Coalescent Inference Using Serially Sampled, High-Throughput Sequencing Data from Intrahost HIV Infection

Kevin Dialdestoro,* Jonas Andreas Sibbesen,^{†,1} Lasse Maretty,^{†,1} Jayna Raghwani,[‡] Astrid Gall,[§]

Paul Kellam,^{§,**} Oliver G. Pybus,[‡] Jotun Hein,^{*} and Paul A. Jenkins^{††,††,2}

*Department of Statistics and [†]Department of Zoology, University of Oxford, Oxford OX1 3LB [Statistics] OX1 3PS [Zoology], United Kingdom, [‡]The Bioinformatics Centre, Department of Biology, University of Copenhagen, 2200 Copenhagen, Denmark,

[§]Wellcome Trust Sanger Institute, Wellcome Trust Genome Campus, Hinxton, Cambridge CB10 1SA, United Kingdom, ^{**}Medical Research Council/University College London Centre for Medical Molecular Virology, Division of Infection and Immunity, University College London, London WC1E 6BT, United Kingdom, and ^{††}Department of Statistics and ^{‡‡}Department of Computer Science, University of Warwick, Coventry CV4 7AL, United Kingdom

ORCID ID: 0000-0001-8538-4995 (J.R.)

ABSTRACT Human immunodeficiency virus (HIV) is a rapidly evolving pathogen that causes chronic infections, so genetic diversity within a single infection can be very high. High-throughput “deep” sequencing can now measure this diversity in unprecedented detail, particularly since it can be performed at different time points during an infection, and this offers a potentially powerful way to infer the evolutionary dynamics of the intrahost viral population. However, population genomic inference from HIV sequence data is challenging because of high rates of mutation and recombination, rapid demographic changes, and ongoing selective pressures. In this article we develop a new method for inference using HIV deep sequencing data, using an approach based on importance sampling of ancestral recombination graphs under a multilocus coalescent model. The approach further extends recent progress in the approximation of so-called *conditional sampling distributions*, a quantity of key interest when approximating coalescent likelihoods. The chief novelties of our method are that it is able to infer rates of recombination and mutation, as well as the effective population size, while handling sampling over different time points and missing data without extra computational difficulty. We apply our method to a data set of HIV-1, in which several hundred sequences were obtained from an infected individual at seven time points over 2 years. We find mutation rate and effective population size estimates to be comparable to those produced by the software BEAST. Additionally, our method is able to produce local recombination rate estimates. The software underlying our method, Coalescensor, is freely available.

KEYWORDS coalescent; importance sampling; HIV evolution; conditional sampling distribution; recombination

HUMAN immunodeficiency virus (HIV) is a rapidly evolving pathogen that causes a chronic infection for an individual's lifetime. As a consequence, the genetic diversity within a single infection can be very high. Important clinical variables, such as the rate of progression to AIDS and the set point viral load, are related to the diversity and evolution of

the within-patient viral population (Shankarappa *et al.* 1999; Ross and Rodrigo 2002; Williamson 2003; Edwards *et al.* 2006; Lemey *et al.* 2007; Pybus and Rambaut 2009), and so genetic data from these populations are of medical relevance in addition to providing insight into molecular evolutionary processes. However, population genomic inference from HIV sequence data can be challenging as result of high rates of mutation and recombination within a small RNA genome of ~10 kb. Furthermore, natural selection is expected to play an important role in shaping within-host HIV genetic diversity (Rouzine and Coffin 1999; Neher and Leitner 2010; Batorsky *et al.* 2011; Pennings *et al.* 2014). For example, phylogenies constructed from serially sampled intrahost population sequences are typically characterized by

Copyright © 2016 by the Genetics Society of America
doi: 10.1534/genetics.115.177931

Manuscript received June 5, 2015; accepted for publication January 31, 2016;
published Early Online February 5, 2016.

Supporting Information is available online at www.genetics.org/lookup/suppl/doi:10.1534/genetics.115.177931/-DC1.

¹These authors contributed equally to this work.

²Corresponding author: Department of Statistics, University of Warwick, Coventry CV4 7AL, United Kingdom. E-mail: p.jenkins@warwick.ac.uk

a ladder-like topology (Shankarappa *et al.* 1999), indicating a rapid and continual turnover of genetic lineages as result of continual host immune selection. Although the relationship between census viral population size (*i.e.*, viral loads) over the course of infection and virus kinetics is well understood for HIV (Nowak and May 2000), relatively little is known about the effective population size during a single infection.

