bit** (AMI Id: ami-5e837b37)
Red Hat Enterprise Linux version 6.1, EBS-boot, 64-bit architecture.
Root Device Size: 6 GB


**SUSE Linux Enterprise Server 11 64-bit** (AMI Id: ami-e4a3578d)
SUSE Linux Enterprise Server 11 Service Pack 1 basic install, EBS boot, 64-bit architecture with Amazon EC2 AMI Tools preinstalled; Apache 2.2, MySQL 5.0, PHP 5.3, Ruby 1.8.7, and Rails 2.3.
Root Device Size: 15 GB

 Free tier eligible if used with a micro instance. See [AWS free tier](#) for complete details and terms.

Continue

[Submit feedback](#)

Make sure the 'My AMIs' tab is selected.



Request Instances Wizard Cancel

CHOOSE AN AMI | INSTANCE DETAILS | CREATE KEY PAIR | CONFIGURE FIREWALL | REVIEW

Choose an Amazon Machine Image (AMI) from one of the tabbed lists below by clicking its **Select** button.

Quick Start | **My AMIs** | Community AMIs | AWS Marketplace

Viewing: Owned By Me Search 1 to 5 of 5 Items

AMI ID	Root Device	Name	Platform	
 ami-04fe226d	ebs	971102534555/Neuchatel_CloudBioLinux_v3	 Other Linux	Select 
 ami-0cfa4465	ebs	971102534555/Stacks_CloudBioLinux_v5	 Other Linux	Select 
 ami-9668c8ff	ebs	971102534555/Stacks_CloudBioLinux_v4	 Other Linux	Select 
 ami-c9c34ca0	ebs	971102534555/2013 Short Read Genomics	 Other Linux	Select 
 ami-cc2ef3a5	ebs	971102534555/Neuchatel_Windows_JoinMap	 Windows	Select 

The next step in this process is to select an AMI. The AMI is a “snapshot” of the VM we prepared for you for this class. Select the **2013 Short Read Genomics Workshop**. **Please do not use any other AMI for this workshop!**

Click “Select”.

The next page is “Instance Details” where your first choice is the number and type of instances. For this class your default choice will be micro instance. However, we will need to select a ‘Large instance’ for this workshop. Depending on how many CPUs and how much memory you need you can choose a different instance if you were running this outside the class. See <http://aws.amazon.com/ec2/instance-types/> for information about the instance types and <http://aws.amazon.com/ec2/pricing/> for information about pricing per hour.

Tip:

The “Availability Zone” option can be any of the options, but you should “stick to it” for the class. If you create any S3 volumes (see the additional information in the supplementary tutorial) they are specific to a given availability zone so you'll need to have your instances and EBS volumes in the same zone.

Your next choice is “Launch Instances” or “Request Spot Instances”. The latter lets you bid on how much you want to spend on your EC2 node. This is a fantastic innovation, it works really well if you have tasks that need to get done but not urgently so you can afford to wait and, therefore, will only run your jobs during “cheap” off-peak hours. But, for this workshop we'll use the regular

“Launch Instances”. Use a Large instance and pick either us-east-1a, us-east-1b or us-east-1c as the Availability Zone.

Request Instances Wizard Cancel

CHOOSE AN AMI **INSTANCE DETAILS** CREATE KEY PAIR CONFIGURE FIREWALL REVIEW

Provide the details for your instance(s). You may also decide whether you want to launch your instances as "on-demand" or "spot" instances.

Number of Instances: Instance Type:

Launch Instances

EC2 Instances let you pay for compute capacity by the hour with no long term commitments. This transforms what are commonly large fixed costs into much smaller variable costs.

Launch into:

- EC2

Availability Zone:

Request Spot Instances

[Back](#) Continue

Click “Continue”. You can use the defaults on the next page, click “Continue” again.

On the next page you'll be given an opportunity to name your VM. Make sure you give this VM a name that includes your username so you can later find it in the list of instances. For example “StudentKonrad” since the username is “StudentKonrad”. Make sure you type your name into the 'Value' column and not the 'Key' column.

Request Instances Wizard Cancel

CHOOSE AN AMI **INSTANCE DETAILS** CREATE KEY PAIR CONFIGURE FIREWALL REVIEW

Add tags to your instance to simplify the administration of your EC2 infrastructure. A form of metadata, tags consist of a case-sensitive key/value pair, are stored in the cloud and are private to your account. You can create user-friendly names that help you organize, search, and browse your resources. For example, you could define a tag with key = Name and value = Webserver. You can add up to 10 unique keys to each instance along with an optional value for each key. For more information, go to [Using Tags](#) in the *EC2 User Guide*.

Key (127 characters maximum)	Value (255 characters maximum)	Remove
Name	StudentKonrad	✖
		✖

[Add another Tag](#). (Maximum of 10)

Continue

Request Instances Wizard

Cancel

✓

✓

CHOOSE AN AMI

INSTANCE DETAILS

CREATE KEY PAIR

CONFIGURE FIREWALL

REVIEW

Public/private key pairs allow you to securely connect to your instance after it launches. To create a key pair, enter a name and click **Create & Download your Key Pair**. You will then be prompted to save the private key to your computer. Note, you only need to generate a key pair once - not each time you want to deploy an Amazon EC2 instance.

☐ Choose from your existing Key Pairs

☐ Create a new Key Pair

☒ Proceed without a Key Pair

I do not want a keypair installed on this instance.

NOTE:You will not be able to connect to this instance unless you already know the password built in to this AMI.

< Back

Continue >

The next step lets you select the “keypair” used to let you log into this machine. This keypair is a file you download and supply to your SSH client in order to connect to the running server. By default most servers on Amazon EC2 don’t use passwords, but ours does.

Select “Proceed without a Key Pair”. Note, that for this workshop we will be using the NX client which is pre-installed on the image you are using so you will not need to use the key. However, if you ever use a different Amazon AMI from other sources you may need to create a new key pair for yourself.

Request Instances Wizard

Cancel

CHOOSE AN AMI

INSTANCE DETAILS

CREATE KEY PAIR

CONFIGURE FIREWALL

REVIEW

Security groups determine whether a network port is open or blocked on your instances. You may use an existing security group, or we can help you create a new security group to allow access to your instances using the suggested ports below. Add additional ports now or update your security group anytime using the Security Groups page.

☒ Choose one or more of your existing Security Groups

sg-1137e179 - default

sg-3935e351 - quick-start-1

(Selected groups: sg-1137e179)

☐ Create a new Security Group

< Back

Continue >

In the next step “Configure Firewall” select “Choose one or more of your existing security groups”.


Select 'Default'

Select “Continue”.

This will open the following ports:


- 22 (ssh)
- 80 (web server)

The next step is just a review. If it looks OK hit “Launch”. At this point you can watch your instance boot under the “Instances” section in the “EC2” tab.



At this point you wait just a couple minutes for the AMI instance to come online. Above you can see the instance is running, give it a couple minutes to finish its boot cycle. It's booting somewhere on a virtualized cluster node in Virginia! If you click on “Instances” on the left-hand panel you'll get a list of your running VMs (you'll see yours listed along with the other students, did you give it a name you can recognize?).

Once it turns green, you should click on the AMI and copy the Public DNS address to the clipboard (shown below is an example where the Public DNS is `ec2-184-72-71-5.compute-1.amazonaws.com`). We will need it in a moment.



Step 3– Log into the Running VM's Remote Desktop with NX

If you need a windowing environment rather than just an SSH terminal this is possible with CloudBioLinux! This is a great option if you want to use a GUI application (Graphical User Interface like Apollo or Artemis). It's very cool to see a remote desktop with Firefox and every other GUI application rendered quick and snappy over the Internet!

Here are the steps to get remote NX login working, take a look at the document <http://bit.ly/fA3ykQ> for more information. *Note that these instructions will only work for this particular VM AMI. Many AMIs will not have the NX server installed and therefore you will not be able to connect using the NX-client.*

1 - Download an NX Client

See <http://www.nomachine.com/download.php> for the various client installers. You should download the NX client specific to your machine (i.e. Windows, Linux or Mac). Make sure you download the NX Client and not the NX-server. Note that Mac OSX is incompatible with NX client 3.5. You should download NX Player preview version 4 if you have Mac OSX. See additional instructions if you have a Mac.




2 - Launch the Client and Connect

Once it is installed, start the NX Connection Wizard on your laptop. On Windows the start menu the icon will look like:




At the next screen you should name the connection Genomics Workshop and set the host-name to be the public DNS address of your Amazon AMI.



Note, if you ever stop or restart your AMI, the Public DNS address will change. Make sure you update your NX client with the new DNS!

Click 'Next >' and at the next screen ensure you change 'KDE' to 'GNOME'



Click 'Next >' and 'Next >' again on the next screen. Double click on the NX shortcut which will be on your Desktop.


Finally enter the login username “ubuntu” and password “ubuntu” (without the quotation marks).



Once connected you should get a nice desktop with links to the Terminal and to the Firefox web-browser. If a window pops up asking if you would like to run any updates, just close it.

One final task is to change your password from ubuntu to something less easy to guess. Double-click on the 'Terminal' icon and type the command 'passwd' (without the quote marks). Enter 'ubuntu' as the current password (again without the quotes) and then type in a new password of your choice. Note that for security, no characters will appear as you type. Please make it difficult for someone to guess your new password! You will then need to re-enter it to confirm.

Every time you log in via NX after this you will need to enter 'ubuntu' as the username and the password you have just entered. Please do not forget it!



To close the connection just close the window, you can choose “Disconnect” if you want to log back in later and pick up where you left off. However, if you stop or restart the AMI via the console, any unsaved information will be lost. **If you do stop or restart, remember to check and change the Public DNS address using the AWS console and change the DNS address in NX.**

At this point you've logged into the machine.

Step 4 – Stopping and starting the Instance

When you're done with the workshop it's very important to turn off the VM to avoid unnecessary charges. Just log back into the console, right click the instance and select “Terminate”. It will ask you to confirm. You can then watch the status change from “shutting down” to “terminated”. **Note – this will destroy all work done to date.** If you wish to keep your data, use the 'Stop' option instead. In this state you will not be charged for computing time, but will still be charged for storage.

Very important!

If you stop and then start you instance your Public DNS address may change. If this happens you will need to use the new DNS address with the NX-client or SSH.

Although much of what we have just done may not make much sense yet, most of you will feel totally comfortable and confident working on an EC2 node running Linux within a few hours. It's really amazing how quickly the fact that this is a remote computer will fade away. It may be hundreds of miles away but it will act just like a local computer, especially if you connect via NX.

Move on to the genomics tutorial!

2013 Workshop on Genomics, Cesky Krumlov

Part 3: Short read genomics: Remapping

Instructors:

- Konrad Paszkiewicz k.h.paszkiewicz at exeter ac uk

Objectives:

By the end of the workshop you will be expected to:

- Interpret FASTQ quality metrics
- Remove poor quality data
- Trim adaptor/contaminant sequences from FASTQ data
- Count the number of reads before and after trimming and quality control
- Align reads to a reference sequence to form a SAM file (Sequence AlignMent file) using BWA
- Convert the SAM file to BAM format (Binary AlignMent format)
- Identify and select high quality SNPs and Indels using SAMtools
- Identify missing or truncated genes with respect to the reference genome
- Identify SNPs which overlap with known coding regions

3.1 Introduction

In this section of the workshop we will be analysing a (hypothetical) strain of *E.coli* which is involved in urinary tract infections and causes internal hemorrhaging. The strain we start with is treatable with standard antibiotics. We want to obtain a list of single nucleotide polymorphisms (SNPs), insertions/deletions (Indels) and any genes which have been deleted.

In later sections you will repeat this analysis on your own using two other strains which have developed resistance. The final section will ask you to compare the results from all three datasets.

3.2 Quality control

In this section of the workshop we will be learning about evaluating the quality of an Illumina sequencing run. The process described here can be used with any FASTQ formatted file from any platform (e.g 454, Illumina, Ion Torrent, PacBio etc).

2nd (and 3rd) generation sequencers produce vast quantities of data. A single Illumina HiSeq lane will produce over 10Gb of data. However, the error rates of these platforms are 10-100x higher than Sanger sequencing. They also have very different error profiles. Unlike Sanger sequencing, where the most reliable sequences tend to be in the middle, NGS platforms tend to be most reliable near the beginning of each read.

Quality control usually involves:

- Calculating the number of reads before quality control
- Calculating GC content, identifying over-represented sequences
- Remove or trim reads containing adaptor sequences
- Remove or trim reads containing low quality bases
- Calculating the number of reads after quality control
- Rechecking GC content, identifying over-represented sequences

Quality control is necessary because:

- CPU time required for alignment and assembly is reduced
- Data storage requirements are reduced
- Reduce potential for bias in variant calling and/or de-novo assembly

Quality scores:

Most quality scores are calculated using the Phred scale (*Ewing B, Green P: Basecalling of automated sequencer traces using phred. II. Error probabilities. Genome Research 8:186-194 (1998)*). Each base call has an associated base call quality which estimates chance that the base call is incorrect.

Q10 = 1 in 10 chance of incorrect base call

Q20 = 1 in 100 chance of incorrect base call

Q30 = 1 in 1000 chance of incorrect base call

Q40 = 1 in 10,000 chance of incorrect base call

For most 454, SOLiD and Illumina runs you should see quality scores between Q20 and Q40. Note that these are only estimates of base-quality based on calibration runs performed by the

manufacturer against a sample of known sequence with (typically) a GC content of 50%. Extreme GC biases and/or particular motifs or homopolymers can cause the quality scores to become unreliable. Accurate base qualities are an essential part in ensuring variant calls are correct. As a rough and ready rule we generally assume that with Illumina data anything less than Q20 is not useful data and should be excluded.

FASTQ format:

A FASTQ entry consists of 4 lines

```
@D3P26HQ1:110:d0eh1acxx:8:1101:1116:2122 1:N:0:
AGGTGTCTCCTACAACCAAGCTACAACAGAGCAATGGGCTATCTGGTGGGATTAAAGGGGTGAAATGCATCCCCCTTAAATNAAAGTGTTTT
+
ADDADCFFHHHDHGHIII<GIICH4FGCIHIEGFHGHGIIIGDHFDFG?DEHH>FGIG=E@GGADDDCCCC@A>ABB>BBC:A>A#,228(4>:??B
```

1. A header line beginning with '@' containing information about the name of the sequencer, and the position at which the originating cluster was located and whether it passed purity filters.
2. The DNA sequence of the read
3. A header line or line beginning with just '+'
4. Quality scores for each base encoded in ASCII format

Typical FASTQ formatted file:

To reduce storage requirements, the FASTQ quality scores are stored as single characters and converted to numbers by obtaining the ASCII quality score and subtracting either 33 or 64. For example, the above FASTQ file is Sanger formatted and the character '@' has an ASCII value of 33. Therefore the corresponding base would have a Phred quality score of $33-33=Q0$ (i.e. totally unreliable). On the other hand a base with a quality score denoted by '@' which has an ASCII value of 64 would have a Phred quality score of $64-33=Q31$ (i.e. less than 1/1000 chance of being incorrect).

Just to confuse matters, there are several different methods of encoding quality scores in the ASCII format.

	Range	Offset	Type	Range
Sanger standard				
fastq-sanger	33-126	33	PHRED	0 to 93
Solexa/early Illumina				
fastq-solexa	59-126	64	Solexa	-5 to 62
Illumina 1.3+				
fastq-illumina	64-126	64	PHRED	0 to 62

Note that the latest Illumina CASAVA 1.8 pipeline (released June 2011), outputs in fastq-sanger rather than Illumina 1.3+. Thus Illumina 1.3+ and other Illumina scoring metrics are unlikely to be encountered if you are using Illumina sequencing data generated after July 2011.

3.2.1 Quality control – evaluating the quality of Illumina data

The first task when one receives sequencing data is to evaluate its quality and determine whether all the cash you have handed over was well-spent! To do this we will use the FastQC toolkit (<http://www.bioinformatics.bbsrc.ac.uk/projects/fastqc/>). FastQC offers a graphical visualisation of QC metrics, but *does not* have the ability to filter data.


Task 1:

From your home directory change into the `genomics_tutorial/strain1` directory. E.g.:

```
cd genomics_tutorial/strain1
```

Note that the FASTQ files are contained within a sub-directory called `illumina_reads`. Change into this directory by typing:

```
cd illumina_reads
```



```
ubuntu@ip-10-28-75-38: ~/genomics_tutorial/strain1/illumina_reads
File Edit View Search Terminal Help
ubuntu@ip-10-28-75-38:~$ ls
Desktop Documents genomics_tutorial perl_tutorial_materials
ubuntu@ip-10-28-75-38:~$ cd genomics_tutorial/
ubuntu@ip-10-28-75-38:~/genomics_tutorial$ ls
reference_sequence strain1 strain2 strain3
ubuntu@ip-10-28-75-38:~/genomics_tutorial$ cd strain1/
ubuntu@ip-10-28-75-38:~/genomics_tutorial/strain1$ ls
illumina_reads
ubuntu@ip-10-28-75-38:~/genomics_tutorial/strain1$ cd illumina_reads/
ubuntu@ip-10-28-75-38:~/genomics_tutorial/strain1/illumina_reads$ ls
strain1_read1.fastq strain1_read2.fastq
ubuntu@ip-10-28-75-38:~/genomics_tutorial/strain1/illumina_reads$
```

Note that this is a paired-end run. As such there are two files, one for read 1 (`strain1_read1.fastq`) and the other for the reverse read 2 (`strain1_read2.fastq`). Reads from the same pair can be identified because they have the same header. Many programs require that the read 1 and read 2 files have the reads in the same order. To view the first few reads we can use the 'more' command:

```
@CP000243.1-6332150-1
ctgccgcaggcgctgcgcaattggaagatcataat
+
CCCC#CCIB2CCCC@=BB#CCCAC7@C$B:C2CAIC=CCC
@CP000243.1-6332148-1
cctgccaaatttgcgcccatatttaagtatcccttaacg
+
IBAI CC (9ABBCBB#CCBC8?BC##)@#CCCB6#BCDBC9
@CP000243.1-6332146-1
attgagaatggcgcaagtaacaatatccagacccgacatg
+
CCCCICBBCC#C7C=C#C3CCBBC;CCICC9=BCCAA#C<
@CP000243.1-6332144-1
ccggcgatatcccaggcaaaatcctcccagctccagccgc
+
C=CCCCICCCBCAB;CCCAC=CCC#CCCCAC>;<#%7C
```

```
more strain1_read1.fastq
```

Press 'q' to exit more. We can do the same for read 2 and confirm that the headers are the same.

```
more strain1_read2.fastq
```


```
@CP000243.1-6332150-2
accgtagaaggcgacaaggccttctatcgaagtcccggctc
+
F>=B4AFBIF>E?B@AH:B???A8F@#*5DCI&#B4??C?
@CP000243.1-6332148-2
cctacaatctgtcaacagaatgtgaaaacgtcaatacagc
+
H>E?CFI@8BB?'@?+AC/-BHBCH5@FD>#BBF(=B&EF
@CP000243.1-6332146-2
gcgccaaagtggcagaaaaactgaatcaggtttgcgctaa
+
BGHBA'AAHB=B#E=9BH3;C>@7+5B@6@EC>B>FBH<A
@CP000243.1-6332144-2
tccgcgagcataccaccatcgtaagcactatcgcgatct
+
HG8B5BEI?A=8CHDBHB>C:H2I4>88??H<8/H>EB=4
```

Again, press 'q' to exit more.

Now, let's start the fastqc program. Type:

fastqc

This should open the fastqc program. Load the strain1_read1.fastq file in the genomics_tutorial/strain1/illumina_reads directory.




