Published in final edited form as:

Nat Methods. 2010 February ; 7(2): 105–106. doi:10.1038/nmeth0210-105.

Mind the gaps

Steven L Salzberg

Center for Bioinformatics and Computational Biology and Department of Computer Science,
University of Maryland, College Park, Maryland, USA. salzberg@umd.edu

Abstract

By grouping short reads derived from the same long genomic fragment, the reads can easily be assembled into fragments that approach the length of capillary sequencing reads.

Amid all the excitement about next-generation sequencing, scientists often neglect to mention the problems that are caused by short read lengths. Genome assemblies produced from short reads are far more fragmented than those produced from long reads, with many more gaps and with relatively poor long-range linking information. Fortunately, technology to produce long reads is still being pursued vigorously in some laboratories. In this issue, a group led by Jay Shendure¹ describes a method to create ‘reads’ as long as 700 base pairs (bp) starting with pairs of reads that are just 20 bp and 76 bp.

For more than a decade after its introduction, automated sequencing steadily produced longer reads until, by the early 2000s, the ABI 3730 line of capillary sequencers routinely produced read lengths over 800 bp. As read lengths increased, costs came down, and genome assembly became easier. It was also easier to capture full-length transcripts for protein-coding genes with the longer reads, which in turn made gene discovery easier and more reliable.

Then along came next-generation sequencers, with read lengths of 25–100 bp, and although costs per megabase got dramatically cheaper, the quality of the assembled genome sequences got worse—in some cases much worse. Read length is once again increasing, but we still have not returned to the read lengths of the early 2000s. Genomes sequenced entirely with short reads, such as the recently published panda genome², are starting to appear. On the one hand, it is truly remarkable that we can assemble a large genome entirely from short reads; on the other hand, the number of gaps in the resulting assembly—over 200,000 in the panda genome—leaves much to be desired.

Similarly, the rapid decline in costs owing to next-generation sequencing has driven an explosion of metagenomics projects. Scientists begin with DNA extracted from an environment such as soil, water or the human gut and sequence the unknown mixture of species contained in it. Some of these environments are remarkably rich in species, containing thousands of unknown bacteria, archaea and viruses^{3,4}. Unlike whole-genome shotgun sequencing projects, which begin with a clonal sample grown from a single bacterial cell, metagenomics samples contains complex populations, in which any one species might be represented by many slightly different variants.

For metagenomics projects, one of the main challenges is to identify the genes in the environment being sequenced. This becomes very difficult when the sequences contain only

fragments of those genes. Fortunately, bacterial genomes tend to be very gene-rich, and fragments of just a bit longer than a kilobase are likely to contain complete genes. If only there were a way to convert short, next-generation reads into longer fragments.

The method developed by members of Shendure's lab¹ can be used to create 'sub-assembled' (SA) reads averaging ~500 bp from a pair of short reads: a 20-bp 'tag' read and a 76-bp 'breakpoint' read. In this method (Fig. 1), first DNA is sheared and size-selected to pull out fragments of ~550 bp. These fragments are PCR-amplified, tagged, concatenated and sheared randomly to produce from each original fragment a set of shorter fragments. The shorter fragments are then sequenced from both ends using an Illumina Genome Analyzer II. The 20-bp sequences from tagged ends are used to group together the reads, and the breakpoint reads, all of which came from copies of the same original DNA fragment, are then assembled together. The resulting sub-assemblies, or SA reads, can then be treated as long reads and assembled using conventional genome assembly software.

To demonstrate that their method works, Hiatt *et al.*¹ applied it to two different problems: a conventional, whole-genome shotgun assembly of a bacterial genome and an assembly of a metagenomics sample. In the first case, they applied their method to a strain of *Pseudomonas aeruginosa*, a pathogenic bacterium associated primarily with lung infections. When combined with a modest amount of paired-end sequence data, the result was comparable to what one might expect from a conventional sequencing project: they assembled the 6.1 Mb genome into just 32 scaffolds covering >99% of the genome. For the metagenomics experiment, they used DNA collected from sediment at the bottom of a lake and produced SA reads with a median length over 400 bp. When compared to conventional Sanger sequencing of the same sample, the new technique produced a comparable amount of total sequence (after assembly), although somewhat shorter contig lengths, with far less sequencing effort.

A reasonable question is how well this method compares to pyrosequencing with the Roche 454 sequencer, which can currently produce read lengths of ~400 bp, similar in length to the SA reads. In supplementary data, Hiatt *et al.*¹ report that their cost (per megabase) is approximately half of that of 454 sequencing, and they also point out that their SA reads, because they are assembled from multiple overlapping short reads, have a much lower error rate than individual 454 reads. And although they focused on 'long' reads of 500 bp, in principle they could use any length for which next-generation sequencers can capture both ends.

With its very low costs, short-read sequencing will likely dominate the genome sequencing world until something better comes along. But short reads will always leave more gaps than long reads, and no one likes gaps. As the scientific community sequences an unending stream of genomes and metagenomes (and there seems to be an almost limitless number in our biosphere), techniques for producing long DNA sequences will always be needed, until the day arrives when we can grab onto one end of a chromosome and sequence the entire molecule in one go. That day is likely quite far off.

References

1. Hiatt JB, Patwardhan RP, Turner EH, Lee C, Shendure J. Nat. Methods 2010;7:119–122. [PubMed: 20081835]
2. Li R, et al. *Nature* advance online publication. December 13;2009 doi:10.1038/nature08696.
3. DeLong EF. Nat. Rev. Microbiol 2005;3:459–469. [PubMed: 15886695]
4. Turnbaugh PJ, et al. *Nature* 2007;449:804–810. [PubMed: 17943116]

Figure 1.

Turning short reads into long reads using subassembly. The process begins with size-selected fragments of approximately 500 bp. The fragments get unique tags on both ends (red and green), and all fragments are then amplified using PCR. Random shearing breaks each fragment at many different places. The sheared fragments are sequenced from both ends, producing reads that originate all along the fragments. These reads can then be clustered together based on the unique tags and assembled to produce 'reads' that are nearly as long as the original DNA fragment.