

Fast and accurate *de novo* genome assembly from long uncorrected reads

Robert Vaser^{1*}, Ivan Sović^{2*}, Niranjan Nagarajan³, Mile Šikić^{1,4#}

¹*Department of Electronic Systems and Information Processing, University of Zagreb, Faculty of Electrical Engineering and Computing, Unska 3, 10000 Zagreb, Croatia*

²*Centre for Informatics and Computing, Ruđer Bošković Institute, Bijenička 54, 10000 Zagreb, Croatia*

³*Genome Institute of Singapore, Singapore 138672, Singapore*

⁴*Bioinformatics Institute, Singapore 138671, Singapore*

* Contributed equally.

Correspondence should be addressed to: mile.sikic@fer.hr

The assembly of long reads from Pacific Biosciences and Oxford Nanopore Technologies typically requires resource intensive error correction and consensus generation steps to obtain high quality assemblies. We show that the error correction step can be omitted and high quality consensus sequences can be generated efficiently with a SIMD accelerated, partial order alignment based stand-alone consensus module called Racon. Based on tests with PacBio and Oxford Nanopore datasets we show that Racon coupled with Miniasm enables consensus genomes with similar or better quality than state-of-the-art methods while being an order of magnitude faster.

Introduction

With the advent of long read sequencing technologies from Pacific Biosciences (PacBio) and Oxford Nanopore Technologies (ONT), the ability to produce genome assemblies with high contiguity has received a significant fillip. However, to cope with the relatively high error rates (>5%) of these technologies assembly pipelines have typically relied on resource intensive error correction (of reads) and consensus generation (from the assembly) steps (Chin et al. 2013; Loman et al. 2015). More recent methods such as Falcon (Chin et al. 2016) and Canu (Koren et al. 2016) have refined this approach and have significantly improved runtimes but are still computationally demanding for large genomes (Sović et al. 2016a). Recently, Li (Li 2016) showed that long erroneous reads can be assembled without the need for a time-consuming error-correction step. The resulting assembler, Miniasm, is an order of magnitude faster than other long-read assemblers, but produces sequences which can have >10 times as many errors as other methods (Sović et al. 2016a). As fast and accurate long-read assemblers can enable a range of applications, from more routine assembly of mammalian and plant genomes, to structural variation detection, improved metagenomic classification and even online, "read until" assembly (Loose et al. 2016), a fast and accurate consensus module is a critical need. This was also noted by Li (Li 2016), highlighting that fast assembly was only feasible if a consensus module matching the speed of Minimap and Miniasm was developed as well.

Here we address this need by providing a very fast consensus module called Racon (for Rapid Consensus), which when paired with a fast assembler such as Miniasm, enables the efficient construction of genome sequences with high accuracy even without an error correction step.

Racon provides a standalone, platform-independent consensus module for long and erroneous reads and can also be used as a fast and accurate read correction tool.

Results

Racon is designed as a user friendly standalone consensus module that is not explicitly tied to any *de novo* assembly method or sequencing technology. It reads multiple input formats (GFA, FASTA, FASTQ, SAM, MHAP and PAF), allowing simple interoperability and modular design of new pipelines. Even though other stand-alone consensus modules, such as Quiver (Chin et al. 2013) and Nanopolish (Loman et al. 2015) exist, they require sequencer specific input and are intended to be applied *after* the consensus phase of assembly to further polish the sequence. Racon is run with sequencer-independent input (only uses base quality values), is robust enough to work with uncorrected read data and is designed to rapidly generate high-quality consensus sequences. These sequences can be further polished with Quiver or Nanopolish or by applying Racon for more iterations.

Racon takes as input a set of raw backbone sequences, a set of reads and a set of mappings between reads and backbone sequences. Mappings can be generated using any mapper/overlapper which supports either the MHAP, PAF or SAM output formats, such as Minimap (Li 2016), MHAP (Berlin et al. 2015) or GraphMap (Sović et al. 2016b). In our tests, we used Minimap as the mapper as it was the fastest and provided reasonable results. Racon uses the mapping information to construct a partial order alignment graph, using a Single Instruction Multiple Data (SIMD) implementation to accelerate the process (SPOA). More details on Racon and SPOA can be found in the Methods section.

For the purpose of evaluation, we paired Racon with Minimap (as an overlapper) and Miniasm to form a fast and accurate *de novo* assembly pipeline (referred to here as *Miniasm+Racon* for short). Our complete pipeline is thus composed of four steps: (I) Minimap for overlap detection, (II) Miniasm layout for generating raw contigs, (III) Minimap for mapping of raw reads to raw contigs and (IV) Racon for generating high-quality consensus sequences. Steps three and four can be run multiple times to achieve further iterations of the consensus. We then compared the Miniasm+Racon pipeline to other state-of-the-art *de novo* assembly tools for third generation sequencing data (i.e. Falcon and Canu). Note that Falcon and Canu have previously been

benchmarked with other assembly methods such as PBcR and a pipeline from Loman *et al.* (Loman *et al.* 2015) and shown to produce high quality assemblies with improved running times (Sović *et al.* 2016a). Assembly pipelines were evaluated in terms of consensus sequence quality (**Table 1**), runtime and memory usage (**Table 2; Figure 1**), and scalability with respect to genome size (**Figure 2**), on several PacBio and Oxford Nanopore datasets (see Methods).

As can be seen from **Table 1**, all assembly pipelines were able to produce assemblies with high coverage of the reference genome and in a few contigs. Canu, Falcon and the Miniasm+Racon pipeline also constructed sequences with comparable sequence identity to the reference genome, with the iterative use of Racon serving as a polishing step for obtaining higher sequence identity. In addition, the Miniasm+Racon pipeline was found to be significantly faster for all datasets, with a 5-301× speedup compared to Canu and 9-218× speedup compared to FALCON (**Figure 1**).

Racon's speedup was more pronounced for larger genomes and is likely explained by the observation that it scales linearly with genome size (for fixed coverage; **Figure 2**).

The runtime of the Miniasm+Racon pipeline was dominated by the time for the consensus generation step in Racon, highlighting that this step is still the most compute intensive one for small genomes (**Table 3**). However, the results in **Table 3** suggest that for larger genomes the mapping computation stage can catch up in terms of resource usage. Furthermore, if a polishing stage is used, this would typically be more resource intensive.

Comparison of the results of the various assembly pipelines after a polishing stage confirmed that the use of Racon provided better results than just the Miniasm assembly (avg. identity of 99.80% vs 98.06%) and that the Miniasm+Racon assembly matched the best reported sequence quality for this dataset (from the Loman *et al.* pipeline; Sović *et al.* 2016a), while providing a better match to the *actual size* of the reference genome (4641652 bp; **Table 4**). We additionally observed that Nanopolish executed >6× faster on Miniasm+Racon contigs than on raw Miniasm assemblies (248.28 CPUh vs. 1561.80 CPUh), and the Miniasm+Racon+Nanopolish approach achieved the same sequence quality as the original Loman *et al.* pipeline, while being much faster.

We also evaluated Racon's use as an error-correction module. The three main conceptual differences between error correction and consensus modes in Racon are that: (I) instead of mapping reads to a reference, overlapping of reads needs to be performed (e.g. using Minimap with overlap parameters instead of defaults), (II) multiple overlaps are not filtered like mappings (when mapping, we only want to find one best mapping position), and (III) the parallelization is not conducted on a per-window level, but on a per-read level to reduce the system time, i.e. each read consumes one thread. We noted that Racon corrected reads had error rates comparable to Falcon and Canu but provided better coverage of the genome (**Table 5**). Overall, Nanocorrect (Loman et al. 2015) had the best results in terms of error rate but it had lower reference coverage and was more than two orders of magnitude slower than Racon.