After a few minutes the program should finish analysing the FASTQ file.

The fastqc program performs a number of tests which determines whether or a green tick (pass), exclamation mark (warning) or red cross (fail) is displayed. However it is important to realise that fastqc has no knowledge of what your library is or should look like. All of its tests are based on a

completely random library with 50% GC content. Therefore if you have a sample which does not match these assumptions, it may and 'fail' the library. For example, if you have a high AT or high GC organism it may fail the per sequence GC content. If you have any barcodes or low complexity libraries (e.g. small RNA libraries) they may also fail some of the sequence complexity tests.


The bottom line is that you need to be aware of what your library is and whether what fastqc is reporting makes sense for that type of library.



In this case we have an *E.coli* library so a percentage GC content of 50% is consistent with our expectations.


Quality scores:

This is one of the most important metrics. If the quality scores are poor, either the wrong FASTQ encoding has been guessed by fastqc (see the title of the chart), or the data itself is poor quality. This view shows an overview of the range of quality values across all bases at each position in the FastQ file. Generally anything with a median quality score greater than Q20 is regarded as acceptable, anything above Q30 is regarded as 'good'. For more details, see the help documentation in fastqc.



Per base sequence content:


For a completely randomly generated library with a GC content of 50% one expects that at any given position within a read there will be a 25% chance of finding an A,C,T or G base. Here we can see that our library satisfies these criteria, although there appears to be some minor bias at the beginning of the read. This may be due to PCR duplicates during amplification or during library preparation. It is unlikely that one will ever see a perfectly uniform distribution. See <http://biosciences.exeter.ac.uk/facilities/sequencing/dataguide/qualitycontrol/> for examples of good vs bad runs as well as the fastqc help for more details.



Sequence duplication levels:

In a library that covers a whole genome uniformly most sequences will occur only once in the final set. A low level of duplication may indicate a very high level of coverage of the target sequence, but a high level of duplication is more likely to indicate some kind of enrichment bias (e.g. PCR over-amplification).

This module counts the degree of duplication for every sequence in the set and creates a plot showing the relative number of sequences with different degrees of duplication.



Task 2:

Do the same for read 2 as we have for read 1. Open fastqc and the read 2 file. Look at the various plots and metrics which are generated. How similar are they?

Note that the number of reads reported in both files is identical. This is because if one read fails to pass the Illumina chastity filter, its partner is automatically excluded too.

Overall, both read 1 and read 2 can be regarded as 'good' data-sets.

3.2.2 Quality control – filtering of Illumina data

In this section we will be actually be filtering the data to ensure any low quality reads are removed and that any sequences containing adaptor sequences are either trimmed or removed altogether. To do this we will use the fastq-mcf program from the ea-utils package (available at <http://code.google.com/p/ea-utils/>). This package is remarkably fast and ensures that after filtering both read 1 and read 2 files are in the correct order.

Note: Typically when submitting raw Illumina data to NCBI or EBI you would submit unfiltered data, so don't delete your original fastq files!

Make sure you are in the ~/genomics_tutorial/strain1/illumina_reads directory. We will execute the fastq-mcf program which performs both adaptor sequence trimming and low quality bases. To remove adaptor sequences, we need to supply the adaptor sequences to the program. A list of the most common adaptors used is given in the file /usr/local/ea-utils/adaptors.fasta:

```
ubuntu@ip-10-28-27-206:~/genomics_tutorial/strain1/illumina_reads$ more /usr/local/ea-utils/adaptors.fasta
>IlluminaPEAdapters1
GATCGGAAGAGCGGTTTCAGCAGGAATGCCGAGACACTCTTTCCCTACACGACGCTCTTCCGATCT
>UniversalAdaptor
AATGATACGGCGACCACCGAGATCTACACTCTTTCCCTACACGACGCTCTTCCGATCT
>Genericadaptor
TGTATTATACAAGAT
```

Task 3:

To run the fastq-mcf program, type the following (all on one line):

```
fastq-mcf /usr/local/ea-utils/adaptors.fasta strain1_read1.fastq  
strain1_read2.fastq -o strain1_read1.filtered.fastq -o  
strain1_read2.filtered.fastq -C 1000000 -q 20 -p 10 -u -x 0.01
```

After a few minutes the filtering should be complete and you should see something similar to:

```
Scale used: 2.2
Filtering Illumina reads on purity field
Phred: 33
Threshold used: 2501 out of 1000000
Files: 2
Total reads: 3247450
Too short after clip: 0
Trimmed 708987 reads (strain1_read1.fastq) by an average of 1.22 bases on quality < 20
Trimmed 1037979 reads (strain1_read2.fastq) by an average of 1.36 bases on quality < 20
```

Note that read 1 has few reads trimmed than read 2. This is generally true as read 2 tends to be of slightly lower quality than read 1.

The filtered reads should be present in files strain1_read1.filtered.fastq and strain1_read2.filtered.fastq

Task 4:

Check the quality scores and sequence distribution in the fastqc program for the two filtered fastq files. You should notice very little change (since comparatively few reads were filtered).

Note that although in this case filtering did very little, most runs are not as 'clean' as this so you may find your filtering removes many more reads in the real-world.

Task 5:

We can perform a quick check (although this by no means guarantees) that the sequences in read 1 and read 2 are in the same order by checking the ends of the two files and making sure that the headers are the same.

```
ubuntu@ip-10-28-24-212:~/genomics_tutorial/strain1/illumina_reads$ tail strain1_read1.filtered.fastq
+
3BABBBBBBABABABABABBAABBBBBBBBABBGGHHHG
@BENM_2011_0526_022:5:9999:500:40/1
TCCGGCTTTATCCCCGATTCTGTGGATAACCGTATTACC
+
)EDDDEDEDDDEDDDEDDDDDEEDDEDEDDDEBBAAA
@BENM_2011_0526_022:6:10000:423:40/1
5ACGCCATTCTGCGTCAGAGCAGACTCAATCTTGCAATA
+
)CCDDCCDCDDCCDDDDCCDDCCDDCCCCDDCCDCHHHHH
ubuntu@ip-10-28-24-212:~/genomics_tutorial/strain1/illumina_reads$ tail strain1_read2.filtered.fastq
+
=GGGGFGGGFFFFGGFFFGFFGGFFGGF; <<<<
@BENM_2011_0526_022:5:9999:500:40/2
TGTCTGCAAGTTCGAATTACCAACAAAGCACTGCTGTTTT
+
[IJIIJIIJIIJIIIIJIIJIIIIJIIJIIIIJII87888
@BENM_2011_0526_022:6:10000:423:40/2
TGCGCCATTATCCGGTAACTATCGTCTTGAGTCCAACCCGG
+
CCDCCDDCCDDCCDDCCDDCCDDCCDDCCDDCCDBBBAB
```

Task 6:

Check the number of reads in each filtered file. They should be the same. To do this use the `grep` command to search for the number of times the header appears. E.g:

```
grep -c "@CP000243" strain1_read1.filtered.fastq
```

Do the same for the `strain1_read2.filtered.fastq` file.

3.3 – Aligning Illumina data to a reference sequence

Now that we have checked the quality of our raw data, we can begin to align the reads against a reference sequence. In this way we can compare how the reference sequence and the strain we have sequenced compare.

To do this we will be using a program called BWA (Burrows Wheeler Aligner *Li H. and Durbin R. (2009) Fast and accurate short read alignment with Burrows-Wheeler Transform. Bioinformatics, 25:1754-60.*). This uses an algorithm called (unsurprisingly) Burrows Wheeler to rapidly map reads to the reference genome. BWA also allows for a certain number of mismatches to account for variants which may be present in strain 1 vs the reference genome. Unlike other alignment packages such as Bowtie (version 1) BWA allows for insertions or deletions as well.

By mapping reads against a reference, what we mean is that we want to go from a FASTQ file listing lots of reads, to another type of file (which we'll describe later) which lists the reads AND where/if it maps against the reference genome. The figure below illustrates what we are trying to achieve here. Along the top in grey is the reference sequence. The coloured sequences below that indicates individual sequences and how they map to the reference. If there is a real variant in a bacterial genome we would expect that (nearly) all the reads would contain the variant at the relevant position rather than the same base as the reference genome. Remember that error rates for any single read on second generation platforms tend to be around 0.5-1%. Therefore a 100bp read is likely to contain at least one error.


Let's look at 2 potential SNPs which are in fact artefacts.

1. Sequencing error:

The region highlighted in green on the right shows that most reads agree with the reference sequence (i.e. C-base). However, 2 reads near the bottom show an A-base. In this situation we can safely assume that the A-bases are due to a sequencing error rather than a genuine variant.

2. PCR duplication:

The highlighted region red on the left shows where there appears to be a variant (either due to sequencing of a diploid genome or non-clonal samples). A C-base is present in the reference and half the reads, whilst an A-base is present in the other reads.



Is this a genuine difference or a sequencing or sample prep error? What do you think? If this was a real sample, would you expect all the reads containing an A to start at the same location?

The answer is no. This 'SNP' is in fact due to PCR duplication. I.e. the same fragment of DNA has been replicated many times more than the average and happens to contain an error. We can filter out such reads during after alignment to the reference (see later).


3.3.1: Indexing a reference genome:

Before we can start aligning reads to a reference genome, the genome sequence needs to be indexed. This means sorting the genome into easily searched chunks.

Task 7: Generating an index file from the reference sequence

Change directory to the reference_sequence directory:

```
cd ~/genomics_tutorial/reference_sequence
```

A terminal window titled 'ubuntu@ip-10-28-24-212: ~/genomics_tutorial/reference_sequence'. The window has a menu bar with 'File', 'Edit', 'View', 'Search', 'Terminal', and 'Help'. The terminal text shows the user running 'ls' in the directory '~/genomics_tutorial/reference_sequences', which lists two files: 'Ecoli_UTI89.fna' and 'Ecoli_UTI89.gff'. The prompt returns to 'ubuntu@ip-10-28-24-212:~/genomics_tutorial/reference_sequences\$'.

In this directory we have 2 files. Ecoli_UTI89.fna is a FASTA file which contains the reference genome sequence. The Ecoli_UTI89.gff file contains the annotation for this genome. We will use this later.

First, let's look at the bwa command itself. Type:

```
bwa
```

This should yield something like:

Program: bwa (alignment via Burrows-Wheeler transformation)
Version: 0.5.7 (r1310)
Contact: Heng Li <lh3@sanger.ac.uk>

Usage: bwa <command> [options]

Command:

index	index sequences in the FASTA format
aln	gapped/ungapped alignment
samse	generate alignment (single ended)
sampe	generate alignment (paired ended)
bwasw	BWA-SW for long queries
fa2pac	convert FASTA to PAC format
pac2bwt	generate BWT from PAC
pac2bwtgen	alternative algorithm for generating BWT
bwtupdate	update .bwt to the new format
pac_rev	generate reverse PAC
bwt2sa	generate SA from BWT and Occ
pac2cspac	convert PAC to color-space PAC
stdsw	standard SW/NW alignment

BWA is actually a suite of programs which all perform different functions. We are only going to use the three during this workshop, bwa index, bwa aln and bwa sampe.

If we type:

bwa index

We can see more options for the bwa index command:

Usage: bwa index [-a bwtsv|div|is] [-c] <in.fasta>

Options: -a STR BWT construction algorithm: bwtsv or is [is]

-p STR prefix of the index [same as fasta name]

-c build color-space index

Warning: '-a bwtsv' does not work for short genomes, while '-a is' and '-a div' do not work not for long genomes. Please choose '-a' according to the length of the genome.

By default bwa index will use the IS algorithm to produce the index. This works well for most genomes, but for very large ones (e.g. vertebrate) you may need to use bwtsv. For bacterial genomes the default algorithm will work fine.

Now we will create a reference index for the genome using BWA:

bwa index Ecoli_UTI89.fna

```
ubuntu@ip-10-28-24-212: ~/genomics_tutorial/reference_sequence
File Edit View Search Terminal Help
ubuntu@ip-10-28-24-212:~/genomics_tutorial/reference_sequence$ bwa index Ecoli_UTI89.fna
[bwa_index] Pack FASTA... 0.08 sec
[bwa_index] Reverse the packed sequence... 0.02 sec
[bwa_index] Construct BWT for the packed sequence...
[bwa_index] 1.16 seconds elapse.
[bwa_index] Construct BWT for the reverse packed sequence...
[bwa_index] 1.12 seconds elapse.
[bwa_index] Update BWT... 0.02 sec
[bwa_index] Update reverse BWT... 0.03 sec
[bwa_index] Construct SA from BWT and Occ... 0.36 sec
[bwa_index] Construct SA from reverse BWT and Occ... 0.36 sec
ubuntu@ip-10-28-24-212:~/genomics_tutorial/reference_sequence$ ls
Ecoli_UTI89.fna      Ecoli_UTI89.fna.ann  Ecoli_UTI89.fna.pac  Ecoli_UTI89.fna.rpac  Ecoli_UTI89.fna.sa
Ecoli_UTI89.fna.amb  Ecoli_UTI89.fna.bwt  Ecoli_UTI89.fna.rbwt  Ecoli_UTI89.fna.rsa   Ecoli_UTI89.gff
ubuntu@ip-10-28-24-212:~/genomics_tutorial/reference_sequence$
```

If you now list the directory contents using the 'ls' command, you will notice that the BWA index program has created a set of new files. These are the index files BWA needs.

3.3.2: Aligning reads to the indexed reference sequence:

Now we can begin to align read 1 and read 2 to the reference genome. First of all change back into the ~/genomics_tutorial/strain1/ directory.


cd ~/genomics_tutorial/strain1/

Now let's create a directory to store our alignment (just to be tidy):

```
mkdir remapping_to_reference
```

And now let's change into that directory:

```
cd remapping_to_reference
```

A terminal window screenshot with a blue title bar containing the text 'ubuntu@ip-10-28-24-212: ~/genomics_tutorial/strain1/remapping_to_ref'. Below the title bar is a menu bar with 'File', 'Edit', 'View', 'Search', 'Terminal', and 'Help'. The terminal shows the following commands and output:
ubuntu@ip-10-28-24-212:~/genomics_tutorial/strain1\$ ls
illumina_reads
ubuntu@ip-10-28-24-212:~/genomics_tutorial/strain1\$ mkdir remapping_to_reference
ubuntu@ip-10-28-24-212:~/genomics_tutorial/strain1\$ ls
illumina_reads remapping_to_reference
ubuntu@ip-10-28-24-212:~/genomics_tutorial/strain1\$ cd remapping_to_reference/
ubuntu@ip-10-28-24-212:~/genomics_tutorial/strain1/remapping_to_reference\$ █

We can now explore the alignment options BWA has to offer. Type:

```
bwa aln
```

Usage: bwa aln [options] <prefix> <in.fq>

Options:

- n NUM max #diff (int) or missing prob under 0.02 err rate (float) [0.04]
- o INT maximum number or fraction of gap opens [1]
- e INT maximum number of gap extensions, -1 for disabling long gaps [-1]
- i INT do not put an indel within INT bp towards the ends [5]
- d INT maximum occurrences for extending a long deletion [10]
- l INT seed length [32]
- k INT maximum differences in the seed [2]
- m INT maximum entries in the queue [2000000]
- t INT number of threads [1]
- M INT mismatch penalty [3]
- O INT gap open penalty [11]
- E INT gap extension penalty [4]
- R INT stop searching when there are >INT equally best hits [30]
- q INT quality threshold for read trimming down to 35bp [0]
- c input sequences are in the color space
- L log-scaled gap penalty for long deletions
- N non-iterative mode: search for all n-difference hits (sloow)
- f FILE file to write output to instead of stdout

We can see that we need to provide BWA with a FASTQ file containing the raw reads (denoted by <in.fq>) to align to a reference file (unhelpfully this is listed as <prefix>). There are also a number of options. The most important are the maximum number of differences in the seed (-k i.e. the first 32 bp of the sequence vs the reference), the number of processors the program should use (-t – our machine has 8 processors), and the file to output the results to (-f).

Note that our reference sequence is in ~/genomics_tutorial/reference_sequence/Ecoli_UTI89.fna

Our filtered reads for read 1 are in

~/genomics_tutorial/strain1/illumina_reads/strain1_read1.filtered.fastq

Our filtered reads for read 2 are in

~/genomics_tutorial/strain1/illumina_reads/strain1_read2.filtered.fastq

So to align read 1 using 8 processors and ignoring bases with quality scores less than 20 and output to file bwa_read1.aln, type (all on one line):

Task 8: Aligning reads to the reference genome


```
bwa aln -t 8 -q 20 -f bwa_read1.aln
~/genomics_tutorial/reference_sequence/Ecoli_UTI89.fna
~/genomics_tutorial/strain1/illumina_reads/strain1_read1.filtered.fastq
```

This will take around 5 minutes. Once completed we can do exactly the same for read 2.

```
bwa aln -t 8 -q 20 -f bwa_read2.aln
~/genomics_tutorial/reference_sequence/Ecoli_UTI89.fna
~/genomics_tutorial/strain1/illumina_reads/strain1_read2.filtered.fastq
```

```
ubuntu@ip-10-28-24-212:~/genomics_tutorial/strain1/remapping_to_reference$ bwa aln -t 8 -q 20 -f bwa_read1.aln
ln ~/genomics_tutorial/reference_sequence/Ecoli_UTI89.fna ~/genomics_tutorial/strain1/illumina_reads/strain1_read1.filtered.fastq
[bwa_aln] 17bp reads: max_diff = 2
[bwa_aln] 38bp reads: max_diff = 3
[bwa_aln] 64bp reads: max_diff = 4
[bwa_aln] 93bp reads: max_diff = 5
[bwa_aln] 124bp reads: max_diff = 6
[bwa_aln] 157bp reads: max_diff = 7
[bwa_aln] 190bp reads: max_diff = 8
[bwa_aln] 225bp reads: max_diff = 9
[bwa_read_seq] 0.0% bases are trimmed.
[bwa_aln_core] calculate SA coordinate... 14.99 sec
[bwa_aln_core] write to the disk... 0.03 sec
[bwa_aln_core] 262144 sequences have been processed.
[bwa_read_seq] 0.0% bases are trimmed.
[bwa_aln_core] calculate SA coordinate...
```

Once complete you can list the contents of the directory and you should see two new files `bwa_read1.aln` and `bwa_read2.aln`. These files store the alignment of read 1 and read 2 to the reference genome. They are in a binary format and only BWA will read these files (i.e. they are not human readable).



```
ubuntu@ip-10-28-24-212: ~/genomics_tutorial/strain1/remapping_to_reference
File Edit View Search Terminal Help
ubuntu@ip-10-28-24-212:~/genomics_tutorial/strain1/remapping_to_reference$ ls
bwa_read1.aln bwa_read2.aln
ubuntu@ip-10-28-24-212:~/genomics_tutorial/strain1/remapping_to_reference$
```

The task is now to marry the alignments of the two reads (i.e. to include the paired-end distance information) To do this we will use the `bwa sampe` command.

```
bwa sampe
```

This should yield the following description:

Usage: `bwa sampe [options] <prefix> <in1.sai> <in2.sai> <in1.fq> <in2.fq>`

Options:

- a INT maximum insert size [500]
- o INT maximum occurrences for one end [100000]
- n INT maximum hits to output for paired reads [3]
- N INT maximum hits to output for discordant pairs [10]
- c FLOAT prior of chimeric rate [1.0e-05]
- P preload index into memory (for base-space reads only)
- s disable Smith-Waterman for the unmapped mate
- f FILE sam file to output results to instead of stdout

Notes: 1. For SOLID reads, <in1.fq> corresponds R3 reads and <in2.fq> to F3.
2. For reads shorter than 30bp, applying a smaller -o is recommended to get a sensible speed at the cost of pairing accuracy.