High coverage “deep” sequencing is now being used to address these problems and to investigate the genomic diversity of HIV, with several data sets already under study (Henn *et al.* 2012; Poon *et al.* 2012; Gall *et al.* 2013). For the study of within-host evolution of HIV patients, deep sequencing serves as a potentially powerful way to infer the evolutionary and ecological dynamics of the viral population in unprecedented detail, particularly since sequencing can be performed at different time points during infection (Drummond *et al.* 2003). This is especially important for studying fast-evolving RNA viruses, where the substitution rates and effective population size may change through time. By utilizing the temporal structure of the sampled sequences, statistical power can greatly improve the precision of population demographic and evolutionary estimates (Pybus *et al.* 2000; Drummond *et al.* 2003). However, from the perspective of population genomic inference, deep sequencing data are unusual: high-throughput sequencing generates large numbers of short sequence reads, with each read representing 400–700 nt of a viral genome. Typically no virus particle is sampled twice, so that only a small fraction of the viral genome is sequenced for any given virion in the population. Genealogies can be estimated, for example, for each ~500-nt subgenomic partition, but each genealogy represents a different random sample of individuals from the same population. This presents peculiar challenges in genomic inference.

To infer parameters of an underlying evolutionary model from deep sequencing data, such as effective population size, mutation rate, and recombination rate, new theoretical and statistical approaches are therefore needed. In this article we work within a coalescent framework and in particular its extensions to allow for serially sampled, or *heterochronous*, sequences (Rodrigo and Felsenstein 1999). (This is in contrast to the usual situation of *isochronous* sampling at a single fixed time.) The coalescent is a powerful and flexible framework for modeling the genealogy of a large, panmictic population, with many further extensions that incorporate changing population size, recombination, and recurrent mutations (see Hein *et al.* 2005, for a textbook introduction). It is a crucial component in the inference of the evolutionary dynamics of fast-evolving RNA viruses, which can be combined with epidemiological data in an approach known as *phylodynamics* (Grenfell *et al.* 2004). A potentially powerful method of inference under complicated coalescent evolutionary models is to proceed by computationally intensive Monte Carlo simulation. The quantity of interest, such as the likelihood for the data, is estimated by averaging over a large number of the possible unobserved gene genealogies that

could have given rise to the observed sequences. Genealogies of high posterior probability can be targeted by Markov chain Monte Carlo (MCMC) or importance sampling (IS). There is now a sizeable literature on these and related approaches, focusing on various complications of the basic coalescent, depending on the species under study. For example, there are methods to account for recombination [MCMC (Kuhner *et al.* 2000; Wang and Rannala 2008; Rasmussen *et al.* 2014) and IS (Griffiths and Marjoram 1996; Fearnhead and Donnelly 2001; McVean *et al.* 2002; Griffiths *et al.* 2008; Jenkins and Griffiths 2011)], changing population size [MCMC (Beaumont 1999; Drummond *et al.* 2002, 2005; Wilson *et al.* 2003; Minin *et al.* 2008) and IS (Griffiths and Tavaré 1994a; Beaumont 2003; Leblois *et al.* 2014)], and heterochronous sequence data [MCMC (Drummond *et al.* 2002, 2005; Minin *et al.* 2008) and IS (Beaumont 2003; Anderson 2005; Fearnhead 2008)].

However, for accurate inference using deep sequencing data from within-host virus populations we need to account for several of these complications simultaneously, and no existing methods are available for this task. In particular, current inference methods in this context have tended to omit the process of recombination (Pybus and Rambaut 2009), limiting our understanding of the extent of the association between recombination and viral adaptation. However, the effective recombination rate is of an order of magnitude comparable to the mutation rate (for example, Shriner *et al.* 2004 found the recombination rate to be 5.5-fold greater than the mutation rate in HIV); furthermore, recombination may play an important role in the evolution of drug resistance (Kellam and Larder 1995; Neher and Leitner 2010). There is also growing evidence that recombination rates are not constant along the genome (Fan *et al.* 2007; Archer *et al.* 2008). Thus, our goal in this article is to develop a coalescent method of inference that can handle *all* of the following: (i) recombination, (ii) high mutation rates, (iii) heterochronous sequences, (iv) missing data, and (v) changing effective population size.