Finally, the default mode of Racon determines the consensus by splicing results of non-overlapping windows together. We additionally attempted to assess the influence of using overlapping windows, where we extended each window by 10% on both of its ends. Neighboring windows were then aligned and clipped to form the final consensus sequence. Using this approach, we re-evaluated the consensus quality (**Supplemental Table S1**), the speed (**Supplemental Table S2**) and the quality of error correction (**Supplemental Table S3**) to the default Racon mode on all datasets. Using the window overlapping mode the consensus quality increased (up to 0.06%) in all but one case. The most significant increase was obtained on the *E. coli* PacBio dataset, where average identity rose from 99.90% to 99.94%. This improvement came at the expense of higher running time ($\approx 10-15\%$) and an additional heuristic (window overlap length), which is why we enabled this mode as a non-default option within the Racon module.

Recently, another standalone consensus module called Sparc (Ye and Ma 2016) emerged. We evaluated Sparc in a similar manner to Racon, with the only difference being in the mapping step (step (III) in the Miniasm+Racon pipeline) where Sparc uses BLASR (Chaisson and Tesler 2012). We named this pipeline Miniasm+Sparc accordingly. The results of the comparison are shown in **Table 6** which contains consensus sequence quality, runtime and memory usage. It can be seen that the consensus quality of the Miniasm+Sparc pipeline is consistently lower than that of the Miniasm+Racon pipeline. In addition, Miniasm+Sparc is up to 10 \times slower on larger genomes. Also, we noticed that Sparc crashed on some contigs due to unknown reasons.

Discussion

The principal contribution of this work is to take the concept of fast, error-correction-free, long read assembly, as embodied by the recently developed program Miniasm, to its logical end. Miniasm is remarkably efficient and effective in taking erroneous long reads and producing contig sequences that are structurally accurate (Sović et al. 2016a). However, assemblies from Miniasm do not match up in terms of sequence quality when compared with the best assemblies that can be produced with existing assemblers. This serves as a significant barrier for adopting this 'light-weight' approach to assembly, despite its attractiveness for greater adoption of *de novo* assembly methods in genomics. In this work we show that the sequence quality of a correction-free assembler can indeed be efficiently boosted to a quality comparable to other resource-intensive state-of-the-art assemblers. This makes the tradeoff offered much more attractive and the concept of a correction-free assembler more practically useful. Racon is able to start from uncorrected contigs and raw reads and still generate accurate sequences efficiently because it exploits the development of a Single Instruction Multiple Data (SIMD) version of the robust partial order alignment framework. This makes the approach scalable to large genomes and general enough to work with data from very different sequencing technologies. With the increasing interest in the development of better third generation assembly pipelines, we believe that Racon can serve as useful plug-in consensus module that enables software reuse and modular design.

We would also like to note that while Miniasm is a very fast tool which produces correct genome structures, its current implementation is not optimized for long repetitive genomes (Li 2016). Miniasm loads all overlaps into RAM which might cause crashes due to insufficient memory resources. Although the focus of this paper is on the Miniasm+Racon pipeline, Racon is designed as a general purpose consensus module which can be coupled with other assembly pipelines as a consensus or polishing step.

Methods

Racon is based on the *Partial Order Alignment* (POA) graph approach (Lee et al. 2002; Lee 2003) and we report the development of a Single Instruction Multiple Data (SIMD) version that

significantly accelerates this analysis. An overview of Racon's steps is given in **Figure 3**. The entire process is also shown in detail in **Algorithm 1**.

To perform consensus calling (or error-correction), Racon depends on an input set of query-to-target mappings (query is the set of reads, while a target is either a set of contigs in the consensus context, or a set of reads in the error-correction context) as well as quality values of the input reads. Racon then loads the mappings and performs simple filtering (**Algorithm 1**, lines 1–3; **Algorithm 2**): (I) at most one mapping per read is kept in consensus context (in error-correction context this particular filtering is disabled), and (II) mappings which have high error-rate (i.e. $|1 - \min(d_q, d_t)/\max(d_q, d_t)| \geq e$, where d_q and d_t are the lengths of the mapping in the query and the target respectively and e is a user-specified error-rate threshold) are removed. For each mapping which survived the filtering process, a fast edit-distance based alignment is performed (Myers 1999) (**Algorithm 1**, lines 4–10). We used the Edlib implementation of the Myers algorithm (Šošić and Šikić 2016). This alignment is needed only to split the reads into chunks which fall into particular non-overlapping windows on the backbone sequence. Chunks of reads with average quality lower than a predefined threshold are removed from the corresponding window. With this quality filter, we are able to use only high quality parts of reads in the final consensus. Each window is then processed independently in a separate thread by constructing a POA graph using SIMD acceleration and calling the consensus of the window. Quality values are used again during POA graph construction, where each edge is weighted by the sum of qualities of its source and destination nodes (bases; logarithmic domain). The final consensus sequence is then constructed by splicing the individual window consensi together (per contig or read to be corrected).

Algorithm 1. The Racon algorithm for consensus generation.

Input: Set of noisy target sequences (e.g. contigs) \mathcal{T} to perform consensus on, set of read sequences \mathcal{Q} where each $q \in \mathcal{Q}$ is a tuple consisting of $(seq, qual)$ (sequence data and quality values), set of mappings (or overlaps) \mathcal{O} between \mathcal{Q} and \mathcal{T} , window size w , minimum allowed quality value qv_{min} , mode m to run the Racon in (can be either "consensus" or "errorcorrection"), and maximum expected error rate e . A mapping/overlap $o \in \mathcal{O}$ is a tuple of 12 values: $(q_{id}, t_{id}, p_{cb}, n_k, q_{rev}, q_s, q_e, q_{len}, t_{rev}, t_s, t_e, t_{len})$, where q_{id} and t_{id} are the query and target sequence IDs, p_{cb} percentage of query bases covered by the overlap, n_k number of exact matching seeds shared between the query and the target used to construct the mapping/overlap, q_{rev} and t_{rev} are 0 if query and target are forward oriented and 1 otherwise, q_s, q_e, t_s, t_e are the start and end coordinates of the mapping/overlap in the query and target, and q_{len}, t_{len} total query and target sequence lengths.