We can see that we need to provide BWA with a reference genome in FASTA format file (unhelpfully this is listed as <prefix>), the 2 alignments of read 1 and read 2 (denoted by <in1.sai> and <in2.sai>) and the two files containing the filtered FASTQ reads (denoted by <in1.fq> and <in2.fq>). There are also a number of options. The most important are the maximum insert size (i.e. the distance between the paired-end reads) and the file to output the results to (-f).

Here we will set the maximum insert-size to be 600bp (determined from the library preparation), and the output file to be alignment.sam. The maximum insert-size enables BWA to flag potential insertions/deletions compared to the reference.

Note that our reference sequence is in `~/genomics_tutorial/reference_sequence/Ecoli_UTI89.fna`

Our filtered reads for read 1 are in

`~/genomics_tutorial/strain1/illumina_reads/strain1_read1.filtered.fastq`

Our filtered reads for read 2 are in

`~/genomics_tutorial/strain1/illumina_reads/strain1_read2.filtered.fastq`

Our bwa alignment files are in the same directory we are in (bwa_read1.aln, bwa_read2.aln). Enter the following on the command-line.

```
bwa sampe -P -a 600 -f alignment.sam
~/genomics_tutorial/reference_sequence/Ecoli_UTI89.fna bwa_read1.aln
bwa_read2.aln
~/genomics_tutorial/strain1/illumina_reads/strain1_read1.filtered.fastq
~/genomics_tutorial/strain1/illumina_reads/strain1_read2.filtered.fastq
```


This will take around 10 minutes.

Once completed the alignment will be stored in the alignment.sam file.

```
ubuntu@ip-10-28-24-212:~/genomics_tutorial/strain1/remapping_to_reference$ time bwa sampe -P -a 600 -f alignment.sam ~/genomics_tutorial/reference/sequence/Ecoli_UTI89.fna bwa_read1.aln bwa_read2.aln ~/genomics_tutorial/strain1/illumina_reads/strain1_read1.filtered.fastq ~/genomics_tutorial/strain1/illumina_reads/strain1_read2.filtered.fastq
[bwa_read_seq] 0.0% bases are trimmed.
[bwa_read_seq] 0.0% bases are trimmed.
[bwa_sai2sam_pe_core] convert to sequence coordinate...
[infer_isize] (25, 50, 75) percentile: (461, 470, 479)
[infer_isize] low and high boundaries: 425 and 515 for estimating avg and std
[infer_isize] inferred external isize from 127342 pairs: 469.552 +/- 14.008
[infer_isize] skewness: -0.415; kurtosis: 0.397; ap_prior: 4.26e-05
[infer_isize] inferred maximum insert size: 553 (5.99 sigma)
[bwa_sai2sam_pe_core] time elapses: 1.80 sec
[bwa_sai2sam_pe_core] changing coordinates of 741 alignments.
[bwa_sai2sam_pe_core] align unmapped mate...
```

3.4: Viewing the alignment

Once the alignment is complete, list the directory contents using the ls command, check that the alignment.sam file is present.



```
ubuntu@ip-10-28-24-212:~/genomics_tutorial/strain1/remapping_to_references$ ls
alignment.sam  bwa_read1.aln  bwa_read2.aln
ubuntu@ip-10-28-24-212:~/genomics_tutorial/strain1/remapping_to_references$
```

The raw alignment is stored in what is called SAM format (Simple AlignMent format). It is in plain text format and you can view it if you wish using the 'less' command. (e.g. less alignment.sam). Do not try to open the whole file in a text editor as you will likely run out of memory!



```
SRR292678.2 141 * 0 0 * * 0 0 GCAGGGGAGCGTGTGAGGCGATTGCTGCTGGTGAATTAAGGACGGTGGCAGAGACTGCCGTATCCGCTCAATAACTGTTATGTCAGGTAT HHHHHHHHHHHHHHHHHHHHHHHHHHHHHHH
PHGPHHHHHHHHEGCGGFFAFDFGGGG=.5;C@=?CEEG04C'>?AA449E:AA?>=,7> = 543972 -470 GAACTGGGCTATACCGTAGATAAAGAGAGCATGTTGTAGGCTGAATTGCATAGCTTCAGCAATTTACGATGATGCGGTAGCTGTTT EDCEEE?EDEEGEGEGEGFGE
SRR292678.3 83 CP000243.1 544352 50 90M = XT:A:U NM:i:4 SM:i:25 AM:i:25 X0:i:1 X1:i:0 XM:i:4 XG:i:0 MD:Z:38C2C35A5T6 GTCAATTAGGCTGGTGGCATATAGGATTAGGTGCTCTTAACGTGGGGGGCGGTACATCCATAACCGCATGCTCTCCGTCGCCGGGC HHHHHHHHHHHHHHHHHHHHHHH
7FGGGBGDFEGGGGGGGGGGGGGGGBFGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGGG 544352 470 GTCAATTAGGCTGGTGGCATATAGGATTAGGTGCTCTTAACGTGGGGGGCGGTACATCCATAACCGCATGCTCTCCGTCGCCGGGC HHHHHHHHHHHHHHHHHHHHHHH
SRR292678.3 163 CP000243.1 543972 50 90M = XT:A:U NM:i:4 SM:i:25 AM:i:25 X0:i:1 X1:i:0 XM:i:4 XG:i:0 MD:Z:405T35T2A40 GFEODGGGIEACECHFFHHCHCA=';@4FEFBEHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHH 4240754 476 AGGATGTTATTTTAAACCCGAGCAGCGTCAGGCGCTGTTTGAACGATTTCATGCCAGCGGAGCGCAATCGTCACCGTTGCGATGGGAT HHHHHHHHHHHHHHHHHHHHHHH
SRR292678.4 99 CP000243.1 4240368 60 90M = XT:A:U NM:i:1 SM:i:37 AM:i:37 X0:i:1 X1:i:0 XM:i:1 XG:i:0 MD:Z:8960 BGGGHGHFHHHGHFHHFHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHH 4240368 -476 AAAGCATTCAAATTTTAAATGCTTTATTTGCAATTTCTCTGAAATACGAAACATTCGCAACATCGATGTACCCATAACGATAACCGG EESBCCDEDFCBEBFHFEC=
DDEA7FHFHFF;EFCBA@CAE;F>AB000A2D:D?FCDAEB09HFEFHH<FHHHHHHFGGDFGGGI XT:A:U NM:i:1 SM:i:37 AM:i:37 X0:i:1 X1:i:0 XM:i:1 XG:i:0 MD:Z:86A3 SRR292678.5 99 CP000243.1 2340587 29 158M = 2340985 488 TCAGCTTTTGGCGACTACAACTCGTACCCTGTCGCGCAGCTCGACCGTAACAGCTACGCCCTGACAGAGCCGAGCGCAATGCCAG HGHGCHHHHEEHFHHHHHHH
GHFHHHHHHGGGGEHFGFFFFH;EBEFFHFFH5FEABEGAFBF2@:A?CCADGEFF>EEECA:8B?@ XT:A:M NM:i:7 SM:i:29 AM:i:29 X0:i:0 XG:i:0 MD:Z:1865T19T2T8T5G28C5 SRR292678.5 147 CP000243.1 2340985 29 90M = 2340587 -488 CGAGCTGATTAAACGATAGTGAAGCCCTGTTTGAATAAGGAAGAGCATGGACAAGGAATTACCGTGGCTGGCGGATAACGCCCACTGGA FF/EFAA>.:BA=GCEB/EEH
DEHHEDEBD7DC@=EEEEDEE.GIEEGGBFHHHFFEFF=EEE=EGDFGCCDD7/CBA//EE=DDDBB XT:A:U NM:i:2 SM:i:29 AM:i:29 X0:i:1 X1:i:0 XM:i:2 XG:i:0 MD:Z:6A10T72
```

Each alignment line has 11 mandatory fields for essential alignment information such as mapping position, and variable number of optional fields for flexible or aligner specific information. For further details as to what each field means see <http://samtools.sourceforge.net/SAM1.pdf>

Task 9: Convert SAM to BAM file

Before we can visualise the alignment however, we need to convert the SAM file to a BAM (Binary AlignMent format) which can be read by most software analysis packages. To do this we will use another suite of programs called samtools. Type:

```
samtools view
```

Usage: `samtools view [options] <in.bam>|<in.sam> [region1 [...]]`

Options:

- b output BAM
- h print header for the SAM output
- H print header only (no alignments)
- S input is SAM
- u uncompressed BAM output (force -b)
- l fast compression (force -b)
- x output FLAG in HEX (samtools-C specific)
- X output FLAG in string (samtools-C specific)
- c print only the count of matching records
- L FILE output alignments overlapping the input BED FILE [null]
- t FILE list of reference names and lengths (force -S) [null]
- T FILE reference sequence file (force -S) [null]
- o FILE output file name [stdout]
- R FILE list of read groups to be outputted [null]
- f INT required flag, 0 for unset [0]
- F INT filtering flag, 0 for unset [0]
- q INT minimum mapping quality [0]
- l STR only output reads in library STR [null]
- r STR only output reads in read group STR [null]
- s FLOAT fraction of templates to subsample; integer part as seed [-1]
- ? longer help

We can see that we need to provide `samtools view` with a reference genome in FASTA format file (-T), the -b and -S flags to say that the output should be in BAM format and the input in SAM, plus the alignment file.

Note that our reference sequence is in `~/genomics_tutorial/reference_sequence/Ecoli_UTI89.fna`

Type (all on one line):

```
samtools view -bS -T  
~/genomics_tutorial/reference_sequence/Ecoli_UTI89.fna  
alignment.sam > alignment.bam
```

This should take around 5 minutes.

Task 10: Sort BAM file

Once this is complete we then need to sort the BAM file so that the reads are stored in the order they appear along the chromosomes (don't ask me why this isn't done automatically....). We can do this using the `samtools sort` command.

```
samtools sort alignment.bam alignment.sorted
```

This will take another 5 minutes or so. Once complete, list the directory contents (again using the `ls`

command). Note that the output file is called alignment.sorted, but actually the program appends .bam to the end of the file (see below). Just to add confusion.

```
ubuntu@ip-10-28-24-212:~/genomics_tutorial/strain1/remapping_to_reference$ samtools sort alignment.bam alignment.sorted
[bam sort_core] merging from 6 files...
ubuntu@ip-10-28-24-212:~/genomics_tutorial/strain1/remapping_to_reference$ ls
alignment.bam alignment.sam alignment.sorted.bam bwa_read1.aln bwa_read2.aln
ubuntu@ip-10-28-24-212:~/genomics_tutorial/strain1/remapping_to_reference$
```


Task 11: Fill the 'MD' field using samtools

This mysterious task is an obscure term whose origins are lost in the mists of time. However, all it means is that we want samtools to look at the BAM file and annotate where it thinks SNPs should be.

On the command-line type:

```
samtools fillmd -b alignment.sorted.bam
~/genomics_tutorial/reference_sequence/Ecoli_UTI89.fna >
alignment.calmd.bam
```

Again, this should take around 5 minutes.



```
ubuntu@ip-10-28-24-212: ~/genomics_tutorial/strain1/remapping_to_reference
File Edit View Search Terminal Help
ubuntu@ip-10-28-24-212:~/genomics_tutorial/strain1/remapping_to_reference$ samtools fillmd -b alignment.
ubuntu@ip-10-28-24-212:~/genomics_tutorial/strain1/remapping_to_reference$ ls
alignment.bam alignment.calmd.bam alignment.sam alignment.sorted.bam bwa_read1.aln bwa_read2.aln
ubuntu@ip-10-28-24-212:~/genomics_tutorial/strain1/remapping_to_reference$
```

Once complete, list the directory contents with the 'ls' command and confirm that the alignment.calmd.bam file has been created.

Task 12: Remove suspected PCR duplicates

Especially when using paired-end reads, samtools can do a reasonably good job of removing potential PCR duplicates (see the first part of this workshop if you are unsure what this means).

Again, samtools has a great little command to do this:

On the command-line type:

```
samtools rmdup alignment.calmd.bam alignment.calmd.rmdup.bam
```

This will take around 3-4 minutes. Again, once complete, list the directory.

```

ubuntu@ip-10-28-24-212:~/genomics_tutorial/strain1/remapping_to_reference$ samtools rmdup alignment.calmd.bam alignment.calmd.rmdup.bam
[bam_rmdup_core] processing reference CP000243.1...
[bam_rmdup_core] 25 unmatched pairs
[bam_rmdup_core] processing reference CP000244.1...
[bam_rmdup_core] 91933 / 4264027 = 0.0216 in library ' '
ubuntu@ip-10-28-24-212:~/genomics_tutorial/strain1/remapping_to_reference$ ls
alignment.bam alignment.calmd.bam alignment.calmd.rmdup.bam alignment.sam alignment.sorted.bam bwa_read1.aln bwa_read2.aln
ubuntu@ip-10-28-24-212:~/genomics_tutorial/strain1/remapping_to_reference$

```

Task 12: Index the BAM file

Most programs used to view BAM formatted data require an index file to locate the reads mapping to a particular location quickly. We'll use the samtools index command to do this.

Type:

```
samtools index alignment.calmd.rmdup.bam
```

```

ubuntu@ip-10-28-24-212:~/genomics_tutorial/strain1/remapping_to_reference$ samtools index alignment.calmd.rmdup.bam
ubuntu@ip-10-28-24-212:~/genomics_tutorial/strain1/remapping_to_reference$ ll
total 782252
drwxr-xr-x 2 ubuntu ubuntu      4096 Jan 18 18:30 ./
drwxr-xr-x 4 ubuntu ubuntu      4096 Jan 17 14:09 ../
-rw-r--r-- 1 ubuntu ubuntu 800991101 Jan 18 17:56 alignment.calmd.rmdup.bam
-rw-r--r-- 1 ubuntu ubuntu 15336 Jan 18 18:30 alignment.calmd.rmdup.bam.bai
-rw-r--r-- 1 ubuntu ubuntu    313 Jan 18 18:01 mappingstats.txt
ubuntu@ip-10-28-24-212:~/genomics_tutorial/strain1/remapping_to_reference$ █

```

We should obtain a .bai file (known as a BAM-index file).

Task 13: Obtain mapping statistics

Finally we can obtain some summary statistics.

```
samtools flagstat alignment.calmd.rmdup.bam > mappingstats.txt
```

This should only take a few seconds. Once complete view the mappingstats.txt file using a text-editor (e.g. gedit or nano) or the 'more' command.

```
6488111 + 0 in total (QC-passed reads + QC-failed read)
0 + 0 duplicates
6303345 + 0 mapped (97.15%:-nan%)
6488111 + 0 paired in sequencing
3245025 + 0 read1
3243086 + 0 read2
6287131 + 0 properly paired (96.90%:-nan%)
6287139 + 0 with itself and mate mapped
16206 + 0 singletons (0.25%:-nan%)
2 + 0 with mate mapped to a different chr
1 + 0 with mate mapped to a different chr (mapQ>=5)
```

So here we can see we have 6488111 reads in total, none of which failed QC. 97.15% of reads mapped to the reference genome and 96.90% mapped with the expected 500-600bp distance between them. 16206 read could not have their read-pair mapped.

2 reads have mapped to a different chromosome than their pair (1 has a mapping quality > 5 – this is a Phred scaled quality score much as we say in the FASTQ files). If there were any such reads they would likely due to repetitive sequences (e.g phage insertion sites) or an insertion of plasmid or phage DNA into the main chromosome.

Task 14: Cleanup

We have a number of leftover intermediate files which we can now remove to save space.

Type (all on one line):

```
rm alignment.bam alignment.calmd.bam alignment.sam  
alignment.sorted.bam *.aln
```

You should just have the alignment.calmd.rmdup.bam, the index file and the mappingstats.txt file left after this.


Well done! You have now mapped, filtered and sorted your first whole genome data-set!
Let's take a look at it!

Task 15: Load the Integrative Genomics Viewer


Within your amazon instance (make sure you don't accidentally load firefox on your desktop PC), load the Firefox web browser by double clicking the icon on the desktop. Go to <http://www.broadinstitute.org/software/igv/download>

You may find that you need to register to access the download page. Please do so if you need to.

Once you have registered you should reach the downloads page (similar to the screenshot below). Click on the 'Launch with 2GB' option. If Firefox asks you whether or not to Open or Save the file, click on 'Open'. For larger genomes you may need to use the 10Gb version (though you'll need a machine with at least this much RAM).



After a minute or so the IGV viewer should appear:




Task 16: Import the *E.coli* UTI89 reference genome to IGV

By default IGV does not contain our reference genome. We'll need to import it.

Click on 'Genomes ->Create genome file...'

Enter the following details into and click on 'Save'.



Import Genome

ID * (unique id, e.g. hg18)

Name *

☐ Fasta file is a directory

Fasta file * ...

Cytoband

Gene file ...


Alias file ...

Supply a sequence URL if defining a web-hosted genome (optional, not common). See user guide for more details.

Sequence URL

** required.
The sequence file (required) can be a FASTA file, a directory of FASTA files, or a zip of FASTA files. Optionally, specify a cytoband file to display the chromosome ideogram and an annotation file to display the gene track. See the documentation for descriptions of supported annotation formats.*


IGV will ask where it can save the genome file. Your home directory will be fine. Click 'Save' again.




Note that the genome has now been imported. Two chromosomes are listed.

Task 17: Load the BAM file

Load the alignment.calmd.rmdup.bam file. Note that IGV requires the .bai index file to also be in the same directory.




Once loaded your screen should look similar to the following. Note that you can load BAM files if you wish to compare different samples or the results of different mapping programs.




Note that at the moment, you cannot see any reads. We first need to select a chromosome and zoom in. To save memory, IGV will only display reads at a certain magnification (you can alter this in the View -> Preferences -> Alignments menu).

Select the CP000243.1 chromosome.



Either use the +/- keys to zoom in or use the zoom bar at the top right of the screen. At the bottom you can see a representation of the genes that have been annotated for the reference strain. The



following is a snapshot of the genome at position 25171850-2543949 – a 25kb window on the genome.

The gray graph at the top of the figure indicates the coverage of the genome by the grey reads below. The more reads mapping to a certain location, the higher the peak on the graph. You'll see a coloured line of blue, green or red in this coverage plot if there are any SNPs (single-nucleotide polymorphisms) present (there are none in the plot below). If there are any regions in the genome which are not covered by the reads, you will see these as gaps in the coverage graph. Sometimes these gaps are caused by repetitive regions, others are caused by genuine insertions/deletions in your new strain with respect to the reference.

Below the coverage graph is a representation of each read as it is mapped to the genome. Any areas of mismatch either due to inconsistent distances between paired-end reads or due to mutations are highlighted by a colour. If there is a consistent column of colour, this is indicative of a SNP or Indel. The brighter the colour, the higher the base-calling quality is estimated to be.

3.5: SNPs and Indels

The following 3 tasks are open-ended. Please take your time with these. Read the examples on the following page if you get stuck.

Task 18: Read about the alignment display format

Visit <http://www.broadinstitute.org/software/igv/AlignmentData>

Task 19: Manually Identify a region without reads mapping

What is this region? Is there a gene close-by? What sort of genes seem to result in areas of low/zero coverage? Why do you think this is? (think about repetitive sequences, what does BWA do if a read maps to multiple locations?)

Task 20: Identify SNPs and Indels manually

Zoom into the region at the start of chromosome CP000243.1. Can you find any SNPs? Which genes (if any) are they in? How reliable do they look? (Hint – look at the number of reads mapping, their orientation and how bright the base-calls are).