Here we take a novel importance sampling approach to tackle this problem. It is based on recent developments in the efficient computation of *conditional sampling distributions* (CSDs) (Paul and Song 2010; Paul *et al.* 2011; see also Sheehan *et al.* 2013): the probability distribution of an additionally sampled haplotype, conditioned on the sampled haplotypes we have seen already. These are crucial in the design of efficient IS algorithms: Stephens and Donnelly (2000) noted an equivalence between designing an IS proposal distribution and approximating the (unknown) CSD. This work was later formalized by De Iorio and Griffiths (2004a), which subsequently allowed Griffiths *et al.* (2008) and Paul and Song (2010) to derive an accurate approximate CSD in the presence of crossover recombination. Finally, work by Paul *et al.* (2011) resulted in an efficient approximation of the CSD for multiple loci via a hidden Markov model (HMM).

Our work is most closely related to that of Griffiths *et al.* (2008), whose focus was a model able to account for recombination, but was not designed with heterochronous deep

sequencing data in mind and thus exhibited only properties i and ii above. Extending this method to include properties iii, iv, and v raises numerous methodological challenges that we address in further detail below. Briefly, we replace the CSD of Griffiths *et al.* (2008) with that of the sequentially Markov model of Paul *et al.* (2011). The latter is more efficient to compute and easily extended to blocks of completely linked sites. To allow for samples taken at different times we introduce an explicit time variable, which in turn determines when samples are inserted into the reconstruction of a genealogy backward in time. Further, and in contrast to the imputation approaches of Fearnhead and Donnelly (2001) and Griffiths *et al.* (2008), our model allows for haplotypes to be only *partially* specified, assigning allelic states only at a subset of loci where possible. This dramatically reduces the state space of the model, and it also provides a convenient way of handling missing data.

Despite all of these contributions, application to large high-throughput data sets remains challenging, particularly for full-likelihood methods. To retain tractability, recent research has turned to approximations of the model, or of the likelihood, or both. The approach we take here is a principled method to find an accurate *full*-likelihood solution first by restricting attention to a *two-locus model*. Our two-locus model is then extended to multilocus data by analyzing selected pairs of loci separately and then aggregating these pairwise inferences, either by taking the median of the inferred estimates or by combining their likelihood surfaces via a pseudolikelihood [in particular, via a pairwise composite likelihood (McVean *et al.* 2002; Larribe and Fearnhead 2011)]. As we describe below, we performed a simulation study to demonstrate that our method can recover model parameters accurately.

The inferential process involves simulating ancestral trees backward in time, and, in agreement with Griffiths and Tavaré (1994b), Nielsen (1997), and Jasra *et al.* (2011), as we get closer to the most recent common ancestor (MRCA), coalescence times increase greatly. This adds undesirably high variance to the inference and extensive central processing unit (CPU) time. To circumvent this, we further employ the *Time Machine* strategy developed by Jasra *et al.* (2011): stopping the simulation before the MRCA is reached while controlling for the bias, acquiring sizable computational saving and reduced variance.

To investigate the performance of our method on empirical data, we analyzed HIV-1 RNA samples that had been extracted and sequenced at seven time points over a period of 2 years, from a patient enrolled in the control arm of the Short Pulse Antiretroviral Therapy at Seroconversion (SPARTAC) study. This patient received no antiviral drugs. Nine regions from the whole HIV-1 genome alignments of these heterochronous data were then analyzed using our model and for comparison by the Bayesian MCMC coalescence approach implemented in BEAST (Drummond and Rambaut 2007; Drummond *et al.* 2012). Both analyses provide strong agreement in the mutation rate, effective population size, and time to the most recent common ancestor. However, in addition, our approach also estimates recombination rates.

A C++ implementation of the algorithm is available from <https://github.com/OSSCB2013> under the name Coalescentator. The program can process serially sampled, deep sequencing data from viral populations, to infer mutation rates, recombination rates, the effective population size, and times to recent ancestors.

Materials and Methods

Model and notation

In this section we describe our notation and genealogical model. We begin by formulating a two-locus model; that is, given a pair of loci, call them A and B , recombination may occur at any position along the sequence separating them. We regard a *locus* as a fixed stretch of nucleotides, which may contain several polymorphic sites. Recombination within a locus is ignored.

Demographic model: Suppose we have data $\mathcal{D} = (\mathcal{D}_{-K}, \mathcal{D}_{-K+1}, \dots, \mathcal{D}_0)$ collected at time points $t_{-K} > t_{-K+1} > \dots > t_0 = 0$, with t_0 being the most recent collection time and $t_{-1}, t_{-2}, \dots$ extending into the past. We assume the effective population size (N_e) to be constant between collection times, but it may change at these times. This piecewise-constant model is similar to that of Pybus *et al.* (2000). Except through their effects on N_e , we otherwise ignore the effects of natural selection and population substructure.