Output: Set of corrected target sequences \mathcal{T}_c

Function RACON($\mathcal{T}, \mathcal{Q}, \mathcal{O}, w, qv_{min}, m, e$) **begin**

```

1   $\mathcal{O}_f \leftarrow FilterErroneousMappings(\mathcal{O}, e)$ 
2  if  $m \neq "errorcorrection"$  then
3     $\mathcal{O}_f \leftarrow FilterUniqueMappings(\mathcal{O}_f)$ 
4   $\mathcal{A} \leftarrow$  empty hash table of sets ▷ Aligned mappings/overlaps grouped by  $t_{id}$ 
5  foreach  $o \in \mathcal{O}_f$  do ▷ Align filtered mappings/overlaps
6     $q_{seq} \leftarrow \mathcal{Q}[o.q_{id}].seq[o.q_s...o.q_e]$ 
7     $q_{qual} \leftarrow \mathcal{Q}[o.q_{id}].qual[o.q_s...o.q_e]$ 
8     $t_{seq} \leftarrow \mathcal{T}[o.t_{id}][o.t_s...o.t_e]$ 
9     $aln \leftarrow$  bit-vector global alignment between  $q_{seq}$  and  $t_{seq}$ 
10    $\mathcal{A}[o.t_{id}] \leftarrow \mathcal{A}[o.t_{id}] \cup (o, aln, q_{seq}, q_{qual}, t_{seq})$ 
11   $\mathcal{T}_c \leftarrow \emptyset$  ▷  $\mathcal{A}$  set for consensus sequences
12  foreach  $t_{id} \in \mathcal{A}$  do ▷ Process each target sequence individually
13     $\mathcal{A}_t \leftarrow \mathcal{A}[t_{id}]$ 
14     $t_c \leftarrow$  empty string
15    for  $w_{id} = 0$  to  $\lceil \mathcal{T}[t_{id}]/w \rceil$  step  $w$  do ▷ Process a target in non-overlapping windows
16       $w_s \leftarrow w_{id} \cdot w$  ▷ Window start position on target seq
17       $w_e \leftarrow (w_{id} + 1) \cdot w - 1$  ▷ Window end position on target seq
18       $s_{tseq} \leftarrow \mathcal{T}[t_{id}][w_s...w_e]$  ▷ Noisy target subsequence for SPOA
19       $s_{tqual} \leftarrow$  an array of  $|s_{tseq}|$  elements initialized to 0 ▷ Dummy QVs
20       $\mathcal{S} \leftarrow \emptyset$ 
21      foreach  $(o, aln, q_{seq}, q_{qual}, t_{seq}) \in \mathcal{A}_t$  do ▷ Get subsequences of all reads
22         $(s_{qseq}, s_{qqual}) \leftarrow$  extract a substring of  $s_{qseq}$  which aligns between  $w_s$  and  $w_e$  on target with
          corresponding quality values  $q_{qqual}$ 
23        if  $average(s_{qqual}) \geq qv_{min}$  then
24           $\mathcal{S} \leftarrow \mathcal{S} \cup (s_{qseq}, s_{qqual})$ 
25       $\mathcal{G} \leftarrow BuildGraph(s_{tseq}, s_{tqual})$  ▷ As described in (Lee et al. 2002)
26      foreach  $(s_{qseq}, s_{qqual}) \in \mathcal{S}$  do
27         $\mathcal{G} \leftarrow AlignSequenceToGraphNW(\mathcal{G}, s_{qseq}, s_{qqual}, 5, -4, -8, -6)$ 
28       $cons \leftarrow GenerateConsensus(\mathcal{G})$  ▷ As described in (Lee 2003)
29       $t_c \leftarrow$  concatenate  $cons$  to  $t_c$ 
30     $\mathcal{T}_c \leftarrow \mathcal{T}_c \cup t_c$ 
31  return  $\mathcal{T}_c$ 

```

Algorithm 2. Functions for filtering mappings/overlaps in Racon.

Input: A set of mappings (or overlaps) \mathcal{O} and expected maximum error rate e . Mappings/overlaps are defined the same as in Algorithm 1.

Output: Set of filtered mappings/overlaps \mathcal{O}_f

Function FILTERERRONEOUSMAPPINGS(\mathcal{O}, e) **begin**

```

1   $\mathcal{O}_f \leftarrow \emptyset$ 
2  foreach  $o \in \mathcal{O}$  do
3       $o_{max} \leftarrow \max(|o.q_e - o.q_s|, |o.t_e - o.t_s|)$   $\triangleright$  Maximum mapping/overlap length
4       $o_{min} \leftarrow \min(|o.q_e - o.q_s|, |o.t_e - o.t_s|)$   $\triangleright$  Minimum mapping/overlap length
5      if  $(1 - o_{min}/o_{max}) > e$  then  $\triangleright$  Max mapping/overlap indel error rate
6          continue
7       $\mathcal{O}_f \leftarrow \mathcal{O}_f \cup o$ 
8  return  $\mathcal{O}_f$ 

```

Input: A set of mappings (or overlaps) \mathcal{O} . Mappings/overlaps are defined the same as in Algorithm 1.

Output: Set of filtered mappings/overlaps \mathcal{O}_f

Function FILTERUNIQUEMAPPINGS(\mathcal{O}) **begin**

```

9   $\mathcal{H} \leftarrow$  empty hash table
10 foreach  $o \in \mathcal{O}$  do
11     if  $o.q_{id} \in \mathcal{H}$  then  $\triangleright$  Keep only one mapping/overlap per query (read)
12         if  $o.n_k > \mathcal{H}[o.q_{id}].n_k$  then
13              $\mathcal{H}[o.q_{id}] \leftarrow o$ 
14         else
15              $\mathcal{H}[o.q_{id}] \leftarrow o$ 
16  $\mathcal{O}_f \leftarrow \emptyset$ 
17 foreach  $q_{id} \in \mathcal{H}$  do
18      $\mathcal{O}_f \leftarrow \mathcal{O}_f \cup \mathcal{H}[q_{id}]$ 
19 return  $\mathcal{O}_f$ 

```

Partial order alignment and SIMD vectorization.

POA performs Multiple Sequence Alignment (MSA) through a directed acyclic graph (DAG), where nodes are individual bases of input sequences, and weighted, directed edges represent whether two bases are neighboring in any of the sequences. Weights of the edges represent the multiplicity (coverage) of each transition. Alternatively, weights can be set according to the base qualities of sequenced data. The alignment is carried out directly through dynamic programming (DP) between a new sequence and a pre-built graph. While the regular DP for pairwise alignment has time complexity of $O(3nm)$, where n and m are the lengths of the sequences being aligned, the sequence to graph alignment has a complexity of $O((2n_p + 1)n|V|)$, where n_p is the average number of predecessors in the graph and $|V|$ is the number of nodes in the graph (Lee et al. 2002).

Consensus sequences are obtained from a built POA graph by performing a topological sort and processing the nodes from left to right. For each node v , the highest-weighted in-edge e of

weight e_w is chosen, and a score is assigned to v such that $scores[v] = e_w + scores[w]$, where w is the source node of the edge e (Lee 2003). The node w is marked as a predecessor of v , and a final consensus is generated by performing a traceback from the highest scoring node r . In case r is an internal node (r has out edges), Lee (Lee 2003) proposed the idea of branch completion, where all scores for all nodes except $scores[r]$ would be set to a negative value, and the traversal would continue from r as before, with the only exception that nodes with negative scores could not be added as predecessors to any other node.

One important advantage of POA is its speed, with its linear time complexity in the number of sequences (Lee et al. 2002). However, implementations of POA in current error-correction modules, such as Nanocorrect, are prohibitively slow for larger datasets. In order to increase the speed of POA while retaining its robustness, we explored a Single Instruction Multiple Data (SIMD) version of the algorithm (SPOA).

SPOA (**Figure 4; Algorithm 3**) is inspired by the Rognes and Seeberg Smith-Waterman intra-set parallelization approach (Rognes and Seeberg 2000). It places the SIMD vectors parallel to the query sequence (the read), while placing a graph on the other dimension of the DP matrix (**Figure 4**). In our implementation, the matrices used for tracking the maximum local-alignment scores ending in gaps are stored entirely in memory (**Algorithm 3**, line 8 and 10). These matrices are needed to access scores of predecessors of particular nodes during alignment. Unlike regular Gotoh alignment (Gotoh 1982), for each row in the POA DP matrix all its predecessors (via in-edges of the corresponding node in graph) need to be processed as well (**Algorithm 3**, line 17). All columns are then processed using SIMD operations in a query-parallel manner and the values of Gotoh's vertical matrix (**Algorithm 3**, line 20) and a partial update to Gotoh's main scoring matrix (**Algorithm 3**, line 24) are calculated. SIMD operations in **Algorithm 3** process 8 cells of the DP matrix at a time (16-bit registers). A temporary variable is used to keep the last cell of the previous vector for every predecessor (**Algorithm 3**, lines 21–23), which is needed to compare the upper-left diagonal of the current cell to the cell one row up. Processing the matrix horizontally is not performed using SIMD operations due to data dependencies (each cell depends on the result of the cell to the left of it), and are instead processed linearly (**Algorithm 3**, lines 25–33). SPOA uses shifting and masking to calculate every particular value of a SIMD vector individually (**Algorithm 3**, lines 29–31). After the alignment is completed, the

traceback is performed (**Algorithm 3**, line 39) and integrated into the existing POA graph (**Algorithm 3**, line 40).