Zoom in to maximum magnification at the site of the SNP. Can you determine whether a SNP results in a synonymous (i.e. silent) or non-synonymous change in the amino acid? Can you use PDB (<http://www.rcsb.org/pdb/home/home.do>) or other resources to determine whether or not this occurs in a catalytic site or other functionally crucial region? (Note this may not always be possible).

What effect do you think this would have on the cell?

Example: Identifying SNPs manually

To help you there is an example pasted below for position 15,450-15,540bp on chromosome CP000243.1.

The first thing to note is that only discrepancies with respect to the reference are shown. If a read is entirely the same as its reference, it will appear entirely grey. Blue and red blocks indicate the presence of an 'abnormal' distance between paired-end reads. Note that unless this is consistent across most of the reads at a given position, it is not significant.



Here we have a G->C SNP. This changes the codon from CAG->GAG (remember this gene is on the reverse strand as indicated by the arrows) and results in a Gln->Thu mutation in the final protein product. As this is a cell-killing gene, this could be a very important change. Note that there is a single T base listed in one of the reads towards the bottom of the screen. Based on the frequency of C-bases vs T-bases at this position we can say that the T-base is likely to be a sequencing error and ignore it.


One additional check is that the SNPs occur on both read 1 and read 2. We can check this by looking at the direction of the grey reads. We can see that approximately half of the bases reporting the G->C mutation occur in read 1 (forward arrow), and half in read 2 (reverse arrow). This adds confidence to the base-call as it reduces the likelihood of this SNP being the result of a PCR duplication error.

Note that sequencing errors in Illumina data are quite common (look at the odd bases showing up in the screen above). We rely on depth of sequencing to average out these errors. This is why people often mention that a minimum median coverage of 20-30x across the genome is required for accurate SNP-

calling with Illumina data. This is not necessarily true for simple organisms such as prokaryotes, but for diploid and polyploid organisms it becomes important because each position may have one, two or many alleles changed.

Example: Identifying Indels manually

To help you there is an example pasted below for position 11,200-11,250bp on chromosome CP000243.1.



Much the same guidelines apply for indels as they do for SNPs. Here we have an A->G SNP which is preceded by the insertion of a C base in our sample compared to the reference. Again, we can see how much confidence we have that the insertion is real by the brightness of the colours and by ensuring that the indel is found on both read 1 and read 2.

The insertion is signified by the presence of a purple bar. Hover your mouse over it to get more details.

We can hover our mouse over the reference sequence to get details of the gene. We can see that it is a cytochrome C chaperone. Given that this indel changes the reading frame of the protein it could have very significant effects.

One can research the effect the effect of a SNP or Indel may have by finding the relevant gene at <http://www.uniprot.org> (or google 'yaaW uniprot' in this case).

It should be clear from this quick exercise that trying to work out where SNPs and Indels are manually is a fairly tedious process if there are many mutations. As such, the next section will look at how to obtain spread-sheet friendly summary details of these.

Recap: SNP/Indel identification

1. Only changes from the reference sequence are displayed in IGV
2. Genuine SNPs/Indels should be present on both read 1 and read 2
3. Genuine SNPs/Indels should be supported by a good (i.e. 20-30x) depth of coverage
4. Very important mutations (i.e. ones relied upon in a paper) should be confirmed via PCR/Sanger sequencing.

3.5: Automated analyses

Viewing alignments is useful when convincing yourself or others that a particular mutation is real rather than artefactual and for getting a feel for short read sequencing datasets. However, if we want to quickly and easily find variants we need to be able to generate lists of variants, in which gene they occur (if any) and what effect they have. We also need to know which (if any) genes are missing (i.e. have zero coverage).

3.5.1: Automated variant calling

To call variants we can use a number of packages (e.g. VarScan, GATK). However here, we will show you how to use the bcftools package which comes with samtools. First we need to generate a pileup file which contains only locations with the variants and pass this to bcftools.

Task 21: Identify SNPs and Indels using automated variant callers

Make sure you are in the ~/genomics_tutorial/strain1/remapping_to_reference directory. Type the following:

```
samtools mpileup
```

You should see a screen similar to the following

Usage: samtools mpileup [options] in1.bam [in2.bam [...]]

Input options:

- 6 assume the quality is in the Illumina-1.3+ encoding
- A count anomalous read pairs
- B disable BAQ computation
- b FILE list of input BAM files [null]
- C INT parameter for adjusting mapQ; 0 to disable [0]
- d INT max per-BAM depth to avoid excessive memory usage [250]
- E extended BAQ for higher sensitivity but lower specificity
- f FILE faidx indexed reference sequence file [null]

- G FILE exclude read groups listed in FILE [null]
- I FILE list of positions (chr pos) or regions (BED) [null]
- M INT cap mapping quality at INT [60]
- r STR region in which pileup is generated [null]
- R ignore RG tags
- q INT skip alignments with mapQ smaller than INT [0]
- Q INT skip bases with baseQ/BAQ smaller than INT [13]

Output options:

- D output per-sample DP in BCF (require -g/-u)
- g generate BCF output (genotype likelihoods)
- O output base positions on reads (disabled by -g/-u)
- s output mapping quality (disabled by -g/-u)
- S output per-sample strand bias P-value in BCF (require -g/-u)
- u generate uncompress BCF output

SNP/INDEL genotype likelihoods options (effective with ``-g'` or ``-u'`):

- e INT Phred-scaled gap extension seq error probability [20]
- F FLOAT minimum fraction of gapped reads for candidates [0.002]
- h INT coefficient for homopolymer errors [100]
- I do not perform indel calling
- L INT max per-sample depth for INDEL calling [250]
- m INT minimum gapped reads for indel candidates [1]
- o INT Phred-scaled gap open sequencing error probability [40]
- P STR comma separated list of platforms for indels [all]

Notes: Assuming diploid individuals.

If you are running this on your own datasets, please make sure you set the `-d` parameter if you have high coverage (i.e. > 100x coverage) per sample.

As the samtools mpileup command outputs an unfriendly output, we will pass it directly to the bcftools view command using the linux pipe (`|`). Type the following:

```
samtools mpileup -uf ~/genomics_tutorial/reference_sequence/Ecoli_UTI89.fna  
alignment.calmd.rmdup.bam | bcftools view -bv - > var.raw.bcf
```

This may take 15 minutes or so and will generate a BCF (Binary Call Format) file containing the raw unfiltered variant calls in a binary format.

This is not readable by humans, so let's use the bcftools view command and use the linux pipe (`|`) with the `vcfutils.pl varFilter` command. We can see what the options are for this program by typing in:

vcfutils.pl varFilter

Usage: vcfutils.pl varFilter [options] <in.vcf>

Options: -Q INT minimum RMS mapping quality for SNPs [10]
-d INT minimum read depth [2]
-D INT maximum read depth [10000000]
-a INT minimum number of alternate bases [2]
-w INT SNP within INT bp around a gap to be filtered [3]
-W INT window size for filtering adjacent gaps [10]
-1 FLOAT min P-value for strand bias (given PV4) [0.0001]
-2 FLOAT min P-value for baseQ bias [1e-100]
-3 FLOAT min P-value for mapQ bias [0]
-4 FLOAT min P-value for end distance bias [0.0001]
-e FLOAT min P-value for HWE (plus F<0) [0.0001]
-p print filtered variants

Note: Some of the filters rely on annotations generated by SAMtools/BCFtools.

We will use the -d option to limit variant calls to those positions where there are at least 10 reads. Type:

```
bcftools view var.raw.bcf | vcfutils.pl varFilter -d 10 > var.flt.vcf
```

Once complete, view the file using the 'more' command. You should see something similar to: (lines beginning with # are just comment lines explaining the output)

```
#CHROM POS ID REF ALT QUAL FILTER INFO FORMAT alignment.calmd.rmdup.bam
CP000243.1 10739 . ATTT ACTTT 214 . INDEL;DP=32;VDB=0.0803;AF1=1;AC1=2;DP4=0,0,8,11;MQ=51;FQ=-91.5 GT:PL:GQ 1/1:255,57,0:99
CP000243.1 11229 . T TC 214 . INDEL;DP=34;VDB=0.1028;AF1=1;AC1=2;DP4=0,0,14,11;MQ=44;FQ=-110 GT:PL:GQ 1/1:255,75,0:99
CP000243.1 15495 . G C 222 . DP=50;VDB=0.1028;AF1=1;AC1=2;DP4=0,0,25,17;MQ=52;FQ=-150 GT:PL:GQ 1/1:255,123,0:99
CP000243.1 35301 . G A 222 . DP=47;VDB=0.0948;AF1=1;AC1=2;DP4=0,0,14,26;MQ=51;FQ=-147 GT:PL:GQ 1/1:255,120,0:99
CP000243.1 126378 . T A 222 . DP=47;VDB=0.1028;AF1=1;AC1=2;DP4=0,0,18,20;MQ=54;FQ=-141 GT:PL:GQ 1/1:255,114,0:99
CP000243.1 131200 . C CA 214 . INDEL;DP=58;VDB=0.1028;AF1=1;AC1=2;DP4=0,0,19,15;MQ=54;FQ=-137 GT:PL:GQ 1/1:255,102,0:99
CP000243.1 210716 . G C 211 . DP=45;VDB=0.1028;AF1=1;AC1=2;DP4=0,0,19,19;MQ=47;FQ=-141 GT:PL:GQ 1/1:244,114,0:99
CP000243.1 314746 . A T 222 . DP=34;VDB=0.0948;AF1=1;AC1=2;DP4=0,0,12,14;MQ=53;FQ=-105 GT:PL:GQ 1/1:255,78,0:99
CP000243.1 326644 . G C 222 . DP=43;VDB=0.0803;AF1=1;AC1=2;DP4=0,0,18,19;MQ=47;FQ=-138 GT:PL:GQ 1/1:255,111,0:99
CP000243.1 336724 . A G 218 . DP=49;VDB=0.1028;AF1=1;AC1=2;DP4=0,0,20,19;MQ=45;FQ=-144 GT:PL:GQ 1/1:251,117,0:99
CP000243.1 444808 . G A 153 . DP=32;VDB=0.0624;AF1=1;AC1=2;DP4=0,0,12,10;MQ=44;FQ=-93 GT:PL:GQ 1/1:186,66,0:99
CP000243.1 444809 . C A 138 . DP=33;VDB=0.0803;AF1=1;AC1=2;DP4=0,0,12,10;MQ=41;FQ=-93 GT:PL:GQ 1/1:171,66,0:99
```

You can see the chromosome, position, reference and alternate allele as well as a quality score for the SNP. This a VCF file (Variant Call File). This is a standard developed for the 1000 genomes project. The full specification is given at <http://www.1000genomes.org/wiki/Analysis/Variant%20Call%20Format/vcf-variant-call-format-version-41>

The lines starting DP and INDEL contain various details concerning the variants. For haploid organisms, most of these details are not necessary.

Variant qualities:

Typically one should only accept variant calls over a certain quality threshold. Typically a threshold of 60 is used (i.e. a 1 in 1000000 chance of a mis-called variant). Here you can see that all these variants would pass these thresholds. However, for future reference, we can use the Linux 'awk' command to

filter the data on the quality column (i.e. column 6):

```
awk '($6>=60)' var.flt.vcf > out.snps.vcf4
```

Again viewing the final output file out.snps.vcf4 using a text-editor or the 'more' command should yield:

#CHROM	POS	ID	REF	ALT	QUAL	FILTER	INFO	FORMAT	alignment.calmd.rmdup.bam
CP000243.1	10739	.	ATTT	ACTTT	214	.	INDEL;DP=32;VDB=0.0803;AF1=1;AC1=2;DP4=0,0,8,11;MQ=51;FQ=-91.5	GT:PL:GQ	1/1:255,57,0:99
CP000243.1	11229	.	T	TC	214	.	INDEL;DP=34;VDB=0.1028;AF1=1;AC1=2;DP4=0,0,14,11;MQ=44;FQ=-110	GT:PL:GQ	1/1:255,75,0:99
CP000243.1	15495	.	G	C	222	.	DP=50;VDB=0.1028;AF1=1;AC1=2;DP4=0,0,25,17;MQ=52;FQ=-150	GT:PL:GQ	1/1:255,123,0:99
CP000243.1	35301	.	G	A	222	.	DP=47;VDB=0.0948;AF1=1;AC1=2;DP4=0,0,14,26;MQ=51;FQ=-147	GT:PL:GQ	1/1:255,120,0:99
CP000243.1	126378	.	T	A	222	.	DP=47;VDB=0.1028;AF1=1;AC1=2;DP4=0,0,18,20;MQ=54;FQ=-141	GT:PL:GQ	1/1:255,114,0:99
CP000243.1	131200	.	C	CA	214	.	INDEL;DP=58;VDB=0.1028;AF1=1;AC1=2;DP4=0,0,19,15;MQ=54;FQ=-137	GT:PL:GQ	1/1:255,102,0:99
CP000243.1	210716	.	G	C	211	.	DP=45;VDB=0.1028;AF1=1;AC1=2;DP4=0,0,19,19;MQ=47;FQ=-141	GT:PL:GQ	1/1:244,114,0:99
CP000243.1	314746	.	A	T	222	.	DP=34;VDB=0.0948;AF1=1;AC1=2;DP4=0,0,12,14;MQ=53;FQ=-105	GT:PL:GQ	1/1:255,78,0:99
CP000243.1	326644	.	G	C	222	.	DP=43;VDB=0.0803;AF1=1;AC1=2;DP4=0,0,18,19;MQ=47;FQ=-138	GT:PL:GQ	1/1:255,111,0:99
CP000243.1	336724	.	A	G	218	.	DP=49;VDB=0.1028;AF1=1;AC1=2;DP4=0,0,20,19;MQ=45;FQ=-144	GT:PL:GQ	1/1:251,117,0:99
CP000243.1	444808	.	G	A	153	.	DP=32;VDB=0.0624;AF1=1;AC1=2;DP4=0,0,12,10;MQ=44;FQ=-93	GT:PL:GQ	1/1:186,66,0:99
CP000243.1	444809	.	C	A	138	.	DP=33;VDB=0.0803;AF1=1;AC1=2;DP4=0,0,12,10;MQ=41;FQ=-93	GT:PL:GQ	1/1:171,66,0:99

This forms our definitive list of variants for this sample.

Task 22: Compare the variants found using this method to those you found in section 3.5

Can you see any variants which may have been missed? Often variants within a few bp of indels are filtered out as they could be spurious SNPs thrown up by a poor alignment. This is especially the case if you use non-gapped aligners such as Bowtie.

(Hint – check out the SNP/Indel combination in the example in section 3.5 again)

3.5.2: Quickly locating genes which are missing compared to the reference


We can use a command from the BEDTools package (<http://code.google.com/p/bedtools/>) to identify annotated genes which are not covered by reads across their full length.

Type the following one line:

```
coverageBed -abam alignment.calmd.rmdup.bam -b  
~/genomics_tutorial/reference_sequence/Ecoli_UTI89.gff > gene_coverage.txt
```

This should only take a minute or so. Once complete we can view the file in LibreOffice spreadsheet.


Launch 'LibreOffice Calc' from the Applications->Office menu.



Click 'Insert' and then select 'Sheet From File...'. Select the gene_coverage.txt file in the ~/genomics_tutorial/strain1/remapping_to_reference directory.

When the program asks you which delimiters to use, select 'Tab'. Make sure the others are deselected.

You may need to right click on the column labels for column I and select 'Column width'. Reduce this to 5”.



Column M represents the coverage of the particular GFF entry. A value of 1 is fully covered by reads. Less than one indicates regions with no reads mapping. Let's sort on column M in ascending order so

that anything missing with missing chunks shows up at the top. Select all columns and click on Data->Sort...

The screenshot shows the LibreOffice Calc interface with a spreadsheet containing genomic data. A 'Sort' dialog box is open, showing 'Sort Criteria' and 'Options' tabs. The 'Sort by' dropdown is set to 'Column M', and the 'Then by' dropdown is set to '- undefined -'. The 'Options' tab shows 'Ascending' selected for the first sort criterion.

	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S
1	CP000243.1	Genbank	gene	1047539	1050085	0	0	0	ID=CP000243.1	3226	2547	2547	1						
2	CP000243.1	Genbank	CDS	1047539	1050082	0	0	0	ID=CP000243.1	3226	2544	2544	1						
3	CP000243.1	Genbank	misc_feature	1047566	1050046	0	0	0	ID=CP000243.1	3156	2481	2481	1						
4	CP000243.1	Genbank	misc_feature	1047698	1049575	0	0	0	ID=CP000243.1	2354	1878	1878	1						
5	CP000243.1	Genbank	gene	2096593	2098170	0	0	0	ID=CP000243.1	1981	1578	1578	1						
6	CP000243.1	Genbank	CDS	2096593	2098167	0	0	0	ID=CP000243.1	1978	1575	1575	1						
7	CP000243.1	Genbank	misc_feature	2096593	2098155	0	0	0	ID=CP000243.1	1962	1563	1563	1						
8	CP000243.1	Genbank	misc_feature	2096929	2098128	0	0	0	ID=CP000243.1	1532	1200	1200	1						
9	CP000243.1	Genbank	gene	3142897	3146448	0	0	0	ID=CP000243.1	4405	3552	3552	1						
10	CP000243.1	Genbank	CDS	3142900	3146448	0	0	0	ID=CP000243.1	4399	3549	3549	1						
11	CP000243.1	Genbank	misc_feature	3142900	3146442	0	0	0	ID=CP000243.1	439									
12	CP000243.1	Genbank	misc_feature	3143011	3146439	0	0	0	ID=CP000243.1	42									
13	CP000243.1	Genbank	gene	4194207	4194638	0	0	0	ID=CP000243.1	5									
14	CP000243.1	Genbank	CDS	4194207	4194635	0	0	0	ID=CP000243.1	5									
15	CP000243.1	Genbank	misc_feature	4194219	4194635	0	0	0	ID=CP000243.1	5									
16	CP000243.1	Genbank	gene	129246	131909	0	0	0	ID=CP000243.1	31									
17	CP000243.1	Genbank	CDS	129246	131906	0	0	0	ID=CP000243.1	31									
18	CP000243.1	Genbank	misc_feature	129246	131906	0	0	0	ID=CP000243.1	31									
19	CP000243.1	Genbank	gene	261024	262508	0	0	0	locus_tag=UT189	18									
20	CP000243.1	Genbank	CDS	261027	262508	0	0	0	locus_tag=UT189	18									
21	CP000243.1	Genbank	misc_feature	261033	262451	0	0	0	locus_tag=UT189	17									
22	CP000243.1	Genbank	misc_feature	261033	262310	0	0	0	locus_tag=UT189	15									
23	CP000243.1	Genbank	gene	392314	393405	0	0	0	ID=CP000243.1	13									
24	CP000243.1	Genbank	CDS	392317	393405	0	0	0	ID=CP000243.1	13									
25	CP000243.1	Genbank	misc_feature	392377	393399	0	0	0	ID=CP000243.1	13									
26	CP000243.1	Genbank	gene	654052	656775	0	0	0	ID=CP000243.1	32									
27	CP000243.1	Genbank	CDS	654055	656775	0	0	0	ID=CP000243.1	32									
28	CP000243.1	Genbank	misc_feature	654055	656634	0	0	0	ID=CP000243.1	31									
29	CP000243.1	Genbank	misc_feature	654661	655386	0	0	0	ID=CP000243.1	8									
30	CP000243.1	Genbank	gene	785442	786683	0	0	0	ID=CP000243.1	16									
31	CP000243.1	Genbank	CDS	785445	786683	0	0	0	ID=CP000243.1	16									
32	CP000243.1	Genbank	misc_feature	785448	786680	0	0	0	ID=CP000243.1	16									
33	CP000243.1	Genbank	misc_feature	786222	786659	0	0	0	ID=CP000243.1	5									
34	CP000243.1	Genbank	misc_feature	911530	948388	0	0	0	note=possible	450									
35	CP000243.1	Genbank	gene	915637	918444	0	0	0	locus_tag=UT189	33									
36	CP000243.1	Genbank	CDS	915637	918441	0	0	0	locus_tag=UT189	33									