To reduce the parameter space, viral load estimates at the sampling times can be used as a guide for constraining the relative magnitudes of the effective population sizes. In the simplest scenario (*e.g.*, when viral load estimates are unavailable or unreliable or we decide not to use them), a single, constant effective population size can be fitted, and in this article we focus on the inference of a single effective population size parameter. In *Results*, we investigate the robustness of our inference to this assumption.

Mutation model: Let l_A and l_B denote the lengths (in nucleotides) of loci A and B , respectively, and set $l = l_A + l_B$ to be their total length. We assume that all nucleotides conform to a diallelic model, with alleles labeled arbitrarily as $\{0, 1\}$. Each nucleotide mutates independently at the same rate and with symmetric transitions such that the mutation transition matrix for each nucleotide is $\mathbf{P} = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$. (In other words, P_{ij} denotes the probability that, if the i th allele undergoes a mutation event, it changes to the j th allele.) Let $\bar{\mu}$ be the mutation rate per generation time per site, which we assume to be constant across sites, and let $\bar{\theta} = 4N_e\bar{\mu}$ be the population-scaled mutation rate per generation per site when the effective population size is N_e . Then $\theta = l\bar{\theta}$ is the population-scaled mutation rate per generation time across both loci, with individual locus mutation rates given by $\theta_A = l_A\bar{\theta}$ and $\theta_B = l_B\bar{\theta}$. The overall mutation transition matrix for locus ℓ is

$$\mathbf{Q} = \sum_{k=1}^{l_t} \frac{1}{l_t} \mathbf{I} \otimes \mathbf{I} \otimes \dots \otimes \mathbf{P} \otimes \dots \otimes \mathbf{I}, \quad (1)$$

where $\mathbf{I}$ is the identity matrix and for each k th summand, $\mathbf{P}$ appears in the k th position in the direct product (Griffiths and Tavaré 1994b). In other words, when a mutation occurs at locus ℓ , the resulting haplotype is chosen uniformly from among those differing from the parental haplotype at precisely one of the l_t sites. Because of the possibility of a high rate of mutation, this model allows for a site to have undergone more than one mutation event in its genealogical history. In practice we simulate mutation events only at polymorphic sites.

Recombination model: Suppose loci A and B are separated by a region of length d nucleotides, and assume a uniform recombination rate across this region. Let $\bar{r}$ be the recombination rate per generation time per site, which we assume to be constant across sites. Therefore, if $\bar{\rho} = 4N_e\bar{r}$ denotes the population-scaled recombination rate per generation time per site, then $\rho = d\bar{\rho}$ is the population-scaled recombination rate per generation time between loci A and B . We assume only crossover-type recombination, which implies an approximately linear relationship between short physical distance and recombination rate.

Sample notation: We sample n haplotypes at a given collection time; by *haplotype* we mean a sequence of nucleotides that we will later subdivide into a sequence of contiguous blocks of loci. Due to missing data, we observe a haplotypes specified only at locus A , b haplotypes specified only at locus B , and c haplotypes specified at both loci, so that $a + b + c = n$. Reads only partially overlapping with a locus are treated as missing at this locus. A haplotype $(i, j) \in \{0, 1\}^{l_a} \times \{0, 1\}^{l_b}$ is seen with multiplicity c_{ij} , so that $\sum_{i,j} c_{ij} = c$. Haplotypes with data missing at locus A or B are respectively denoted $(i, *)$ and $(*, j)$, with respective multiplicities a_i and b_j , and $\sum_i a_i = a$, $\sum_j b_j = b$. Let $\mathbf{c} = (c_{ij})$, $\mathbf{a} = (a_i)$, and $\mathbf{b} = (b_j)$; the complete data set is thus written compactly as $\mathbf{n} = (\mathbf{a}, \mathbf{b}, \mathbf{c})$. As we reconstruct a genealogy for the sample backward in time, recombination events create lineages each ancestral only at one of the two loci. The notation we have introduced to allow for missing data also allows us to leave unspecified the allelic types of these nonancestral recombinant loci: backward in time, a recombination event replaces a type (i, j) with $(i, *)$ and $(*, j)$.