SIMD intrinsics decrease the time complexity for alignment from $O((2n_p + 1) n|V|)$ to roughly $O((2n_p/k + 1) n|V|)$, where k is the number of variables that fit in a SIMD vector. SPOA supports Intel SSE version 4.1 and higher, which embed 128 bit registers. Both short (16 bits) and long (32 bits) integer precisions are supported (therefore k equals 8 and 4 variables, respectively). 8 bit precision is insufficient for the intended application of SPOA and is therefore not used. Alongside global alignment displayed in **Algorithm 3**, SPOA supports local and semi-global alignment modes, in which SIMD vectorization is implemented as well.

Algorithm 3. Pseudocode for the SPOA algorithm. The displayed function aligns a sequence to a pre-constructed POA graph using SIMD intrinsics. Capitalized variables are SIMD vectors. Alignment mode is Needleman-Wunsch.

Input: A POA graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where \mathcal{V} is a set of vertices and \mathcal{E} a set of edges, a sequence s_{seq} to align to the graph, its quality values s_{qual} and alignment parameters $match$ (m), $mismatch$ (mm), gap_{open} (q) and gap_{extend} (r)

Output: A POA graph $\mathcal{G}_a = (\mathcal{V}_a, \mathcal{E}_a)$ with incorporated s_{seq} and s_{qual}

Function ALIGNSEQUENCETOGRAPHNW($\mathcal{G}, s_{seq}, s_{qual}, m, mm, q, r$) **begin**

```

1  b ← big negative value
2  x ← |V|
3  y ← |sseq|/8                                ▷ 8 variables per SIMD vector
4  Σ ← {A, C, T, G}                             ▷ Alphabet
5  QQ ← [q, q, q, q, q, q, q, q]
6  RR ← [r, r, r, r, r, r, r, r]
7  P[|Σ|][y] ← sequenceProfile(sseq, m, mm, Σ)    ▷ As described in (Rognes 2002)
8  H[x][y] ← initialize first row vectors for NW alignment, set all other vectors values to b
9  fcv[y] ← initialize first column for NW alignment
10 F[x][y] ← initialize all vectors values to b
11 M[8] ← [[b, 0, 0, 0, 0, 0, 0, 0], [b, b, 0, 0, 0, 0, 0, 0], ..., [b, b, b, b, b, b, b, b]]    ▷ Masks needed to properly update H
12 S ← [b, b, b, b, b, b, b, b]                  ▷ Score vector
13 max_score ← b
14 max_i ← -1
15 topologicalSort(G)
16 for i = 0 to (x - 1) do
17   for p in V[i].predecessors do
18     X ← [fcv[p], 0, 0, 0, 0, 0, 0, 0]
19     for j = 0 to (y - 1) do                    ▷ Solve the elements of vertical Gotoh matrix
20       F[i][j] ← MAX(F[i][j], H[p][j] + QQ + RR, F[p][j] + RR)
21       T1 ← (H[p][j] RSHIFT 7)
22       T2 ← (H[p][j] LSHIFT 1) OR X
23       X ← T1
24       H[i][j] ← MAX(H[i][j], T2 + P[V[i].base][j], F[i][j])
25   E ← [0, 0, 0, 0, 0, 0, 0, 0, fcv[i]]
26   for j = 0 to (y - 1) do                      ▷ Solve the elements of horizontal Gotoh matrix
27     E ← ((H[i][j] LSHIFT 1) OR (E RSHIFT 7)) + QQ + RR
28     T3 ← E
29     for k = 0 to 7 do                          ▷ Parallelization cannot be used due to dependencies
30       T3 ← (T3 LSHIFT 1) + (RR LSHIFT (k + 1))
31       E ← MAX(E, M[k] OR T3)
32     H[i][j] ← MAX(H[i][j], E)
33     E ← MAX(H[i][j], E - QQ)
34   if |V[i].successors| = 0 then
35     S ← MAX(S, H[i][y - 1])
36     if max_score < maxElement(SCORE) then
37       max_score ← maxElement(SCORE)
38     max_i ← i
39 A ← alignment path traced back from the cell with the highest score (H[max_i][y - 1]) by checking all possible predecessor cells. As elements from SIMD vectors can not be accessed directly, SIMD vectors are loaded into normal vectors when needed.
40 AddAlignmentToGraph(G, A, sseq, squal)        ▷ As described in (Lee et al. 2002)
41 return G

```

Implementation and reproducibility

Racon and SPOA are both implemented in C++. All tests were run using Ubuntu based systems with two 6-core Intel(R) Xeon(R) E5645 CPUs @ 2.40GHz with Hyperthreading, using 12 threads where possible. The versions of various methods used in the comparisons reported here are shown in **Table 7**. On PacBio datasets, Canu was run using the "-pacbio-raw" option, Falcon using the config file presented in **Supplemental Table S4**, and Sparc using "k 1 c 2 g 1 t 0.2". On ONT datasets, Canu was run using the "-nanopore-raw" option, Falcon using the config file presented in **Supplemental Table S5**, and Sparc using "k 2 g 2 t 0.2". To provide initial alignments for Sparc, BLASR was run using "-bestn 1 -m 5 -minMatch 19". Racon was run with the default parameters on all datasets except *S. cerevisiae* S288c ONT R9 where we increased the base quality value threshold to 20 using the parameter "--bq 20" (the default value is 10). To provide the initial raw assemblies for both Racon and Sparc, Minimap (as an overlapper) and Miniasm were used. Minimap was run using the "-Sw5 -L100 -m0" options on all datasets (PacBio and ONT), and Miniasm with default parameters. For error-correction purposes, Canu was run using "-correct -nanopore-raw", Falcon with an additional config entry "target = pre-assembly", and Racon with the "--erc" option.

Datasets

Six publicly available PacBio and Oxford Nanopore datasets were used for evaluation. These are listed in **Table 8**.

Evaluation methods

The quality of called consensus sequences was evaluated primarily using Dnadiff (Delcher et al. 2003). The parameters we took into consideration for comparison include: total number of bases in the query, aligned bases on the reference, aligned bases on the query and average identity. In addition, we measured the time and memory required to perform the entire assembly process by each pipeline.

The quality of error-corrected reads was evaluated by aligning them to the reference genome using GraphMap (Sović et al. 2016b) with settings "-a anchorgotoh", and counting the match, mismatch, insertion and deletion operations in the resulting alignments.

Software Availability

Racon and SPOA are available open source under the MIT license at <https://github.com/isovic/racon.git> and <https://github.com/rvaser/spoa.git>. In addition, the used version of source codes can be found in the Supplemental Material of this paper.

Acknowledgements

This work has been supported in part by Croatian Science Foundation under the project UIP-11-2013-7353. IS is supported in part by the Croatian Academy of Sciences and Arts under the project "Methods for alignment and assembly of DNA sequences using nanopore sequencing data". NN is supported by funding from A*STAR, Singapore.