The screenshot shows the LibreOffice Calc interface with a spreadsheet containing genomic data. The formula bar shows a complex formula: `=locus_tag=UT189_P001;transl_table=11;product=hypothetical protein;protein_id=ABE10577.1;db_xref=GI:91075697;exon_number=1`. The spreadsheet data is as follows:

	A	B	C	D	E	F	G	H	I	J	K	L	M
1	CP000244.1	Genbank	start_codon	1	3	0	0	0	locus_tag=UT189_P001;transl_table=11;product=hypothetical protein;protein_id=ABE10577.1	1	2	3	0.666667
2	CP000244.1	Genbank	CDS	1	588	0	0	0	locus_tag=UT189_P001;transl_table=11;product=hypothetical protein;protein_id=ABE10577.1	477	587	588	0.9982993
3	CP000244.1	Genbank	gene	1	591	0	0	0	locus_tag=UT189_P001	480	590	591	0.9983079
4	CP000243.1	Genbank	gene	1047539	1050085	0	0	0	ID=CP000243.1;torA;locus_tag=UT189_C1061	3226	2547	2547	1
5	CP000243.1	Genbank	CDS	1047539	1050082	0	0	0	ID=CP000243.1;torA;unknown_transcript_1;Parent=CP000243.1;torA;locus_tag=UT189_C1061	3226	2544	2544	1
6	CP000243.1	Genbank	misc_feature	1047566	1050046	0	0	0	ID=CP000243.1;torA;unknown_transcript_2;Parent=CP000243.1;torA;locus_tag=UT189_C1061	3156	2481	2481	1
7	CP000243.1	Genbank	misc_feature	1047698	1049575	0	0	0	ID=CP000243.1;torA;unknown_transcript_3;Parent=CP000243.1;torA;locus_tag=UT189_C1061	2354	1878	1878	1
8	CP000243.1	Genbank	gene	2096593	2098170	0	0	0	ID=CP000243.1;ybtE;locus_tag=UT189_C2187	1981	1578	1578	1
9	CP000243.1	Genbank	CDS	2096593	2098167	0	0	0	ID=CP000243.1;ybtE;unknown_transcript_1;Parent=CP000243.1;ybtE;locus_tag=UT189_C2187	1978	1575	1575	1
10	CP000243.1	Genbank	misc_feature	2096593	2098155	0	0	0	ID=CP000243.1;ybtE;unknown_transcript_2;Parent=CP000243.1;ybtE;locus_tag=UT189_C2187	1962	1563	1563	1
11	CP000243.1	Genbank	misc_feature	2096929	2098128	0	0	0	ID=CP000243.1;ybtE;unknown_transcript_3;Parent=CP000243.1;ybtE;locus_tag=UT189_C2187	1532	1200	1200	1
12	CP000243.1	Genbank	gene	3142897	3146448	0	0	0	ID=CP000243.1;trpR;locus_tag=UT189_C3221	4405	3552	3552	1

There appear to be very few missing regions. In fact there is only one missing base near the origin of replication for this bacteria. Because the aligner doesn't know that the genome is circular (and thus to allow reads to span the beginning and end of the chromosome), the very first base is missing from the alignment. It is however more than likely to be present in the reads.

3.5.3: Evaluating the impact of variants

So, aside from a few variants we've discovered, this strain of *E.coli* does not seem to have lost anything and seems very closely related to the reference genome. So let's take a closer look at the variants. We'd like to obtain a list of genes in which these variants occur and whether they result in amino acid changes.

To do this we'll use a custom perl script developed by David Studholme and Konrad Paszkiewicz.

We'll just need the reference annotation, sequence and the VCF file containing the SNPs.


Task 23: Determine the effect of variants

Type (all on one line):

```
/usr/local/scripts/snp_comparator.pl 10
~/genomics_tutorial/reference_sequence/Ecoli_UTI89.fna
~/genomics_tutorial/reference_sequence/Ecoli_UTI89.gff out.snps.vcf4 >
snp_report.txt
```

This program takes the information from the reference sequence and annotation, and the VCF snp files and determines whether the variant occurs within a gene, and if so the effect of each mutation.

Once complete, view the snp_report.txt file using the more command:



Chrom	Pos	Ref	out.snps.vcf4	Gene description	Status
CP000243.1	10739	ATTT	ACTTT	UTI89_C0012 putative oxidoreductase	,Indel
CP000243.1	11229	T	TC	UTI89_C0012 putative oxidoreductase	UTI89_C0013 positive regulator for sigma 32 heat shock promoters ,Indel
CP000243.1	126378	T	A	UTI89_C0122 hypothetical protein	,non-silent agt -> agA;
CP000243.1	131200	C	CA	UTI89_C0127 pyruvate dehydrogenase E1 component	,Indel
CP000243.1	15495	G	C	UTI89_C0018	,non-silent cag -> Gag;
CP000243.1	210716	G	C	UTI89_C0199 DNA polymerase III alpha subunit	,non-silent gcc -> Ccc;
CP000243.1	314746	A	T	UTI89_C0304 putative ferredoxin	,non-silent tgt -> Agt;
CP000243.1	326644	G	C	UTI89_C0317 putative LysR-like transcriptional regulator	,non-silent tgt -> tCt;
CP000243.1	336724	A	G	UTI89_C0323 2%2C5-diketo-D-gluconic acid reductase A	,non-silent cag -> cGg;
CP000243.1	35301	G	A	UTI89_C0037 Carbamoyl-phosphate synthase large chain	,non-silent gcg -> Acg;
CP000243.1	444808	G	A	UTI89_C0431 protein-export membrane protein SecF	,silent ttg -> ttA;
CP000243.1	444809	C	A	UTI89_C0431 protein-export membrane protein SecF	,non-silent cag -> Aag;

In later workshops we will see how we can use this program to compare results between different strains.

Task 24: Check each variant in IGV

N.B. If a variant doesn't seem to match what the snp_report file says, check the reverse reading frames.

Task 25: Check each non-synonymous SNP and Indel for a possible link to haemorrhaging phenotypes (often known as EHEC)

Use the web for this. Search UniProt (<http://www.uniprot.org>) or other databases for evidence that any of these proteins may be linked to the outbreak.

That concludes the first part of the course. You have successfully, QC'd, filtered, remapped and analysed a whole bacterial genome! Well done!

In the next installment we will be looking at how to extract and assemble unmapped reads. This will enable us to look at material which may be present in the strain of interest but not in the reference sequence.

2013 Workshop on Genomics, Cesky Krumlov

Part 4: Short read genomics: Assembly of unmapped reads

Instructors:

- Konrad Paszkiewicz k.h.paszkiewicz at exeter ac uk

Objectives:

By the end of the workshop you will be expected to be able to:

- Extract reads which do not map to the reference sequence
- Assemble these reads de novo using Velvet
- Generate summary statistics for the assembly
- Identify potential genes within the assembly
- Search for matches within the Swissprot database via BLAST and against the Pfam database
- Visualize the taxonomic distribution of BLAST hits
- Perform gene prediction and annotation using RAST

4.1 Introduction

In this section of the workshop we will continue the analysis of a strain of *E.coli* which is involved in urinary tract infections. In the previous section we cleaned our data, checked QC metrics, mapped our data and obtained a list of variants and an overview of any missing regions.

Now, we will examine those reads which did not map to the reference genome. We want to know what these sequences represent. Are they novel genes, plasmids or just contamination?

To do this we will extract unmapped reads, evaluate their quality, prepare them for de novo assembly, assemble them using Velvet, generate assembly statistics and then produce some annotation via Pfam, BLAST and RAST.

4.2 Extraction and QC of unmapped reads

Task 1: Extract the unmapped reads

First of all make sure you are in the ~/genomics_tutorial/strain1 directory (hint: use the cd command). Then create a directory called denovo_assembly in which we will do our de novo assembly and analysis.

```
mkdir denovo_assembly/  
cd denovo_assembly/
```

Now we will use the bam2fastq program (<http://www.hudsonalpha.org/gsl/software/bam2fastq.php>) to extract from the BAM file just those reads which did NOT map to the reference genome. The bam2fastq program has a number of options, most of which are self-explanatory. Type (all on one line):

```
bam2fastq --no-aligned -o unaligned\#.fastq  
../remapping_to_reference/alignment.calmd.rmdup.bam
```

The --no-aligned option means only extract reads which did not align. The -o unaligned\# means dump read 1 into a file called unaligned_1.fastq and read 2 into a file unaligned_2.fastq. Below we can see that the program has successfully created the two files.

```
ubuntu@ip-10-28-24-212:~/genomics_tutorial/strain1/denovo_assembly$ bam2fa:  
This looks like paired data from lane 0.  
Output will be in unaligned_1 and unaligned_2  
6488111 sequences in the BAM file  
184766 sequences exported  
WARNING: 16206 reads could not be matched to a mate and were not exported  
ubuntu@ip-10-28-24-212:~/genomics_tutorial/strain1/denovo_assembly$ ls  
unaligned_1  unaligned_2  
ubuntu@ip-10-28-24-212:~/genomics_tutorial/strain1/denovo_assembly$
```

Note that some reads were singletons (i.e. one read mapped to the reference, but the other did not). These will not be included in this analysis.

Task 2: Check that the number of entries in both fastq files is the same. Also check that the last few entries in the read 1 and read 2 files have the same header (i.e. that they have been correctly paired).

Task 3: Evaluate QC of unmapped reads

Use the fastqc program to look at the statistics and QC for the unaligned_1.fastq and unaligned_2.fastq files.

Do these look reasonably good? Remember, some reads will fail to map to the reference because they are poor quality, so the average scores will be lower than the initial fastqc report we did in the remapping workshop. The aim here is to see if it looks as though there are reads of reasonable quality which did not map.

Assuming these reads look ok, we will proceed with preparing them for de novo assembly.

4.3 De novo assembly

Task 3: Prepare reads for de novo assembly

We will be using the Velvet (<http://www.ebi.ac.uk/~zerbino/velvet/>) assembler to stitch together the unmapped reads. This will enable us to obtain large enough contigs to determine which genes, plasmids or contaminant material is present (if any!). However, Velvet requires that the reads are all in a single file and arranged such that read 1 is followed by its corresponding read 2 pair.

Fortunately Velvet comes with a script to do this for us. Type (all on one line):

```
perl /usr/local/velvet/shuffleSequences_fastq.pl unaligned_1.fastq  
unaligned_2.fastq unaligned.fastq
```

Once this is complete, (it should only take a few seconds), check to make sure that the entries are indeed arranged as read 1 followed by read 2 in unaligned.fastq (just use more and tail to do a visual check).

Task 4: Learn more about de novo assemblers

To understand more about de-novo assemblers, read the technical note at:

http://bio-ruby.ex.ac.uk/technote_denovo_assembly_ecoli.pdf

N.B. You will also learn more in the next section so don't worry if it doesn't all make sense immediately.

Task 5: Generate the Velvet hash table

Velvet requires a hash table which stores all of the k-mer and paired-end information prior to trying to assemble the data. Note that velvet ignores quality scores, so whether you pass it a FASTA or a FASTQ file doesn't matter. However, it will automatically convert all N bases to A (this is because to save memory it can only represent 4 bases internally). It pays to remove sequences containing Ns at the initial filtering stage to avoid introducing ambiguities.

```
velveth assembly/ 31 -fastq -shortPaired unaligned.fastq
```

```
ubuntu@ip-10-28-24-212:~/genomics_tutorial/strain1/denovo_assembly$ v
[0.000001] Reading FastQ file unaligned.fastq
[0.285978] 168560 reads found.
[0.285996] Done
[0.286087] Reading read set file assembly/Sequences;
[0.317796] 168560 sequences found
[0.496636] Done
[0.496663] 168560 sequences in total.
[0.496733] Writing into roadmap file assembly/Roadmaps...
[0.548479] Inputting sequences...
[0.548506] Inputting sequence 0 / 168560
[1.535148] === Sequences loaded in 0.990783 s
[1.535218] Done inputting sequences
[1.535228] Destroying splay table
[1.610020] Splay table destroyed
ubuntu@ip-10-28-24-212:~/genomics_tutorial/strain1/denovo_assembly$
```

This will perform a de novo assembly using a k-mer length of 31 and put the results into the assembly/ directory. Note that if you have a large number of reads, you may well need to increase the kmer length to reduce memory requirements. This will have the effect of excluding reads shorter than the k-mer size (e.g. selecting a k-mer length of 41 will exclude all reads less than 39bp long).

Task 6: Perform denovo assembly

We can now (at last!) perform the de novo assembly of the unmapped reads. To do this we will use the program velvetg (stands for velvet graph). Velvetg comes with many options, and if you are performing de novo assembly of a full complex genome you may well need to tinker with these parameters to get the best results. (see the Velvet manual for more detail). Type (all on one line):

```
velvetg assembly/ -cov_cutoff auto -exp_cov auto  
-unused_reads yes -very_clean yes
```

```

[0.292825] Removing reference contigs with coverage < 60.503349...
[0.292835] Concatenation...
[0.292841] Renumbering nodes
[0.292847] Initial node count 1
[0.292853] Removed 0 null nodes
[0.292860] Concatenation over!
[0.293974] Writing contigs into ./contigs.fa...
[0.302388] Writing into stats file ./stats.txt...
[0.302590] Writing into graph file ./LastGraph...
[0.311255] Estimated Coverage = 121.006697
[0.311272] Estimated Coverage cutoff = 60.503349
Final graph has 1 nodes and n50 of 3306, max 3306, total 3306, using 19992/168560 reads

```

Note that velvetg runs quickly. Normally this is not the case, but because we have so few reads (and because this example is artificial) it speeds through the assembly in seconds.

The options used with velvetg -cov_cutoff and -exp_cov were set to automatic. -cov_cutoff specifies how frequently a k-mer should appear to avoid being thrown out as noise (i.e. sequencing errors or contamination). -exp_cov sets the expected level of coverage (i.e. how many times each k-mer should be occur). This allows the assembler to make an educated guess as to how many copies of a sequence are present. This is especially useful in repetitive regions.

The -unused_reads flag and -very_clean flag just mean that velvetg should output any reads which were not included in the assembly and that it should delete as many intermediate files as possible once it has finished.

If we change to the assembly directory and list its contents we should see several files:

```

ubuntu@ip-10-28-24-212:~/genomics_tutorial/strain1/denovo_assembly/assembly$ ls
Log PreGraph UnusedReads.fa contigs.fa stats.txt
ubuntu@ip-10-28-24-212:~/genomics_tutorial/strain1/denovo_assembly/assembly$ █

```

Let's take them in turn:

1. Log

```
Mon Jan 30 22:05:45 2012
velvetg auto_data_21 -clean yes -exp_cov 121 -cov_cutoff 56.0061658364741
Version 1.2.02
Copyright 2007, 2008 Daniel Zerbino (zerbino@ebi.ac.uk)
This is free software; see the source for copying conditions. There is NO
warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.
Compilation settings:
CATEGORIES = 2
MAXKMERLENGTH = 101

Final graph has 1 nodes and n50 of 3306, max 3306, total 3306, using 19992/168560 reads
```

This lists the version number of Velvet, as well as some summary statistics regarding the assembly. (read http://en.wikipedia.org/wiki/N50_statistic for more details about what N50 means). We'll come back to this file in the next section.

2. Pregraph

This is a file generated by Velvetg when generating the k-mer graph but is normally not required.

3. UnusedReads.fa

This is a FASTA formatted file containing the reads which could not be assembled into contigs larger than twice the kmer length (i.e. $2 \times 31 = 62\text{bp}$). This could be because they represent highly repetitive sequences, or because the sequences are shorter than the kmer length selected. (Don't worry if your file looks different to the example below).

```
>SEQUENCE_1_length_40
CGTCTTTGGTGTACAGGACGGTCACGATGATGTGGTGCTG
>SEQUENCE_2_length_39
TGCGGGAGAGTCCCACAACACGCAGGACTCTTTTGCTTT
>SEQUENCE_3_length_39
GTGGCAATCTCTATAATGCGATCGCCGTCGGGGTCCTGA
>SEQUENCE_4_length_40
CCCCTTCCACGATCCCTTTGACATATACGACTTGCGCCGC
>SEQUENCE_5_length_40
CCCTACAAAAGTACAAAAGCCAGACACCAGGGTCACGACG
>SEQUENCE_6_length_39
CGCAGCGCCGATGGCGACCTGCTTTTCGCGACCCGTAAC
>SEQUENCE_7_length_40
GTGGTTCCACCACATCCGCTGGTTTATATGTTGCCAATA
>SEQUENCE_8_length_40
AGAACCGACACCATTTTCCGTATGGGTGAACATCGTCCCC
>SEQUENCE_9_length_39
CGCGTCAAAAAGGCGATATCATTTCATCCCTGACCTGGGGG
>SEQUENCE_10_length_40
ACTCTCTTTATTTACCGCATTACGCTCCGTTTCTGGCGTA
```

4. stats.txt

```
ubuntu@ip-10-28-24-176:~/genomics_tutorial/strain1/denovo_assembly/auto_data_21$ more stats.txt
ID      lgth    out    in      long_cov    short1_cov    short1_Ocov    short2_cov    short2_Ocov
1       3306    1      1       0.000000    120.625832    119.901391    0.000000    0.000000
```

We can see that Velvet has managed to assemble 1 supercontig (node). This file is a simple tab-delimited description of the contigs. The column names are pretty much self explanatory. Note however that contig lengths are given in k-mers. To obtain the length in nucleotides of each node you simply need to add $k - 1$, where k is the hash length used in velvet. The in and out columns correspond to the number of arcs on the 5' and 3' ends of the contig respectively. The coverages in columns short1 cov, short1 Ocov, short2 cov, and short2 Ocov are provided in k-mer coverage.

The relation between k-mer coverage C_k and standard (nucleotide-wise) coverage C is $C_k = C * (L - k + 1)/L$ where k is your hash length, and L your read length.

5. contigs.fa

This contains the supercontigs as assembled by velvetg in FASTA format. It constitutes the main output file. Note that this does not contain any annotation – Velvet only deals with assembly of DNA.

Task 7: Count number of unassembled reads in UnusedReads.fa and the number of assembled contigs in contigs.fa. Does it tally with the Log file?

Hint, use the grep -c command

4.3 More denovo assembly statistics

Now that we have a de novo assembly, we can generate some basic statistics for it using a script called fasta_summary.pl. Note that we can use this for any FASTA formatted file. Make sure you are in the ~/genomics_tutorial/strain1/denovo_assembly/assembly directory.

Task 8: Generate summary statistics

```
fasta_summary.pl -i contigs.fa -o assembly_statistics -t contig -r 1
```

```
cd assembly_statistics
```

```
ls
```

```

ubuntu@ip-10-28-24-212:~/genomics_tutorial/strain1/denovo_assembly/assembly$ cd assembly_statistics/
ubuntu@ip-10-28-24-212:~/genomics_tutorial/strain1/denovo_assembly/assembly/assembly_statistics$ ls
histogram_bins.dat      sorted_contigs.fa  sum_reads_vs_read_len.dat      summed_contig_lengths.dat
histogram_bins.dat.png  stats.txt          sum_reads_vs_read_len.dat.png  summed_contig_lengths.dat.png
ubuntu@ip-10-28-24-212:~/genomics_tutorial/strain1/denovo_assembly/assembly/assembly_statistics$

```

Note that this directory contains a file called stats.txt, but that this is a different file from that produced by Velvet. View the *.txt files using a text-editor (e.g. gedit, nano) or the more command.

Statistics for contig lengths:

```

Min contig length:      3,326
Max contig length:      3,326
Mean contig length:     3326.00
Standard deviation of contig length: 0.00
Median contig length:   3,326
N50 contig length:      3,326

```

Statistics for numbers of contigs:

```

Number of contigs:      1
Number of contigs >=1kb: 1
Number of contigs in N50: 1

```

Statistics for bases in the contigs:

```

Number of bases in all contigs: 3,326
Number of bases in contigs >=1kb: 3,326
GC Content of contigs: 43.42 %

```

Simple Dinucleotide repeats:

```

Number of contigs with over 70% dinucleotide repeats: 0.00 % (0 contigs)
AT: 0.00 % (0 contigs)
CG: 0.00 % (0 contigs)
AC: 0.00 % (0 contigs)
TG: 0.00 % (0 contigs)
AG: 0.00 % (0 contigs)
TC: 0.00 % (0 contigs)

```

Simple mononucleotide repeats:

```

Number of contigs with over 50% mononucleotide repeats: 0.00 % (0 contigs)
AA: 0.00 % (0 contigs)
TT: 0.00 % (0 contigs)
CC: 0.00 % (0 contigs)
GG: 0.00 % (0 contigs)

```

Note that Velvet was only able to assemble 3,326bp in 1 contig with a GC content slightly lower than that of *E.coli*. Typically with a real dataset you would see many contigs here with a wide variety of lengths – assembling just a single contig is very rare! In this case, it is because we are using simulated data rather than 'dirty' real data.

You can also use gthumb to view various plots which have the extension .png. In this case they are not very informative, but in future projects they can be very valuable when faced with thousands of contigs.

4.4 Analysing the de novo assembled reads

Now that we have assembled the reads and have a feel for how much (or in this case, how little) data we have, we can set about analysing it. By analysing, we mean identifying which genes are present, which organism they are from and whether they form part of the main chromosome or are an independent unit (e.g. plasmid).

We are going to take a 3-prong approach. The first will simply search the nucleotide sequences of the contigs against the NCBI non-redundant database. This will enable us to identify the species to which a given contig matches best (or most closely). The second will call open reading frames within the contigs and search those against the Swissprot database of manually curated (i.e. high quality) annotated protein sequences. Finally, we will search the open reading frames against the Pfam database of protein families (<http://pfam.sanger.ac.uk>).

Why not just search the NCBI blast database? Well, remember nearly all of our biological knowledge is based on homology – if two proteins are similar they probably share an evolutionary history and may thus share functional characteristics. Metrics to define whether two sequences are homologous are notoriously difficult to define accurately. If two sequences share 90% sequence identity over their length, you can be pretty sure they are homologous. If they share 2% they probably aren't. But what if they share 30%? This is the notorious twilight zone of 20-30% sequence identity where it is very difficult to judge whether two proteins are homologous based on sequence alone.

To help overcome this searching more subtle signatures may help – this is where Pfam comes in. Pfam is a database which contains protein families identified by particular signatures or patterns in their protein sequence. These signatures are modeled by Hidden Markov Models (HMMs) and used to search query sequences. These can provide a high level annotation where BLAST might otherwise fail. It also has the advantage of being much faster than BLAST.

Task 9: Obtain open reading frames

The first task is to call open reading frames within the contigs. These are designated by canonical start and stop codons and are usually identified by searching for regions free of stop codons. We will use the EMBOSS package program `getorf` to call these.

We will use codon table 11 which defines the bacterial codon usage table (<http://www.ncbi.nlm.nih.gov/Taxonomy/Utils/wprintgc.cgi>) and state that the sequences we are dealing with are not circular (they are nowhere near long enough!). We will also restrict the ORFs to just those sequences longer than 300 nucleotides (i.e. 100 amino acids). We will store the results in file `contigs.orf.fa`.

```
getorf -table 11 -circular N -minsize 300 -sequence contigs.fa -outseq contigs.orf.fa
```

If we look at the output file we can see that it is a FASTA formatted file containing the name of the contig on which the ORF occurs, followed by an underscore and a number (e.g. `_1`) to indicate the number of the ORF on that contig. The numbers in square brackets indicate the start and end position of the ORF on the contig (i.e. in nucleotide space). So the first ORF occurs on NODE 1 and is between position 1073 and 1948. Whereas the fourth ORF occurs between positions 606-989 on the reverse

strand. This is a relatively short peptide and is unlikely to be a genuine peptide.

Also note that many ORFs do not start with a Methionine. This is because by default the getorf program calls ORFs between stop codons rather than start and stop codons. Primarily this is to avoid spurious ORFs due to Met residues within a protein sequence and to ensure untranslated regions are captured.

```
>NODE_1_length_3306_cov_120.625832_1 [1073 - 1948]
RVSEGRSGRAPYVSFSPYALLCRSASCPERYISVLFSTSKRVCMLFWSSAASSCSFSHVM
ACSSASSASSFSSSVRLWLFMNPDMLSAVCCCLFIFLFSFCLSSATCDYIAHHFSTVLP
PVFCRRTFQSDNTVTAKKQQCFVGNLQGTGHVQLLYRAYMHAITITLVRASPRHPMSD
TEPCFMTKRSGSNTGRRAMSHPVRLTAEEDQEIRKRAAECGKTVSGFLRAAALGKKVNSL
TDDRALKEVMRLGALQKKLFIDGKRVGDREYAEVLIAIKEYHRALLSRLMAD
>NODE_1_length_3306_cov_120.625832_2 [3324 - 2566] (REVERSE SENSE)
YTIPKFCDDDRANLFEFLSEEGITITEDNNNDLNCKHQYIMTTSYGDRVRAKIYKRGSIQ
FQGYLQIASWINDFMCILNMKEVVEQKNKEFNVDIKKETIESELHSLPKSIDKIHED
IKKQLSSSLIMKKIDVEMEDYSTYCFSAIRAIEGFIYQILNDVCNPSSSKNLGEYFTENK
PKYIIREIHQETINGEIAEVLCECYTYWHENRHGLFHMKPGIADTKTINKLESIAIIDTD
CQLKDGVARLKL
>NODE_1_length_3306_cov_120.625832_3 [2575 - 2189] (REVERSE SENSE)
VEIMKKDKKYQIEAIKNKDKTLFIVYATDIYSPSEFFSKIESDLKKKSKGDVFFDLIIP
NGGKKDRYVYTSFNGGKFSSYTLNKVTKTDEYNDLSELSASFCKNFDKINVNLLSKATS
FALKKGIPI
>NODE_1_length_3306_cov_120.625832_4 [989 - 606] (REVERSE SENSE)
SRKEHVSQRQSTGRSQHRRRFSIGSAPLTSITKIDAEARGGETRQDYKDTRRFPPEAPSC
ALLFRPCHLPDTCPPFSLREAWRFLIAHAGISVRCRSFAPSWALCTNPPFSPTAAPYPV
TIVLSPTR
>NODE_1_length_3306_cov_120.625832_5 [387 - 1] (REVERSE SENSE)
INRSIKSTNSYSVDKTRSVHIQKKQRIDQMAQNPFKALNINIDKIESALTQNGVTNYSS
NVKNERETHISGYTGIDFLIKLMPSSGNTTIGRASGQNNTYFDEIALIIKENCLYSDTK
NFEYTIPKF
```

Task 10: Search open reading frames against NCBI non-redundant database

The first thing we can do with these open reading frames is to search them against the NCBI non-redundant database of protein sequences to see what they may match.

Here we will perform a BLAST search using the non-redundant (nr) database, using the blastp program and store the results in contigs.orf.blastp. We'll apply an e-value (expectation value) (<http://www.ncbi.nlm.nih.gov/BLAST/tutorial/Altschul-1.html>) cutoff of 1e-06 to limit ourselves to statistically significant hits (i.e. in this case 1 in 1 million likelihood of a hit to a database of this size by a sequence of this length). The -b and -v flags tell blastall to only display the top 10 results for each hit, the -a that it should use 4 CPU cores and -I that it should include general identifier (GI) numbers in the output.

Type (all on one line):

```
blastall -d nr -p blastp -i contigs.orf.fa -o contigs.orf.blastp -e 1e-06 -a 4 -I T -b 10
```

This should take a few minutes. Once complete view the contigs.orf.blastp file in your favourite text-editor (e.g. nano, gedit etc).

Task 11: Review the BLAST format

Have a read of <http://cs124.cs.ucdavis.edu/appendices/BlastResults.html>

```

Sequences producing significant alignments:
Score      E
(bits)     Value

ref|YP_002221413.1| plasmid mobilization [Salmonella enterica su... 476    e-132
gb|EGD63862.1| plasmid mobilization [Escherichia coli 0157:H7 st... 424    e-117
gb|AEG39558.1| Plasmid mobilization protein [Escherichia coli NA... 412    e-113
ref|NP_277065.1| hypothetical protein pCRP3p04 [Citrobacter rode... 223    2e-56
gb|AER93362.1| plasmid mobility protein [Escherichia coli]         219    3e-55

>ref|YP_002221413.1| plasmid mobilization [Salmonella enterica subsp. enterica serovar
    Virchow str. SL491]
gb|ACH87666.1| plasmid mobilization [Salmonella enterica subsp. enterica serovar
    Virchow str. SL491]
    Length = 272

Score = 476 bits (1226), Expect = e-132, Method: Compositional matrix adjust.
Identities = 234/272 (86%), Positives = 237/272 (87%)

Query: 21  LCRSASCPERYISVLFSTSKRVCMLFWXXXXXXXXXXXXXXXXXXXXXXXXXXXXVRLWLF 80
          +CRSASCPERYISVLFSTSKRVCMLFW                                VRLWLF

```

We can see that the first ORF has a hit to a plasmid mobilization protein. It seems possible that this single 3kb stretch of DNA sequence which has been assembled could in fact be part of or a whole plasmid.

However, this could simply be a common protein which is shared by many plasmids and/or genomes. Given that this protein is labelled as a 'plasmid mobilization' protein, the chances are that it shares a considerable portion of its sequence with other plasmids. As such the protein sequence is not the best method to search for the source of this sequence. A better way is to compare the nucleotide sequence and use the added variation caused by synonymous and non-coding regions to identify the likely source of this extra DNA.

Task 12: Search contigs against NCBI non-redundant database

The following command executes a nucleotide BLAST search (blastn) of the sequences in the contigs.fa file against the non-redundant database. Again we restrict ourselves to 10 results per hit and an e-value cutoff of 1e-06.

```
blastall -d nt -p blastn -i contigs.fa -o contigs.fa.blastn -e 1e-06 -b 10 -v 10 -I T -a 4
```

Opening the results file should yield:

```
Query= NODE_1_length_3306_cov_120.625832
      (3326 letters)

Database: All GenBank+EMBL+DDBJ+PDB sequences (but no EST, STS,
GSS, environmental samples or phase 0, 1 or 2 HTGS sequences)
      15,769,693 sequences; 40,107,367,207 total letters

Searching.....done

Sequences producing significant alignments:
```

	Score (bits)	E Value
gi 356601232 gb CP002797.2 Escherichia coli NA114, complete genome	2878	0.0
gi 37811584 gb AY333433.1 Salmonella typhimurium plasmid pU302S...	2573	0.0
gi 67551170 gb AY913943.1 Escherichia coli plasmid pColE1-H22, ...	2407	0.0
gi 353234058 emb HE603113.1 Escherichia coli HUSEC41 plasmid pH...	2399	0.0
gi 141327082 gb CP000642.1 Shigella sonnei Ss046 plasmid pSS046...	2391	0.0
gi 7262429 dbj AB040037.1 Escherichia coli plasmid p9705 DNA, c...	2367	0.0
gi 15723943 gb AF311902.1 AF311902 Citrobacter rodentium strain ...	2359	0.0
gi 188504043 gb EU675685.1 Escherichia coli strain 86-24 seroty...	2327	0.0
gi 4589704 dbj AB011548.2 Escherichia coli 0157:H7 str. Sakai p...	2327	0.0
gi 3152962 emb Y14016.1 Escherichia coli plasmid p4821 mobA gen...	2325	0.0

```
>gi|356601232|gb|CP002797.2| Escherichia coli NA114, complete genome
      Length = 4971461

Score = 2878 bits (1452), Expect = 0.0
Identities = 1581/1624 (97%)
Strand = Plus / Minus

Query: 546      gacaaactcctctaaatcagtaaccaatggctgctgccagtggcgataagtcgtgtctta 605
               |||||
Sbjct: 4859256 gacaaactcctctaaatcagttaccagtgctgctgccagtggcgataagtcgtgtctta 4859197
```

The plasmid encoded is in-fact *Escherichia coli* 0157:H7 plasmid pOSAK1/pSP70 . It is NOT the top hit. This is because plasmids are often very similar (particularly near transposase sites). Although BLAST does by default filter out low complexity sequence, in such a small plasmid, it may have difficulty picking out the correct plasmid. In this example we can get a clue by looking at the amount of coverage the plasmid gets from the assembled contigs. We find that only *Escherichia coli* 0157:H7 plasmid pOSAK1 has almost full coverage, the remainder (even though they have better scores and e-values) are not the best match from a biological perspective.

Task 15: Find out more about the E.coli pOSAK1 plasmid

Makino, K. et al Complete Nucleotide Sequences of 93-kb and 3.3-kb Plasmids of an Enterohemorrhagic Escherichia coli O157:H7 Derived from Sakai Outbreak . DNA Research 5,(1-9) 1998

Additional checks (not necessary here, but will be for other genomes!):

Task 13: Check that contigs do not appear in the reference sequence

In theory, the unmapped reads used to generate the contigs should not assemble into something which will map against the genome. However, it is always possible (especially with more complex genomes), that this might happen. To double check:

```
bl2seq -i contigs.fa -j ~/genomics_tutorial/reference_sequence/Ecoli_UTI89.fna -p  
blastn -e 1e-06
```

Here we use bl2seq (part of the BLAST package) to compare two sequences against each other. Unlike the previous examples where we have searched against a database of sequences, here we are doing a simple search of the contigs against the reference genome sequencing. Running this, you should find that no hits are found with an e-value less than 1e-06. Therefore we can say with some degree of confidence that it is novel plasmid in this strain.

Task 14: Run open reading frames through pfam_scan

Pfam is a database of protein families. They are grouped together using a number of criteria based on their function. For more information read <http://en.wikipedia.org/wiki/Pfam>. Pfam is grouped into several databases depending on the level of curation. Pfam-A is high-quality manual curation and consists of around 12,500 families. Pfam-B is full of automated predictions which may be informative but should not be relied upon without additional evidence. Pfam will also search for signatures of active-sites if you specify the correct flag.

Here we want to search the Pfam database of Hidden Markov Models to see which protein families are contained within this contig. You'll notice that this runs considerably faster than BLAST. Here we search using the contigs.orf.fa file against the Pfam databases in /usr/local/PfamScan and output the results to contigs.orf.pfam. We'll use 4 cpu cores for the search and state that we want to search PfamB entries as well as active site residues.

```
pfam_scan.pl -fasta contigs.orf.fa -dir /usr/local/PfamScan/ -outfile  
contigs.orf.pfam -cpu 4 -pfamB -as
```

```

# <seq id> <alignment start> <alignment end> <envelope start> <envelope end> <hmm acc> <hmm name> <type> <hmm start> <hmm end> <hmm length> <bit score> <E-value> <significance> <clan> <predicted_active_site_residues>

NODE_1_length_3306_cov_120.625832_1 123 237 108 238 PB002473 Pfam-B_2473 Pfam-B 31 146 147 28.6 1.2e-06 NA NA
NODE_1_length_3306_cov_120.625832_1 194 238 188 255 PB011895 Pfam-B_11895 Pfam-B 7 51 121 26.5 3.5e-06 NA NA
NODE_1_length_3306_cov_120.625832_1 203 236 202 238 PF01402.16 RHH_1 Domain 4 37 39 28.8 6.4e-07 1 CL0057
NODE_1_length_3306_cov_120.625832_2 1 248 1 250 PB018707 Pfam-B_18707 Pfam-B 97 345 347 320.1 1.5e-95 NA NA
NODE_1_length_3306_cov_120.625832_4 4 127 1 128 PB004153 Pfam-B_4153 Pfam-B 55 190 270 75.1 6e-21 NA NA
NODE_1_length_3306_cov_120.625832_5 25 129 16 129 PB018707 Pfam-B_18707 Pfam-B 8 102 347 88.9 3.2e-25 NA NA

```

Family: **RHH_1 (PF01402)**

27 architectures 3978 sequences 2 interactions 1486 species 18 structures

Summary: Ribbon-helix-helix protein, copG family

Pfam includes annotations and additional family information from a range of different sources. These sources can be accessed via the tabs below.

No Wikipedia article **Pfam** Interpro

This tab holds the annotation information that is stored in the Pfam database. As we move to using Wikipedia as our main source of annotation, the contents of this tab will be gradually replaced by the Wikipedia tab.

Ribbon-helix-helix protein, copG family [Add annotation](#)


The structure of this protein repressor, which is the shortest reported to date and the first isolated from a plasmid, has a homodimeric ribbon-helix-helix arrangement [2]. The helix-turn-helix-like structure is involved in dimerisation and not DNA binding as might have been expected [2].

Literature references

1. Acebo P, Garcia de Lencastre M, Rivas G, Andreu JM, Espinosa M, del Solar G. Proteins 1998;32:248-261: Structural features of the plasmid pMV158-encoded transcriptional repressor CopG, a protein sharing similarities with both helix-turn-helix and beta-sheet DNA binding proteins. [PubMed:9714164](#)
2. Gomis-Roth FX, Sol M, Acebo P, Panagiotou A, Guasch A, Ertja R, Gonzalez A, Espinosa M, del Solar G, Coll M. EMBO J 1998;17:7404-7415: The structure of plasmid-encoded transcriptional repressor CopG unliganded and bound to its operator. [PubMed:9857196](#)

Clan

This family is a member of clan [Met_repress \(CL0057\)](#), which has a total of **21 members**.



Example structure: [RHH_1 \(PF01402\)](#) - Apo-NMR. View a different structure: [1Q5V](#)

Nearly all the results we obtain are hits against Pfam-B, but one is against Pfam-A. Let's take a look at Pfam domain PF01402 (RHH_1). Not an illuminating description, but a quick google search brings up the entry in the Pfam database (or just go to <http://pfam.sanger.ac.uk> and enter the ID in the search box).

Definitely some form of DNA binding protein, but not much information beyond that. Check out some of the other entries in Pfam just to get a feel for the sort of detail you can expect. The other Pfam-B matches do not tell you much that is useful.

4.5 Analysing the results in RAST

By now you should be able to see that analysing results for de novo assembled reads of any sort can be difficult and time-consuming. Bear in mind that we have only been faced with a single contig of 3kb. Quite often you may find yourself dealing with hundreds, if not thousands of contigs. Some will be a few 100kb long. Others may only be 200-300bp. How should we go about analysing these in a more efficient manner? There are a number of options here. For eukaryotes I would suggest looking at MAKER (<http://www.yandell-lab.org/software/maker.html>). For prokaryotes the situation is somewhat easier and we can use a web-based service known as RAST. This is not the only service (Xbase is another), but it is one of the most common.

RAST is a website where you upload the results of your de novo assembly and RAST will attempt to provide annotation in commonly used GFF and Genbank formats. This can be used to load up the annotation in Artemis or Apollo. Alternatively RAST has its own in-built viewer.