Author contributions

IS, MS and RV devised the original concept for Racon. IS developed and implemented several prototypes for Racon. RV suggested and implemented the SIMD POA version. IS implemented a version of Racon which uses SPOA for consensus generation and error-correction. IS and RV tested Racon and SPOA. MS and NN provided helpful discussions and guidance for the project. MS helped define applications and devised tests for the paper; IS ran the tests, and collected and formatted the results. IS and RV wrote the paper with input from all authors. MS and RV designed the figures for the paper. NN proof-read and corrected parts of the paper. All authors read the paper and approve of the final version. MS coordinated the project and provided computational resources.

Disclosure declaration

MS has received a complimentary place at an Oxford Nanopore Technologies organized symposium. NN has received travel expenses and accommodation from Oxford Nanopore to speak at two organized symposia.

References

- Berlin K, Koren S, Chin C-S, Drake JP, Landolin JM, Phillippy AM. 2015. Assembling large genomes with single-molecule sequencing and locality-sensitive hashing. *Nature biotechnology* 33(6): 623–630.
- Chaisson MJ and Tesler G. 2012. Mapping single molecule sequencing reads using basic local alignment with successive refinement (BLASR): application and theory. *BMC Bioinformatics* 13, 238.
- Chin C-S, Alexander DH, Marks P, Klammer AA, Drake J, Heiner C, Clum A, Copeland A, Huddleston J, Eichler EE et al. 2013. Nonhybrid, finished microbial genome assemblies from long-read SMRT sequencing data. *Nature Methods* 10(6): 563–569.
- Chin C-S, Peluso P, Sedlazeck FJ, Nattestad M, Concepcion GT, Clum A, Dunn C, O'Malley R, Figueroa-Balderas R, Morales-Cruz A et al. 2016. Phased Diploid Genome Assembly with Single Molecule Real-Time Sequencing. *Nature Methods* doi: 10.1038/nmeth.4035
- Delcher AL, Salzberg SL, Phillippy AM. 2003. Using MUMmer to identify similar regions in large sequence sets. *Curr Protoc Bioinformatics* Chapter 10, Unit 10 3.
- Gotoh O. 1982. An improved algorithm for matching biological sequences. *J Mol Biol* 162, 705–8.
- Istace B, Friedrich A, d'Agata L, Faye S, Payen E, Beluche O, Caradec C, Davidas S, Cruaud C, Liti G et al. 2016. de novo assembly and population genomic survey of natural yeast isolates with the Oxford Nanopore MinION sequencer. *bioRxiv* doi: 10.1101/066613
- Koren S, Walenz BP, Berlin K, Miller JR, Phillippy AM. 2016. Canu: scalable and accurate long-read assembly via adaptive k-mer weighting and repeat separation. *bioRxiv* doi: 10.1101/071282.
- Lee C. 2003. Generating consensus sequences from partial order multiple sequence alignment graphs. *Bioinformatics* 19(8): 999–1008.
- Lee C, Grasso C, Sharlow MF. 2002. Multiple sequence alignment using partial order graphs. *Bioinformatics* 18(3): 452–464.
- Li H. 2016. Minimap and Miniasm: fast mapping and de novo assembly for noisy long sequences. *Bioinformatics* 32(14): 2103–2110.
- Loman NJ, Quick J, Simpson JT. 2015. A complete bacterial genome assembled de novo using only nanopore sequencing data. *Nature Methods* 12(8): 733–735.

Loose M, Malla S, Stout M. 2016. Real-time selective sequencing using nanopore technology. *Nature Methods* doi:10.1038/nmeth.3930.

Myers G. 1999. A Fast Bit-Vector Algorithm for Approximate String Matching Based on Dynamic Programming. *Journal of the ACM* 46(3): 395–415.

Rognes T, Seeberg E. 2000. Six-fold speed-up of Smith-Waterman sequence database searches using parallel processing on common microprocessors. *Bioinformatics* 16(8): 699–706.

Sović I, Križanović K, Skala K, Šikić M. 2016a. Evaluation of hybrid and non-hybrid methods for de novo assembly of nanopore reads. *Bioinformatics* doi: 10.1093/bioinformatics/btw237.

Sović I, Šikić M, Wilm A, Fenlon SN, Chen S, Nagarajan N. 2016b. Fast and sensitive mapping of error-prone nanopore sequencing reads with GraphMap. *Nature Communications* 7.

Šošić M, Šikić M. 2016. Edlib: A C/C++ library for fast, exact sequence alignment using edit distance. bioRxiv doi: <http://dx.doi.org/10.1101/070649>.

Ye C, Ma Z. 2016. Sparc: a sparsity-based consensus algorithm for long erroneous sequencing reads. *PeerJ* 4:e2016 <https://doi.org/10.7717/peerj.2016>.

Tables

Table 1. Assembly and consensus results across 6 datasets of varying genome length and sequencing data type.

		Miniasm+Racon 1 iteration	Miniasm+Racon 2 iterations	Canu	Falcon
<i>Lambda</i> ONT R7.3 30×	Ref. genome size [bp]	48502	48502	48502	48502
	Total bases [bp]	47916	47872	25077	7212
	Ref. chromosomes [#]	1	1	1	1
	Contigs [#]	1	1	1	1
	Aln. bases ref. [bp]	48438 (99.87%)	48434 (99.86%)	25833 (53.26%)	7483 (15.43%)
	Aln. bases query [bp]	47916 (100.00%)	47872 (100.00%)	25077 (100.00%)	7212 (100.00%)
	Avg. Identity	97.59	97.97	96.87	95.77
<i>E. coli</i> K-12 ONT R7.3 54×	Ref. genome size [bp]	4641652	4641652	4641652	4641652
	Total bases [bp]	4637173	4632058	4601503	4580230
	Ref. chromosomes [#]	1	1	1	1
	Contigs [#]	1	1	1	1
	Aln. bases ref. [bp]	4640867 (99.98%)	4641323 (99.99%)	4631173 (99.77%)	4627613 (99.70%)
	Aln. bases query [bp]	4636689 (99.99%)	4632055 (100.00%)	4601365 (100.00%)	4580230 (100.00%)
	Avg. Identity	99.13	99.32	99.28	98.84
<i>S. cerevisiae</i> S288C ONT R9 59×	Ref. genome size [bp]	12157105	12157105	12157105	12157105
	Total bases [bp]	12172019	12167797	12224535	11643917
	Ref. chromosomes [#]	17	17	17	17
	Contigs [#]	42	42	45	41
	Aln. bases ref. [bp]	12104541 (99.57%)	12110095 (99.61%)	12120070 (99.70%)	11885904 (97.77%)
	Aln. bases query [bp]	12108082 (99.48%)	12115796 (99.57%)	12196684 (99.77%)	11643482 (100.00%)
	Avg. Identity	97.88	98.04	98.61	98.22
<i>E. coli</i> K-12 PacBio P6C4 160×	Ref. genome size [bp]	4641652	4641652	4641652	4641652
	Total bases [bp]	4653199	4645508	4664416	4666788
	Ref. chromosomes [#]	1	1	1	1
	Contigs [#]	1	1	1	1
	Aln. bases ref. [bp]	4641501 (100.00%)	4641439 (100.00%)	4641652 (100.00%)	4641652 (100.00%)
	Aln. bases query [bp]	4653111 (100.00%)	4645508 (100.00%)	4664416 (100.00%)	4666788 (100.00%)
	Avg. Identity	99.63	99.90	99.99	99.90
<i>S. cerevisiae</i> W303 PacBio P4C2 127×	Ref. genome size [bp]	12157105	12157105	12157105	12157105
	Total bases [bp]	12071278	12051573	12402332	12003077
	Ref. chromosomes [#]	17	17	17	17
	Contigs [#]	30	30	29	44
	Aln. bases ref. [bp]	12023607 (98.90%)	12025677 (98.92%)	12127627 (99.76%)	11932488 (98.15%)
	Aln. bases query [bp]	12046299 (99.79%)	12027338 (99.80%)	12363941 (99.69%)	11910549 (99.23%)
	Avg. Identity	99.43	99.72	99.86	99.70
<i>C. elegans</i> PacBio P6C4 81×	Ref. genome size [bp]	100272607	100272607	100272607	100272607
	Total bases [bp]	106353704	106392402	106687886	105858394
	Ref. chromosomes [#]	6	6	6	6
	Contigs [#]	77	77	134	242
	Aln. bases ref. [bp]	100017898 (99.75%)	99979140 (99.71%)	100166301 (99.89%)	99295695 (99.03%)
	Aln. bases query [bp]	101711974 (95.64%)	101741297 (95.63%)	102928910 (96.48%)	102008289 (96.36%)
	Avg. Identity	99.44	99.73	99.89	99.74

Table 2. Time/memory consumption for all datasets. Results are presented in the format “CPU time [min] / Maximum memory [GB]”.