Task 15: Log in to RAST

Go to <http://rast.nmpdr.org/> Log-in with the details RAST provided you before you started this series of workshops. If you do not have one, you may need to wait several days for your login to be issued by RAST. Please skip ahead and come back to this section.



RAST Rapid Annotation using Subsystem Technology version 4.0
The NMPDR, SEED-based, prokaryotic genome annotation service.
For more information about The SEED please visit theSEED.org.

Info: You have been logged out.
Info: To monitor RAST's load and view other news and statistics for RAST and the SEED, please visit "[The Daily SEED](#)."

RAST (Rapid Annotation using Subsystem Technology) is a fully-automated service for annotating bacterial and archaeal genomes. It provides the number of more or less complete bacterial and archaeal genome sequences is constantly rising, the need for high quality automated we provide RAST as a free service to the community. It leverages the data and procedures established within the [SEED framework](#) to provide quality genome sequences AND the analysis of draft genomes. The service normally makes the annotated genome available within 12-24 hours.

Please note that while the SEED environment and SEED data structures (most prominently [FGfams](#)) are used to compute the automatic annotation the SEED. Once annotation is completed, genomes can be downloaded in a variety of formats or viewed online. The genome annotation provided to be able to contact you once the computation is finished and in case user intervention is required, we request that users register with email.

If you use our service, please cite:
The RAST Server: Rapid Annotations using Subsystems Technology.
Aziz RK, Bartels D, Best AA, DeJongh M, Disz T, Edwards RA, Formsma K, Gerdes S, Glass EM, Kubal M, Meyer F, Olsen GJ, Olson R, Ost Vorstein V, Wilke A, Zagnitko O.
BMC Genomics, 2008, [[article](#)]

This project has been funded in whole or in part with Federal funds from the National Institute of Allergy and Infectious Diseases, National Institutes of Health.

Login
Password

Task 16: Upload the assembled contigs and annotate using RAST

Click on Your jobs->Upload a Genome



RAST Rapid Annotation using Subsystem Technology version 4.0
The NMPDR, SEED-based, prokaryotic genome annotation service.
For more information about The SEED please visit theSEED.org.

[Home](#) [Your Jobs](#)

Upload a Genome

A prokaryotic genome in one or more contigs should be uploaded in either a single [FASTA](#) format file or in a Genbank format file. Our pipeline will use the taxonomy identifier as a handle for the genome. Therefore it is important to include the genus, species and strain in the following upload workflow.

Please note, that only if you submit all relevant contigs (i.e. all chromosomes, if more than one, and all plasmids) that comprise the genomic information of your organism of interest in one job. Features like *Metabolic pathways* are not included.

Confidentiality information: Data entered into the server will not be used for any purposes or in fact integrated into the main SEED environment, it will remain on this server for 120 days or until deleted by the submitter.

If you use the results of this annotation in your work, please cite:
The RAST Server: Rapid Annotations using Subsystems Technology.
Aziz RK, Bartels D, Best AA, DeJongh M, Disz T, Edwards RA, Formsma K, Gerdes S, Glass EM, Kubal M, Meyer F, Olsen GJ, Olson R, Osterman AL, Overbeek RA, McNeil LK, Paarmann D, Paczian T, Parrello B, Pritchett LC, Zagnitko O.
BMC Genomics, 2008, [[article](#)]

File formats: You can either use [FASTA](#) or Genbank format.

- If in doubt about FASTA, [this service](#) allows conversion into FASTA format.
- Due to limits on identifier sizes imposed by some of the third-party bioinformatics tools that RAST uses, we limit the size of contig identifiers to 70 characters or fewer.
- If you use Genbank, you have the option of preserving the gene calls in the options block below. By default, genes will be recalled.

Please note: This service is intended for complete or nearly complete prokaryotic genomes. For now we are not able to reliably process sequence data of very small size, like small plasmid, phages or fragments.

File Upload:
Sequences File

Upload the contigs.fa file obtained by the de novo assembly of unmapped reads.

Review genome data

We have analyzed your upload and have computed the following information.

Contig statistics

Statistic	As uploaded	After splitting into scaffolds
Sequence size	3326	3326
Number of contigs	1	1
GC content (%)	43.4	43.4
Shortest contig size	3326	3326
Median sequence size	3326	3326
Mean sequence size	3326.0	3326.0
Longest contig size	3326	3326

Please enter or verify the following information about this organism:

Required information:

Taxonomy ID: (leave blank if NCBI Taxonomy ID unknown)

Find the taxonomy id for your organism by searching for its name in the [NCBI taxonomy browser](#).

Taxonomy string:

Domain: ☒ Bacteria ☐ Archaea ☐ Virus

Genus:

Species:

Strain:

Genetic Code: ☒ 11 (Archaea, most Bacteria, most Virii, and some Mitochondria)
☐ 4 (Mycoplasmaea, Spiroplasmaea, Ureoplasmaea, and Fungal Mitochondria)

We know this is an *E.coli* genome so we can enter 562 as the Taxonomy ID and click on 'Look up taxonomy at NCBI'. If you're dealing with a different organism, be sure to change this number. RAST will automatically split any scaffolds (i.e. contigs with bits missing in the middle – denoted by Ns).

Upload a Genome

Complete Upload

By answering the following questions you will help us improve our ability to track problems in processing your genome:

Optional information:

Sequencing Method ☐ Sanger ☐ Mix of Sanger and Pyrosequencing ☐ Pyrosequencing ☒ other

Coverage

Number of contigs

Average Read Length (leave blank if unknown)

Please consider the following options for the RAST annotation pipeline:

RAST Annotation Settings:

Select gene caller	RAST	Please select which type of gene calling you would like RAST to perform. Note that
Select FIGfam version for this run	Release45	Choose the version of FIGfams to be used to process this genome.
Automatically fix errors?	☒ Yes	The automatic annotation process may run into problems, such as gene candidates requires deleting some gene candidates), please check this box.
Fix frameshifts?	☐ Yes	If you wish for the pipeline to fix frameshifts, check this option. Otherwise frameshifts
Build metabolic model?	☒ Yes	If you wish RAST to build a metabolic model for this genome, check this option.
Backfill gaps?	☒ Yes	If you wish for the pipeline to blast large gaps for missing genes, check this option
Turn on debug?	☐ Yes	If you wish debug statements to be printed for this job, check this box.
Set verbose level	0	Set this to the verbosity level of choice for error messages.
Disable replication	☐ Yes	Even if this job is identical to a previous job, run it from scratch.

Replicate the settings above and click on 'Finish the upload'.

Your job may take several hours to run. In the meantime, proceed to the next workshop and come back to this later.

Once complete, RAST should email you a message. You can then view the results or download them in standardized formats (e.g. GFF3, Genbank, EMBL etc).



RAST

Rapid Annotation using Subsystem Technology

version 4.0

The NMPDR, SEED-based, prokaryotic genome annotation service.
For more information about The SEED please visit theSEED.org.

[Home](#) | [Your Jobs](#) | [Manage Job #41704](#)

Job Details #41704

» [Browse annotated genome in SEED Viewer](#)

» Available downloads for this job: GFF3 ▼ Download Update download files

» [Share this genome with selected users](#)

» [Back to the Jobs Overview](#)

 Genome Upload has been successfully completed.

Genome ID - Name:	562.698 - Escherichia coli
Job:	#41704
User:	khp204
Date:	Tue Jan 31 08:03:05 2012
Sequencing method:	other
Coverage:	gt8
Number of contigs:	unknown
Read length:	
Genetic code:	11
Include into SEED:	no
Preserve gene calls:	no
Automatically fix errors:	yes
Fix frameshifts:	no
Backfill gaps:	yes

 Rapid Propagation has been successfully completed.

 Quality Check has been successfully completed.

For detailed explanations of the terms used in our quality report, please refer to [our wiki](#).

Number of features:	4
Number of warnings:	0
Number of fatal problems:	0

 Quality Revision has been successfully completed.

If you download the results in GFF3 format and open in a text-editor you can see the annotations. Again they are not particularly enlightening as such, but we can see that this plasmid appears to consist of 4 proteins.

```
##gff-version 3
NODE_1_length_3306_cov_120.625832      FIG      CDS      1162    1356    .      -      1      ID=fig|562.698.peg.1;Name=probable RNAI modulator protein
NODE_1_length_3306_cov_120.625832      FIG      CDS      1313    1951    .      +      2      ID=fig|562.698.peg.2;Name=Plasmid mobilization protein A
NODE_1_length_3306_cov_120.625832      FIG      CDS      2186    2566    .      -      2      ID=fig|562.698.peg.3;Name=14.5 kDa protein
NODE_1_length_3306_cov_120.625832      FIG      CDS      2563    3204    .      -      1      ID=fig|562.698.peg.4;Name=hypothetical protein
```

How does this compare to the original paper which describes this plasmid? We can also look at other features of RAST such as the SEED Viewer. Here we can see each predicted gene along with annotation and further information as to how the evidence was obtained if you click on the Feature ID.

Browse Genome: [Escherichia coli \(562.698\)](#)

Location **Focus** Upload List

contig NODE_1_length_3306_cov_120.625832 (3,326 bp)

start base

window 16,000 bp

Color features by focus

<== draw ==>

export table clear all filters

display 15 items per page
displaying 1 - 4 of 4

Feature ID	Type	Contig	Start	Stop	Length (bp)	Function	Subsystems	Region
fig 562.698.peg.1	CDS	NODE_1_length_3306_cov_120.625832	1356	1162	195	probable RNAI modulator protein	- none -	show
fig 562.698.peg.2	CDS	NODE_1_length_3306_cov_120.625832	1313	1951	639	Plasmid mobilization protein A	- none -	show
fig 562.698.peg.3	CDS	NODE_1_length_3306_cov_120.625832	2566	2186	381	14.5 kDa protein	- none -	show
fig 562.698.peg.4	CDS	NODE_1_length_3306_cov_120.625832	3204	2563	642	hypothetical protein	- none -	show

displaying 1 - 4 of 4

Browse the rest of the RAST server and get a feel for the possibilities the platform may offer you.

When you're ready, move on to (or back to) the de novo assembly part of the workshop!

2013 Workshop on Genomics, Cesky Krumlov

Part 5 Short read genomics: De-novo assembly

Instructors:

- Konrad Paszkiewicz k.h.paszkiewicz at exeter ac uk

Objectives:

By the end of the workshop you will be expected to be able to:

- Perform QC and adaptor-trim Illumina reads
- Assemble these reads de novo using Velvet and VelvetOptimiser
- Generate summary statistics for the assembly
- Identify open reading frames within the assembly
- Search for matches within the Swissprot database via BLAST and against the Pfam database
- Visualize species distribution of potential matches

5.1 Introduction

In this section of the workshop we will continue the analysis of a strain of *E.coli* which is involved in urinary tract infections. In the previous section we extracted those reads which did not map to the reference genome and assembled them. However, it is often necessary to be able to perform a de novo assembly of a genome. In this case, rather than doing any remapping, we will start with the filtered reads we obtained in part 3 of the workshop.

Although it is not strictly necessary for this particular example (because there appears to be very little new material), we will run through the procedure here so that you are able to do this in the next section.

To do this we will use a program called VelvetOptimiser to try to get the best possible assembly for a given genome. We will then generate assembly statistics and then produce some annotation via Pfam and BLAST.

5.2 Task 1 Preparation

Ensure you are in the ~/genomics_tutorial/strain1 directory. We will move the denovo_analysis directory into the remapping_to_reference directory so we don't end up confusing the current de novo assembly-from-scratch, with the previous one using unmapped reads.

```
mv denovo_assembly/ remapping_to_reference/
```

Let's create a new directory to store this denovo-assembly-from-scratch. We'll call it complete_denovo_assembly

```
mkdir complete_denovo_assembly/
```

Listing the contents of the ~/genomics_tutorial/strain1 directory you should now have:

```
ubuntu@ip-10-28-24-176:~/genomics_tutorial/strain1$ ls
complete_denovo_assembly  illumina_reads  remapping_to_reference
ubuntu@ip-10-28-24-176:~/genomics_tutorial/strain1$
```

```
cd complete_denovo_assembly/
```

Assembly theory

Velvet is an assembly algorithm which requires a number of parameters. The three most important are k-mer length, expected coverage and coverage cutoff.

1. **K-mer length.** Rather than store all reads individually which would be unfeasible for Illumina type datasets, Velvet converts each read to a series of kmers and stores each kmer once, along with information about how often it occurs and which other kmers it links to. A short k-mer length (e.g. 21) reduces the chance that data will be missed from an assembly (e.g. due to reads being shorter than the k-mer length, but generally results in higher memory consumption and shorter contigs.
2. **Expected coverage.** This is rough indication of how much sequence information there is relative to the genome size. It is used to try to resolve repetitive regions if paired-end information is missing or insufficient.
3. **Coverage cutoff** Relates to the cutoff Velvet should use to decide that a particular k-mer is noise (i.e. error-laden or low-level contamination) and should be excluded from the assembly.




Generally it is necessary to try a large combination of parameters to ensure that you obtain the 'best' possible assembly for a given dataset. What 'best' actually means in the context of genome assembly is ill-defined. For a genomic assembly you want to try to obtain the lowest number of contigs, with the longest length, with the fewest errors. However, although numbers of contigs and longest lengths are easy to evaluate, it is extremely difficult to know what is or isn't an error when sequencing a genome for the first time.

Description of k-mers:

What are they? Let's say you have a single read:


AACTAACGACGCGCATCAAAA

The set of k-mers obtained from this read with length 6 (i.e. 6-mers) would be obtained by taking the first six bases, then moving the window along one base, taking the next 6 bases and so-on until the end of the read. E.g:



You may well ask, “So what? How does that help”? For a single read, it really doesn't help. However let's say that you have another read which is identical except for a single base:

AACTAACGAC G GCGATCAAAA
 ACTAACGAC T TCGATCAAAA



Rather than represent both reads separately, we need only store the k-mers which differ


frequency plot of contig (node) coverage, weighted by length. In the example below you can see that contigs with a coverage below 7x or 8x occur very infrequently. As such it is probably a good idea exclude those contigs which have coverage less than this – they are likely to be errors.

Description of expected coverage:

In the example below you can see a stretch of DNA with many reads mapping to it. There are two repetitive regions A1 and A2 which have identical sequence. If we try to assemble the reads without any knowledge of the true DNA sequence, we will end up with an assembly that is split into two or more contigs rather than one.

One contig will contain all the reads which did not fall into A1 and A2. The other will contain reads from both A1 and A2. As such the coverage of the repetitive contig will be twice as high as that of the non-repetitive contig.

If we had 5 repeats we would expect 5x more coverage relative to the non-repetitive contig. As such, provided we know what level of coverage we expect for a given set of data, we can use this information to try and resolve the number of repeats we expect.



Task 2: Using VelvetOptimiser

VelvetOptimiser cycles Velvet through various k-mer, expected coverage and coverage cutoff values to see which one gives the 'best' result. Best is defined by the user as either N50 length, length of longest contig, total number of bases in contigs etc. It is distributed with Velvet and is a 'contributed' script (i.e. not written by the original author of Velvet, but thought to be so useful as to be worth bundling with the Velvet release).

Type:

VelvetOptimiser.pl

```
*****
VelvetOptimiser.pl Version 2.2.0

Simon Gladman - CSIRO 2009

*****
Number of CPUs available: 2
Current free RAM: 4.986GB
Velvet OMP compiler setting: 1
Usage: /usr/local/bin/VelvetOptimiser.pl [options] -f 'velveth input line'
--help                This help.
--v|verbose+          Verbose logging, includes all velvet output in the logfile. (default '0').
--s|hashs=i           The starting (lower) hash value (default '19').
--e|hashe=i           The end (higher) hash value (default '101').
--x|step=i            The step in hash search.. min 2, no odd numbers (default '2').
--f|velvethfiles=s     The file section of the velveth command line. (default '0').
--a|amosfile!         Turn on velvet's read tracking and amos file output. (default '0').
--o|velvetgoptions=s  Extra velvetg options to pass through. eg. -long_mult_cutoff -max_coverage etc (default '').
--t|threads=i         The maximum number of simultaneous velvet instances to run. (default '2').
--g|genomesize=f       The approximate size of the genome to be assembled in megabases.
                        Only used in memory use estimation. If not specified, memory use estimation
                        will not occur. If memory use is estimated, the results are shown and then program exits. (default '0').
--k|optFuncKmer=s      The optimisation function used for k-mer choice. (default 'n50').
--c|optFuncCov=s       The optimisation function used for cov cutoff optimisation. (default 'Lbp').
--p|prefix=s          The prefix for the output filenames, the default is the date and time in the format DD-MM-YYYY-HH-MM. (default 'auto').
--d|dir_final=s        The name of the directory to put the final output into. (default '.').

Advanced!: Changing the optimisation function(s)

Velvet optimiser assembly optimisation function can be built from the following variables.
LNbp = The total number of Ns in large contigs
Lbp = The total number of base pairs in large contigs
Lcon = The number of large contigs
max = The length of the longest contig
n50 = The n50
ncon = The total number of contigs
tbp = The total number of basepairs in contigs

Examples are:
'Lbp' = Just the total basepairs in contigs longer than 1kb
'n50*Lcon' = The n50 times the number of long contigs.
'n50*Lcon/tbp+log(Lbp)' = The n50 times the number of long contigs divided
                        by the total bases in all contigs plus the log of the number of bases
                        in long contigs.
```

You should see the options list for Velvet. Note that there are a number of options which we need to set. The first two are the upper and lower bounds for the k-mer search. The lower bound should never be lower than 19 as memory requirements tend to spiral, and the upper bound should not be longer than the read length (i.e. 40) – otherwise no k-mers can be formed at all!

The next parameter we need to feed is the -f parameter which just specifies whether the

reads are paired or unpaired and what format they are in. We can also set -t if we have a large memory machine to do multiple assemblies at the same time to speed things up. If we want to estimate how much memory an assembly will require, we can set the -g flag and indicate in megabases how large we expect the genome to be. We can also set the optimisation functions (i.e. how VelvetOptimiser decides what the 'best' assembly is), both for k-mer size and for coverage cutoff, but we'll leave those alone for now. Feel free to experiment with these at a later date though!

Before we can run VelvetOptimiser.pl we need to interleave the filtered FASTQ files.