	Miniasm+Racon 1 iteration	Miniasm+Racon 2 iterations	Canu	Falcon
<i>Lambda</i> ONT 30×	0.12 / < 0.1	0.25 / < 0.1	2.86 / 1.90	2.29 / 0.854
<i>E. coli</i> K-12 ONT R7.3 54×	25 / 3.32	46 / 3.32	1328 / 4.03	829 / 12.29
<i>S. cerevisiae</i> S288C ONT R9 59×	28 / 7.84	44 / 7.84	13243 / 4.88	9603 / 8.82
<i>E. coli</i> K-12 PacBio P6C4 160×	86 / 9.69	162 / 9.69	773 / 3.60	2908 / 9.93
<i>S. cerevisiae</i> W303 PacBio P4C2 127×	115 / 16.09	215 / 16.09	6375 / 3.65	14808 / 4.78
<i>C. elegans</i> PacBio P6C4 81×	1247 / 50.38	2004 / 50.38	37853 / 10.16	119766 / 7.59

Table 3. Resource usage for various parts of the Miniasm+Racon assembly pipeline. Results are presented in the format “CPU time [min] / Maximum memory [GB]”.

	<i>Lambda</i> ONT	<i>E. coli</i> ONT	<i>S. cerevisiae</i> ONT	<i>E. coli</i> PacBio	<i>S. cerevisiae</i> PacBio	<i>C. elegans</i> PacBio
Minimap overlap	< 0.1 / < 0.1	2.70 / 3.32	10.21 / 7.84	11.23 / 9.69	21.90 / 16.09	473.82 / 50.38
Miniasm	< 0.1 / < 0.1	< 0.1 / < 0.1	0.24 / 0.22	0.46 / 0.40	0.56 / 0.46	4.22 / 3.45
Minimap mapping 1st iter.	< 0.1 / < 0.1	0.24 / 0.24	0.83 / 0.27	0.65 / 0.23	1.48 / 0.26	13.82 / 1.06
Racon consensus 1st iter.	0.11 / < 0.1	22.16 / 2.91	17.06 / 2.11	73.21 / 7.88	91.55 / 4.69	755.29 / 22.11
Minimap mapping 2nd iter.	< 0.1 / < 0.1	0.27 / 0.23	0.89 / 0.30	0.72 / 0.24	1.63 / 0.27	15.15 / 0.96
Racon consensus 2nd iter.	0.13 / < 0.1	20.82 / 2.25	14.85 / 1.98	75.75 / 6.40	97.94 / 4.61	741.68 / 20.71
Total CPU time / Max. mem.	0.25 / < 0.1	46 / 3.32	44 / 7.84	162 / 9.69	215 / 16.09	2004 / 50.38

Table 4. Results after polishing assemblies with Nanopolish on the *E. coli* K-12 ONT R7.3 54× dataset.

	Raw Miniasm	Miniasm+Racon 2 iterations	Canu	Falcon	Loman et. al pipeline
Total bases [bp]	4696482	4641756	4631443	4624811	4695512
[Total bases - Genome size] [bp]	54830	104	10209	16841	53860
Aligned bases ref. [bp]	4635941 (99.88%)	4641312 (99.99%)	4633324 (99.82%)	4627571 (99.70%)	4641325 (99.99%)
Aligned bases query [bp]	4687686 (99.81%)	4641756 (100.00%)	4631361 (100.00%)	4624811 (100.00%)	4695463 (100.00%)
Avg. Identity	98.06	99.80	99.80	99.78	99.80

Table 5. Comparison of error-correction modules on *E. coli* K-12 ONTR7.3 54× dataset. Values presented in the table are median values of the error and match rate estimates.

	CPU time [h]	Coverage	Insertion rate (%)	Deletion rate (%)	Mismatch rate (%)	Match rate (%)	Error rate (I+D+M) (%)
Raw	-	53.55×	5.23	2.83	4.89	89.81	12.95
Racon	10	50.23×	0.52	0.58	0.23	99.25	1.33
Nanocorrect	8100	44.74×	0.14	0.43	0.03	99.83	0.60
Falcon	22	46.95×	0.04	1.11	0.06	99.90	1.21
Canu	5	35.53×	0.06	1.25	0.08	99.85	1.39

Table 6. Quality, time and memory comparison of Racon and Sparc as consensus modules in a *de novo* assembly pipeline composed of Minimap as an overlapper and Miniasm as a layout module. Both consensus tools were run on identical set of raw contigs.

		Miniasm+Racon 1 iteration	Miniasm+Racon 2 iterations	Miniasm+Sparc 1 iteration	Miniasm+Sparc 2 iterations
<i>Lambda</i>	Avg. Identity	97.59	97.97	95.04	96.87
ONT R7.3	CPU Time [min]	0.12	0.25	< 0.1	0.14
30×	Max memory [GB]	< 0.1	< 0.1	< 0.1	< 0.1
<i>E. coli</i> K-12	Avg. Identity	99.13	99.32	98.10	99.16
ONT R7.3	CPU Time [min]	25	46	28	54
54×	Max memory [GB]	3.32	3.32	3.37	3.37
<i>S. cerevisiae</i> S288C	Avg. Identity	97.88	98.04	97.77	98.02
ONT R9	CPU Time [min]	28	44	210	424
59×	Max memory [GB]	7.84	7.84	7.84	7.84
<i>E. coli</i> K-12	Avg. Identity	99.63	99.90	99.34	99.76
PacBio P6C4	CPU Time [min]	86	162	102	187
160×	Max memory [GB]	9.69	9.69	9.72	9.72
<i>S. cerevisiae</i> W303	Avg. Identity	99.43	99.72	99.09	99.57
PacBio P4C2	CPU Time [min]	115	215	297	587
127×	Max memory [GB]	16.09	16.09	16.09	16.09
<i>C. elegans</i>	Avg. Identity	99.44	99.73	99.16	99.57
PacBio P6C4	CPU Time [min]	1247	2004	4032	8915
81×	Max memory [GB]	50.38	50.38	51.80	51.80

Table 7. Versions of various tools used in this paper.

Tool	Commit/Version	Source
Racon	2f41352ab4aa on branch <i>grpaper</i>	https://github.com/isovic/racon
Minimap	1cd6ae3bc7c7	https://github.com/lh3/minimap.git
Miniasm	17d5bd12290e	https://github.com/lh3/miniasm.git
Canu	ab50ba3c0cf0	https://github.com/marbl/canu.git
FALCON-integrate	8bb2737fd1d7	https://github.com/PacificBiosciences/FALCON-integrate.git
Sparc	ae89503a1613	https://github.com/yechengxi/Sparc
Nanocorrect	b09e93772ab4	https://github.com/jts/nanocorrect.git
Nanopolish	47dcd7f147c	https://github.com/jts/nanopolish.git
MUMmer	3.23	http://mummer.sourceforge.net/

Table 8. Description of the six datasets used for evaluation.

Dataset	Description
<i>Lambda</i> phage ONT R7.3	ENA submission ERA476754, with 113× coverage of the NC_001416 reference genome (48502bp). Link: ftp://ftp.sra.ebi.ac.uk/vol1/ERA476/ERA476754/oxfordnanopore_native/Lambda_run_d.tar.gz . This dataset was subsampled to coverages of 30× and 81× for testing.
<i>E. coli</i> K-12 ONT R7.3	54× pass 2D coverage of the genome (U00096.3, 4.6Mbp). Link: http://lab.loman.net/2015/09/24/first-sqk-map-006-experiment/
<i>E. coli</i> K-12 PacBio P6C4	160× coverage of the genome (U00096.3). The dataset was generated using one SMRT Cell of data gathered with a PacBio RS II System and P6-C4 chemistry on a size selected 20kbp library of <i>E. coli</i> K-12. Link: https://s3.amazonaws.com/files.pacb.com/datasets/secondary-analysis/e-coli-k12-P6C4/p6c4_ecoli_RSII_DDR2_with_15kb_cut_E01_1.tar.gz
<i>S. cerevisiae</i> S288c ONT R9	59x 2D (pass and fail) coverage of the genome (http://downloads.yeastgenome.org/sequence/S288C_reference/chromosomes/fasta/). Link: http://www.genoscope.cns.fr/externe/Download/Projets/yeast/datasets/raw_data/S288C/ (Istace et al. 2016).
<i>S. cerevisiae</i> PacBio W303 PacBio P4C2	The dataset is composed of 11 SMRT cells (https://github.com/PacificBiosciences/DevNet/wiki/Saccharomyces-cerevisiae-W303-Assembly-Contigs), of which one was not used in this study because the containing folder (“0019”) was incomplete and the data could not be extracted. The S288C reference (12.1 Mbp) was used for comparison (http://downloads.yeastgenome.org/sequence/S288C_reference/chromosomes/fasta/). Coverage of the dataset with respect to the S288C reference is approx. 127×.
<i>C. elegans</i> PacBio P6C4	Bristol mutant strain, 81× coverage of the genome (gi 449020133). The dataset was generated using 11 SMRT cells P6-C4 chemistry on a size selected 20kbp library. Link: https://github.com/PacificBiosciences/DevNet/wiki/C.-elegans-data-set