Type (all on one line):

```
perl /usr/local/velvet/shuffleSequences_fastq.pl  
~/genomics_tutorial/strain1/illumina_reads/strain1_read1.filtered.fastq  
~/genomics_tutorial/strain1/illumina_reads/strain1_read2.filtered.fastq  
strain1_interleaved.fastq
```

Now we can run VelvetOptimiser. Let's try it first with the -g flag set to 5Mb so that we can estimate the maximum RAM Velvet will use. Type: (all on one line)

```
VelvetOptimiser.pl -s 19 -e 39 -f '-shortPaired -fastq strain1_interleaved.fastq' -t 1 -  
g 5
```

```
ubuntu@ip-10-28-24-176:~/genomics_tutorial/strain1/complete_denovo_assembly$ VelvetOptimiser.pl  
ved.fastq' -t 1 -g 5
```

```
*****
```

```
VelvetOptimiser.pl Version 2.2.0
```

```
Simon Gladman - CSIRO 2009
```

```
*****
```

```
Number of CPUs available: 2
```

```
Current free RAM: 4.986GB
```

```
Velvet OMP compiler setting: 1
```

```
Feb 2 15:41:47 Starting to check input parameters.
```

```
Velveth parameter string OK.
```

```
Feb 2 15:41:47 Finished checking input parameters.
```

```
Logfile name: 02-02-2012-15-41-47_Logfile.txt
```

```
Memory use estimation only! Script will terminate after showing results.
```

```
Velvet details:
```

```
Velvet version: 1.2.02
```

```
Compiled categories: 2
```

```
Compiled max kmer length: 101
```

```
Maximum number of velvetinstances to run: 1
```

```
Will run velvet optimiser with the following paramters:
```

```
Velveth parameter string:
```

```
-shortPaired -fastq strain1_interleaved.fastq
```

```
Velveth start hash values: 19
```

```
Velveth end hash value: 39
```

```
Velveth hash step value: 2
```

```
Read tracking for final assembly off.
```

```
File: strain1_interleaved.fastq has 6759881 reads of length 40
```

```
Total reads: 6.8 million. Avg length: 40.0
```

```
Memory use estimated to be: 1.6GB for 1 threads.
```

```
You should have enough memory to complete this job. (Though this estimate is no guarantee..)
```

So the results indicate that we require around 1.6Gb of RAM to do one assembly at a time. We have a maximum of 4.9Gb of RAM free so although we could try two threads at a time; let's leave it at 1 thread. If the memory requirements are exceeded you may find that your instance stops responding and you'll need to restart VelvetOptimiser with a higher starting kmer (e.g. 25 instead of 19).

```
VelvetOptimiser.pl -s 19 -e 39 -f '-shortPaired -fastq strain1_interleaved.fastq' -t 1
```

```
Velvetg parameter string: auto_data_21 -clean yes -exp_cov 9 -cov_cutoff 5.2522
188895872
Assembly directory: /home/ubuntu/genomics_tutorial/strain1/complete_denovo_assembly/auto_data_21
Velvet hash value: 21
Roadmap file size: 268940253
Total number of contigs: 655
n50: 228461
length of longest contig: 622788
Total bases in contigs: 5151239
Number of contigs > 1k: 75
Total bases in contigs > 1k: 5087404
Paired Library insert stats:
Paired-end library 1 has length: 449, sample standard deviation: 40
Paired-end library 1 has length: 450, sample standard deviation: 40
*****
```

```
Assembly output files are in the following directory:
/home/ubuntu/genomics_tutorial/strain1/complete_denovo_assembly/auto_data_21
```

Once completed after about 30 minutes you will see something similar to the above. Here we can see that a k-mer length of 21 was found to give the highest N50 value (see http://en.wikipedia.org/wiki/N50_statistic for details). A coverage cutoff of 5.25 and expected coverage of 9 were the values which gave the largest number of base pairs in contigs larger than 1kb.

In addition to the statistics, we also get a report of the estimated insert size and standard deviation. It is important to check these (they should be around 300-600 bp) to make sure the library prep went well.

This is a good optimised assembly. It has a high N50 (228,461bp) and although there are a total of 655 contigs, most of the genome seems to be assembled into just 75 contigs.

Check out the contents of the auto_data_21 directory. You'll see that the results are in a file called 'contigs.fa'. The files have exactly the same format as those seen in the previous workshop.

A good check at this point is to map the original reads back to the contigs.fa file and check that all positions are covered by reads. Amazingly it is actually possible for de-novo assemblers to generate contigs to which the original reads will not map.

5.5 Task 3: Map reads back to assembly

Here we will use BWA again to index the contigs.fa file and remap the reads. This is almost identical to the procedure we followed in Part 3, the only difference is that instead of aligning to the reference genome, we are aligning to the contigs.fa file.

Make sure you are in the following directory.

```
~/genomics_tutorial/strain1/complete_denovo_assembly/auto_data_21
```

Let's start by indexing the contigs.fa file-management. Type:

bwa index contigs.fa

```
ubuntu@ip-10-28-24-176:~/genomics_tutorial/strain1/complete_denovo_assembly/auto_data_21$ ls
Graph  Log      Sequences  contigs.fa.amb  contigs.fa.bwt  contigs.fa.rbwt  contigs.fa.rsa  stats.txt
Graph2 PreGraph  contigs.fa  contigs.fa.ann  contigs.fa.pac  contigs.fa.rpac  contigs.fa.sa
ubuntu@ip-10-28-24-176:~/genomics_tutorial/strain1/complete_denovo_assembly/auto_data_21$
```

Once complete we can start to align read 1 back to the contigs. Type (all on one line):

```
bwa aln -t 8 -q 20 -f bwa_read1.aln contigs.fa  
~/genomics_tutorial/strain1/illumina_reads/strain1_read1.filtered.fastq  
tq
```

We won't do read 2 for the sake of convenience and because BWA tends to get horribly confused if reads map to different contigs. (N.B. If we wanted to use paired end reads we could use the Bowtie (v1) aligner instead).

We can then incorporate the paired-end information by running bwa sampe

```
bwa samse -f alignment.sam contigs.fa bwa_read1.aln  
../../illumina_reads/strain1_read1.filtered.fastq
```

This may take longer than the original mapping to the reference.

Once complete we can convert the SAM file to a BAM file:

```
samtools view -bS -T contigs.fa alignment.sam > alignment.bam
```

And then we can sort the BAM file:

```
samtools sort alignment.bam alignment.sorted
```

Once completed, we can index the BAM file:

```
samtools index alignment.sorted.bam
```

We can then (at last!) obtain some basic summary statistics using the samtools flagstat command:

```
samtools flagstat alignment.sorted.bam
```

```
3247450 + 0 in total (QC-passed reads + QC-failed reads)
0 + 0 duplicates
2799941 + 0 mapped (86.22%:-nan%)
0 + 0 paired in sequencing
0 + 0 read1
0 + 0 read2
0 + 0 properly paired (-nan%:-nan%)
0 + 0 with itself and mate mapped
0 + 0 singletons (-nan%:-nan%)
0 + 0 with mate mapped to a different chr
0 + 0 with mate mapped to a different chr (mapQ>=5)
```

We can see here that almost 14% of reads did not map back to the contigs.

Task 4: View assembly in IGV

As in part 3, load up Firefox (make sure it is within your Amazon instance and not on your desktop). Google for the IGV viewer and load up the 2Gb version of the viewer. (On 32-bit windows machines you will only be able to load the 1.6Gb version).

Click File->Import Genome. We are going to import the contigs we have assembled as the reference. Unlike the reference genome though, we have no annotation available.

Make sure you select the contigs.fa file for the complete de novo assembly (not the unmapped reads assembly).

The screenshot shows a configuration window for a de novo assembly project. It contains several input fields and a checkbox. The 'ID' field is labeled 'Strain1Assembly' with a note '(unique id, e.g. hg18)'. The 'Name' field is labeled 'Strain 1 Assembly'. There is a checkbox labeled 'Fasta file is a directory'. The 'Fasta file' field is labeled with an asterisk and contains the path '/home/ubuntu/genomics_tutorial/strain1/complete_denovo_assembly/auto_data_21/contigs.fa'. To the right of this field is a button with three dots. Below it are empty fields for 'Cytoband ...', 'Gene file', and 'Alias file', each with a corresponding button with three dots. A note says 'Supply a sequence URL if defining a web-hosted genome (optional, not common). See user guide for more details.' Below this is a 'Sequence URI' field. A paragraph of text explains that the sequence file is required and can be a FASTA file, a directory of FASTA files, or a zip of FASTA files. It also mentions optional cytoband and gene track files. At the bottom are 'Save' and 'Cancel' buttons.

ID * (unique id, e.g. hg18)

Name *

☐ Fasta file is a directory

Fasta file * ...

Cytoband

Gene file ...


Alias file ...

Supply a sequence URL if defining a web-hosted genome (optional, not common). See user guide for more details.

Sequence URI

** required.
The sequence file (required) can be a FASTA file, a directory of FASTA files, or a zip of FASTA files. Optionally, specify a cytoband file to display the chromosome ideogram and an annotation file to display the gene track. See the documentation for descriptions of supported annotation formats.*

Once saved, click on File->Load From File... select the alignment.calmd.rmdup.bam file. Again, make sure you load the file in the complete_denovo_assembly directory.



Once loaded, explore some of the contigs in IGV. Can you find examples of some which have no reads mapping back to the assembly? Do you notice anything they tend to have in common?



Annotation of de-novo assembled contigs

We will now annotate the contigs using BLAST, Pfam and RAST as with the unmapped contigs.

Task 5: Obtain open reading frames

As before, we'll call open reading frames within the de-novo assembly. Again, we will use codon table 11 which defines the bacterial codon usage table (<http://www.ncbi.nlm.nih.gov/Taxonomy/Utils/wprintgc.cgi>) and state that the sequences we are dealing with are not circular. We will also restrict the ORFs to just those sequences longer than 300 nucleotides (i.e. 100 amino acids) We will store the results in file contigs.orf.fa.

```
getorf -table 11 -circular N -minsize 300 -sequence contigs.fa -outseq contigs.orf.fa
```

The following two tasks are optional. Be warned - the BLAST searches and RAST will take several days! I recommend you skip these and proceed to Task 9.

Task 6 (Optional): Search open reading frames against NCBI non-redundant database

We can also search these open reading frames against the NCBI non-redundant database.

```
blastall -d nr -p blastp -i contigs.orf.fa -o contigs.orf.blastp -e 1e-06 -a 4 -I T -b 10 -v 10
```

Task 7 (Optional): Search contigs against NCBI non-redundant database

The following command executes a nucleotide BLAST search (blastn) of the sequences in the contigs.fa file against the non-redundant database. Again we restrict ourselves to 10 results per hit and an e-value cutoff of 1e-06.

```
blastall -d nt -p blastn -i contigs.fa -o contigs.fa.blastn -e 1e-06 -b 10 -v 10 -I T -a 4
```

Task 8 (Optional): Run the contigs through the RAST server and import the resulting GFF annotation into IGV

Task 9: Run open reading frames through Pfam

As with the unmapped reads we will search the open reading frames against the Pfam HMM database of protein families. Later on we will be able to use these results to identify Pfam domains which are unique to a particular strain.

```
pfam_scan.pl -fasta contigs.orf.fa -dir /usr/local/PfamScan/ -outfile contigs.orf.pfam -cpu 4 -pfamB -as
```

This will take around 5 hours so it is recommended that you leave this running while continuing with the rest of the tutorial. If it is still running when you finish your session for today, leave your instance running overnight, but please be sure to turn it off in the morning!

2013 Workshop on Genomics, Cesky Krumlov

Part 6 **Short read genomics:** **Comparison of results between** **different strains**

Instructors:

- Konrad Paszkiewicz k.h.paszkiewicz at exeter ac uk

Objectives:

In this workshop you will:

By the end of the workshop you will be expected to be able to:

- Run parts 1-3 of the workshop on 2 new datasets
- Use pre-prepared scripts to compare SNPs and Indels between strains
- Generate pseudo-sequences based on synonymous SNPs
- Draw simple trees to illustrate the likely evolutionary relationship between strains
- Compare Pfam matches between strains

Projects!

In the previous sections you have been taken through the steps required to:

1. QC and filter Illumina data
2. Remap Illumina short-read data to a reference sequence
3. View the results in IGV
4. Identify SNPs and Indels in an automated fashion using samtools and bcftools
5. Determine whether SNPs result in synonymous or non-synonymous changes in the corresponding amino acid
6. Extract unmapped reads
7. Assemble unmapped reads and obtain assembly statistics
8. Annotate unmapped reads using Pfam, RAST and/or BLAST
9. Assemble a bacterial genome de-novo using Velvet
10. Obtain assembly statistics
11. Annotate as per the unmapped reads (where computationally feasible).

Now we want you to do the same for another two strains of *E.coli* urinary tract infections which have been isolated. For each strain, make a list of:

1. SNPs, Indels and their effects (from the remapping)
2. Missing genes (from the remapping)
3. Novel plasmids and/or genes (Pfam domains are the easiest way to do this via denovo assembly of unmapped reads)

Once completed, see if you can predict what the phenotype of these bacteria might be. Then proceed to the final part of the tutorial where we will compare the results from all of these strains.

N.B.

It is recommended that you follow the same directory naming convention we have followed here (i.e. one for remapping to the reference, another for assembly of unmapped reads and a final one for the denovo assembly).

You will find the data for the other two strains in the ~/genomics_tutorial/strain2 and ~/genomics_tutorial/strain3. These tasks may take you several days. However, remember that all of the basic procedures are detailed in the previous sections – only the input FASTQ files will have changed. Feel free to refer to these previous tasks to remind yourself of the commands and parameters. By all means feel free to play around with different parameters if you wish, although remember that the results may differ from those you see here.

Just to give you some guidance: You should find that strain 2 yields many more SNPs than strain 1. Strain 2 should be missing a Beta lactamase and Strain 3 Type III secretory pathway and flagellar proteins. You should also find a different plasmid. You may find that some scripts and programs run more slowly because of these extra differences. Also, if you find the de novo assembly process causes your NX session to end, the chances are that Velvet has caused your instance to run out of memory. If this happens, increase the minimum k-mer size in VelvetOptimiser.

Comparing variants between several samples and a reference genome:

Here we will use a script to compare the variants called in each sample. Ensure you are in the ~/genomics_tutorial directory

First of all, let's make a directory to store the results of the comparison:

```
mkdir snp_comparison/
```

```
cd snp_comparison/
```

```
cp ~/genomics_tutorial/strain1/remapping_to_reference/out.snps.vcf4 ./strain1
```

```
cp ~/genomics_tutorial/strain2/remapping_to_reference/out.snps.vcf4 ./strain2
```

```
cp ~/genomics_tutorial/strain3/remapping_to_reference/out.snps.vcf4 ./strain3
```

Note that the last two copy commands may require modification depending on where you have saved the variant call results.

```
perl /usr/local/scripts/snp_comparator.pl 10  
~/genomics_tutorial/reference_sequence/Ecoli_UTI89.fna  
~/genomics_tutorial/reference_sequence/Ecoli_UTI89.gff strain1 strain2 strain3 >  
snp_comparison.txt
```

Looking at the snp_comparison.txt file (either in a text editor, or in a spreadsheet):

```
## Table of SNP and Indel occurrences between these samples. Note that any comma-separated values (e.g. A,C indicate potential heterozygosity and/or sample heterogeneity)
Chrom  Pos    Ref    strain1 strain2 strain3 Gene description      Status
CP000243.1 1000078 C      C      G      C      UTI89_C1007 alkanesulfonate monooxygenase ,silent tcg -> tcC;
CP000243.1 1000196 G      G      A      G      UTI89_C1007 alkanesulfonate monooxygenase ,non-silent gcc -> gTc;
CP000243.1 1000403 G      G      T      G      UTI89_C1007 alkanesulfonate monooxygenase ,non-silent tca -> tAa;
CP000243.1 1000438 G      G      T      G      UTI89_C1007 alkanesulfonate monooxygenase ,non-silent cac -> caA;
```

Here we can see the chromosome ID, the position in bp, the reference base and the base at each position as well as the gene (if any) the variant occurs in as well as the effect (silent, non-silent or indel).

Obtaining a phylogeny based on synonymous SNPs only:

How are the three strains related on the basis of these variants? We can ask a number of questions, but if we are looking at the long-term evolutionary history of the strains we should only look at synonymous (i.e. silent) mutations as these should not confer a significant selective advantage to any strain. Using the data snp_comparison.txt file we can form 'pseudo-sequences' using the script in /usr/local/scripts/snp2tree_fullsequence.pl. These are concatenated bases consisting of only those positions which are silent across all strains. It is essentially the same as turning each column of each strain in the snp_comparison.txt file into a FASTA entry.

```
perl /usr/local/scripts/snp2tree_fullsequence.pl snp_comparison.txt >
synonymous_tree.fasta
```

Examine the contents of the tree.fasta file. We can then treat this file as an alignment (since each base in each sequence is at the same position on the chromosome) and pass it to a phylogeny program called FastTree. FastTree will take an input alignment and output a Newick formatted tree (http://en.wikipedia.org/wiki/Newick_format).

```
FastTree -nt -gtr < synonymous_tree.fasta > synonymous_tree.newick
```

Now we can visually represent this tree by using the newick2pdf program to create a visual representation of the tree based on the synonymous_tree.newick file.

```
newicktopdf synonymous_tree.newick
```

We can then view the resulting PDF file using the evince program or Adobe Acrobat.

```
evince synonymous_tree.pdf
```

Advanced task (optional):

Copy the snp2tree_fullsequence.pl script to this directory and modify it so that it selects positions containing only non-silent mutations (not indels as these modify the alignment). Generate a new alignment and compare the resulting tree against the silent mutations.

Advanced task (optional):

Try to repeat one of the analyses with the new BreSeq pipeline (<http://barricklab.org/twiki/pub/Lab/ToolsBacterialGenomeResequencing/documentation/introduction.html>). This pipeline automatically aligns your reads to a reference genome and calls SNPs along with generation of pretty HTML files detailing the likely effects of mutations. Compare the SNP calls you get with BreSeq with the SNPs you get using the methods used here. You will need to use the GenBank (gbk) files in the reference_sequence directory.

This is an example of the command you may want to execute.

```
breseq -r /home/ubuntu/reference_sequence/Ecoli_UTI89_mainchr.gbk -r
/home/ubuntu/reference_sequence/Ecoli_UTI89_plasmid.gbk
/home/ubuntu/strain1/illumina_reads/strain1_read1.fastq
/home/ubuntu/strain1/illumina_reads/strain1_read2.fastq
```

As always, keep your results organized – put the breseq results in a separate directory.

Comparing Pfam domains found in each strain:

Here we will use a script to compare the various Pfam domains found in each sample. Ensure you are in the ~/genomics_tutorial directory

First of all, lets make a directory to store the results of the comparison:

```
mkdir pfam_comparison/
```

```
cd pfam_comparison/
```

```
cp  
~/genomics_tutorial/strain1/complete_denovo_assembly/contigs.orf.pfam ./strain1
```

```
cp ~/genomics_tutorial/strain2/  
complete_denovo_assembly/contigs.orf.pfam ./strain2
```

```
cp ~/genomics_tutorial/strain3/  
complete_denovo_assembly/contigs.orf.pfam ./strain3
```

Note that the last two copy commands may require modification depending on where you have saved the Pfam search results.

```
perl /usr/local/scripts/pfam_comparison.pl strain1 strain2 strain3 >  
pfam_comparison.txt
```

Examining the pfam_comparison.txt file you should see something similar to:

PF02705	K_trans	strain2.pfam, strain1.pfam,
PF02706	Wzz	strain2.pfam, strain1.pfam,
PF02729	OTCace_N	strain2.pfam, strain1.pfam,
PF02730	AFOR_N	strain2.pfam, strain1.pfam,
PF02733	Dak1	strain2.pfam, strain1.pfam,
PF02734	Dak2	strain2.pfam, strain1.pfam,
PF02737	3HCDH_N	strain2.pfam, strain1.pfam,
PF02738	Ald_Xan_dh_C2	strain2.pfam, strain1.pfam,
PF02739	5_3_exonuc_N	strain2.pfam, strain1.pfam,
PF02742	Fe_dep_repr_C	strain1.pfam,
PF02744	GalP_UDP_tr_C	strain2.pfam, strain1.pfam,

Search the Pfam database to see what some of these differences are (<http://pfam.sanger.ac.uk>)

Concluding remarks:

Well done! If you have reached this far, you deserve a round of applause. You have completed some of the most common tasks in short-read sequencing. You can use the same machine and the same scripts to perform analysis of any short-read dataset! All you need to do is transfer the FASTQ files to the server – you can either do this via Firefox if using the Exeter Sequencing Service (just click on the link in the email), or if you have them on your personal desktop you can use WinSCP (windows) , Fugu or Cyberduck (Mac OSX) or any other SFTP program.

A tutorial can be found at http://www.siteground.com/tutorials/ssh/ssh_winscp.htm