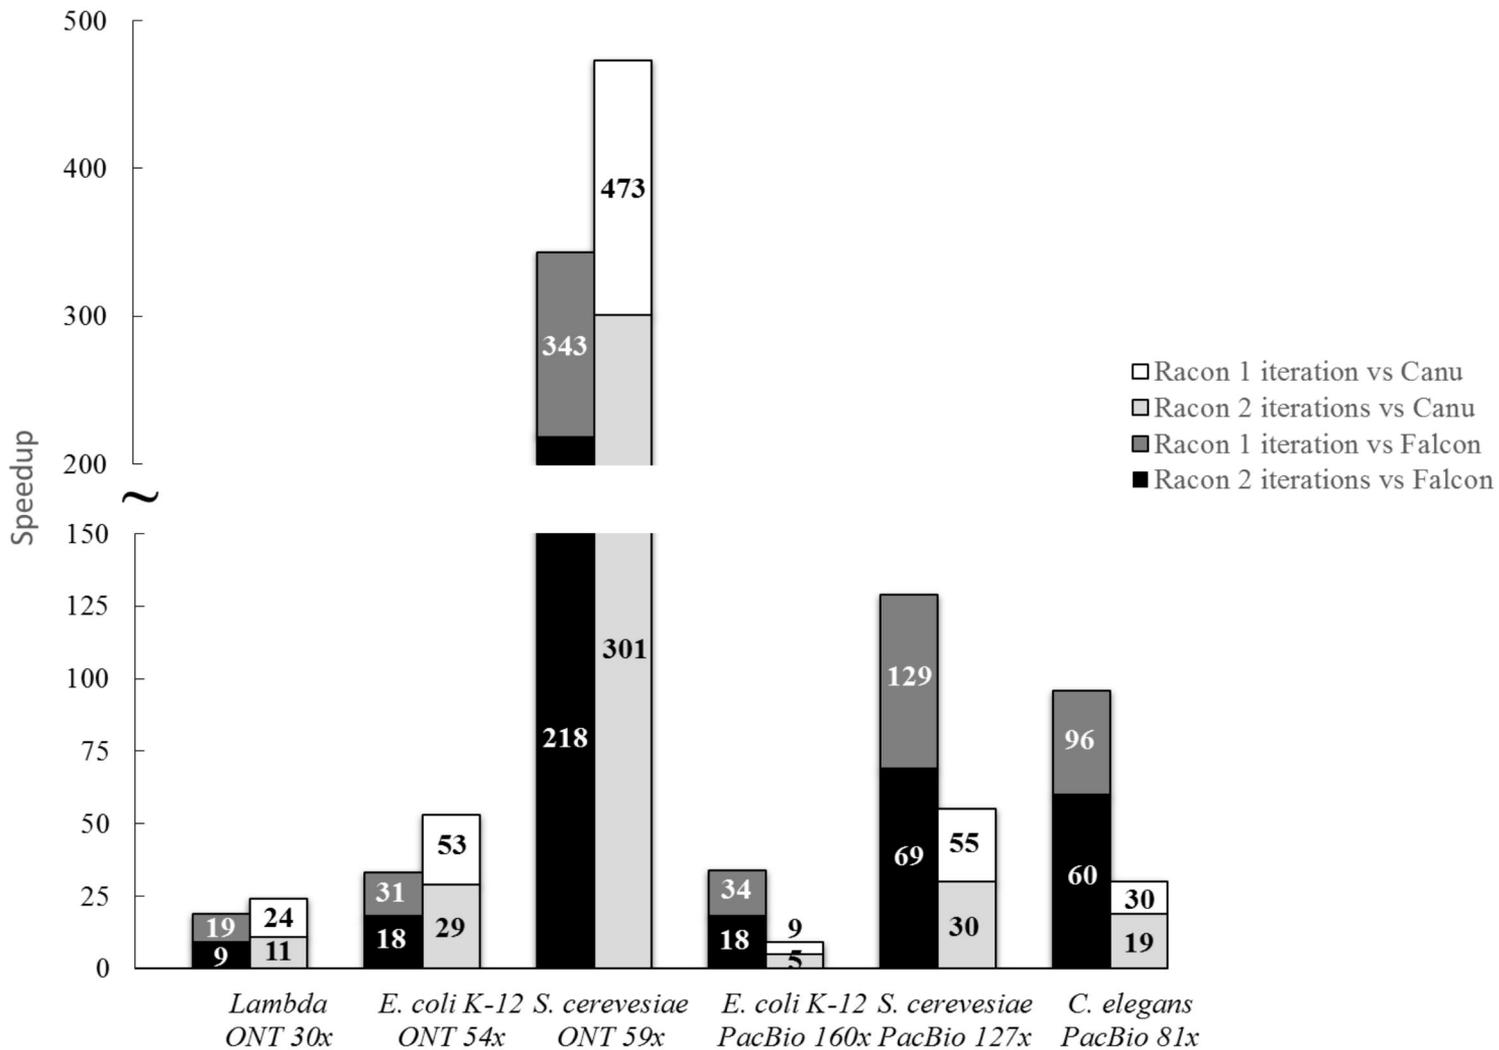
Figures

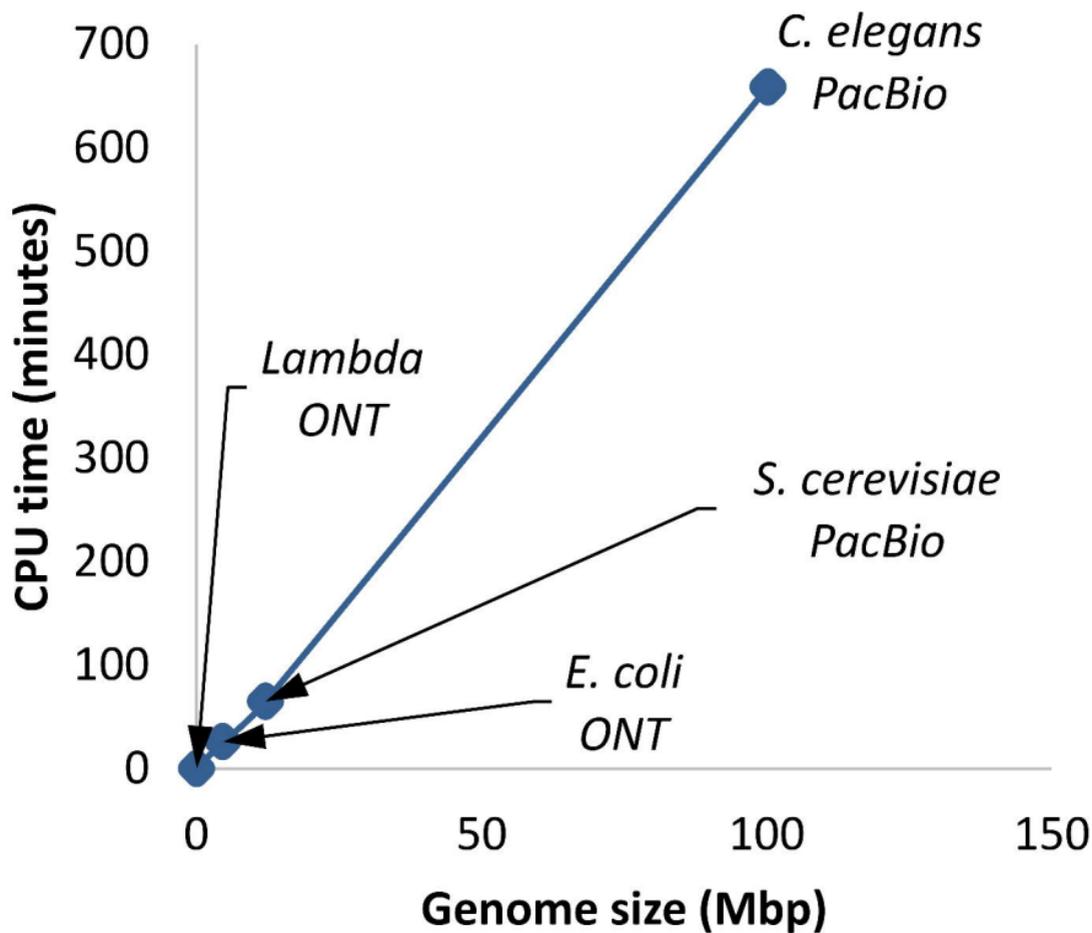
Figure 1. Racon's speedup when compared to Falcon and Canu.

Figure 2. Scalability of Racon as a function of genome size. Read coverage was subsampled to be 81× (limited by the *C. elegans* dataset) and the figure shows results for one iteration of Racon.

Figure 3. Overview of the Racon consensus process.

Figure 4. Depiction of the SIMD vectorization approach used in SPOA.



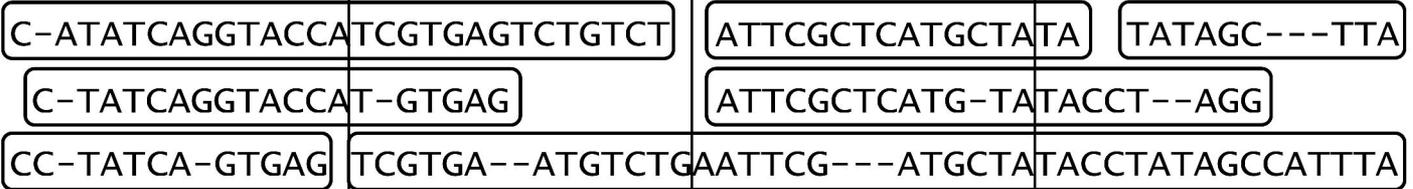


Layout

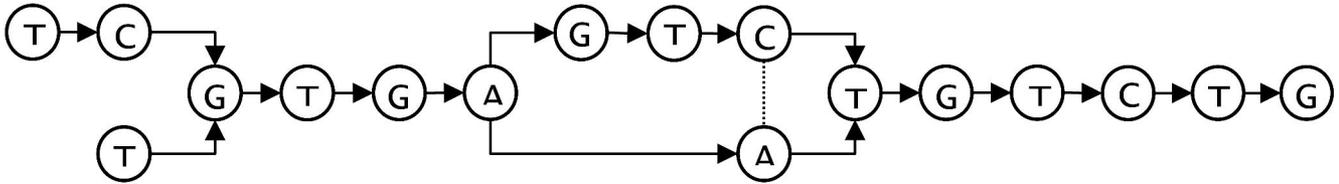
CCATTT CAGGTACCATCGTGAGTCTGTCTGAAATCGCTCATGCTATAGCTATAGCCATTTA

Aligning reads and segmentation

CCATTT CAGGTACCATCGTGAGTCTGTCTGAAATCGCTCATGCTATAGCTATAGCCATTTA



POA graph (SIMD-accelerated)



Segment splicing



Consensus sequence

CCATATCAGGTACCATCGTGAGTCTGTCTGAAATTCGCTCATGCTATACCTATAGCCATTTA



Fast and accurate de novo genome assembly from long uncorrected reads

Robert Vaser, Ivan Sovic, Niranjana Nagarajan, et al.

Genome Res. published online January 18, 2017

Access the most recent version at doi:[10.1101/gr.214270.116](https://doi.org/10.1101/gr.214270.116)

Supplemental Material

<http://genome.cshlp.org/content/suppl/2017/04/07/gr.214270.116.DC1>

P<P

Published online January 18, 2017 in advance of the print journal.

Accepted Manuscript

Peer-reviewed and accepted for publication but not copyedited or typeset; accepted manuscript is likely to differ from the final, published version.

Creative Commons License

This article is distributed exclusively by Cold Spring Harbor Laboratory Press for the first six months after the full-issue publication date (see <http://genome.cshlp.org/site/misc/terms.xhtml>). After six months, it is available under a Creative Commons License (Attribution-NonCommercial 4.0 International), as described at <http://creativecommons.org/licenses/by-nc/4.0/>.

Email Alerting Service

Receive free email alerts when new articles cite this article - sign up in the box at the top right corner of the article or [click here](#).

Advance online articles have been peer reviewed and accepted for publication but have not yet appeared in the paper journal (edited, typeset versions may be posted when available prior to final publication). Advance online articles are citable and establish publication priority; they are indexed by PubMed from initial publication. Citations to Advance online articles must include the digital object identifier (DOIs) and date of initial publication.

To subscribe to *Genome Research* go to:
<http://genome.cshlp.org/subscriptions>