Importance sampling

Parameter estimation in population genetic models requires computation of the likelihood of the observed data $\mathcal{D}$ as a function of the model parameters Φ . We start by describing a method for estimating the likelihood $L(\Phi)$ for heterochronous data at two loci, with $\Phi = (N_e, \bar{\mu}, \bar{r})$. Our method also provides a weighted approximation to the posterior distribution of genealogical histories given these data, so it is straightforward to

address questions of ancestral inference, such as the time to the most recent common ancestor (TMRCA) of the data. Note that under the standard (isochronous) coalescent model, it is not possible to identify mutation and recombination rates separately, as only their respective products with the effective population size can be identified. Serial sampling from a rapidly evolving population (relative to the spacing between samples) allows for the separate estimation of these three parameters, because a separate estimate of N_e is then available. [To understand this, observe that with heterochronous data the times between sample collections appear in the likelihood expression, $L(\Phi)$. The timescale of the coalescent is governed by N_e , which therefore now appears in $L(\Phi)$ —and so can be estimated—separately from its appearance inside the composite parameters $\bar{\theta}$ and $\bar{\rho}$. Further details are given by Rodrigo and Felsenstein 1999, pp. 259–262.]

As it is not possible to write down $L(\Phi) = \mathbb{P}(\mathcal{D}; \Phi)$ analytically, a *latent* genealogy variable, $\mathcal{G}$, is introduced. The likelihood is then calculated by marginalizing over $\mathcal{G}$:

$$L(\Phi) = \int_{\mathcal{G}} \mathbb{P}(\mathcal{D}|\mathcal{G})\mathbb{P}(\mathcal{G}; \Phi) d\mathcal{G}. \quad (2)$$

A naive Monte Carlo estimator for the integral in (2) is given by

$$\tilde{L}(\Phi) = \frac{1}{M} \sum_{i=1}^M \mathbb{P}(\mathcal{D}|\mathcal{G}^{(i)}), \quad (3)$$

where $\mathcal{G}^{(i)}$ represent independent samples from $\mathbb{P}(\mathcal{G}; \Phi)$, the coalescent prior. This estimator has poor properties because $\mathbb{P}(\mathcal{D}|\mathcal{G}^{(i)}) = 0$ with high probability, resulting in $\tilde{L}(\Phi)$ having very high variance (Stephens and Donnelly 2000). Estimation of the integral (2) is therefore generally conducted using either MCMC or IS. The latter approach, which we follow here, was pioneered by Griffiths and Tavaré (1994b). An IS estimator is obtained by introducing an artificial *proposal distribution* $\mathbf{Q}$ and reweighting the samples; the estimator

$$\hat{L}(\Phi) = \frac{1}{M} \sum_{i=1}^M \mathbb{P}(\mathcal{D}|\mathcal{G}^{(i)}) \frac{\mathbb{P}(\mathcal{G}^{(i)}; \Phi)}{\mathbf{Q}(\mathcal{G}^{(i)}; \Phi)}, \quad (4)$$

where $\mathcal{G}^{(i)}$ represent independent samples from $\mathbf{Q}(\mathcal{G}; \Phi)$, has lower variance than (3) provided $\mathbf{Q}(\mathcal{G}; \Phi)$ is chosen carefully. It can be shown that the *optimal* proposal distribution is the posterior $\mathbb{P}(\mathcal{G}|\mathcal{D}; \Phi)$, in which case $\hat{L}(\Phi)$ has variance 0 (Stephens and Donnelly 2000). This distribution is unknown in general, but it can guide our intuition on how to choose $\mathbf{Q}(\mathcal{G}; \Phi)$.

Griffiths and Tavaré (1994b) noted that we can focus solely on genealogies compatible with the observed data by simulating them sequentially and by a reversal of time. To illustrate the idea we introduce some further notation similar to that of Stephens and Donnelly (2000): first, regard the coalescent model as a Markov process forward in time on unordered sets of haplotypes. The process starts at the MRCA of the sample, $H_{-m} \in \{0, 1\}^l$, and ends at an unordered set of

n observed haplotypes, $H_0 \in (\{0, 1\})^n$, corresponding to the most recently observed data. The process visits a sequence of states $\mathcal{H} = (H_{-m}, H_{-m+1}, \dots, H_0)$ as the genealogy is constructed forward in time by mutations, coalescences, and recombinations; the known rates of these transitions are prescribed by the coalescent model. We refer to $\mathcal{H}$ as the *history* of the sample. The importance sampling estimator can then be decomposed as

$$\hat{L}(\Phi) = \frac{1}{M} \sum_{i=1}^M \mathbb{P}(\mathcal{D}|\mathcal{H}^{(i)}) \prod_{k=1}^m \frac{\mathbb{P}(H_{-k+1}^{(i)}|H_{-k}^{(i)}; \Phi)}{\mathbb{Q}(H_{-k}^{(i)}|H_{-k+1}^{(i)}; \Phi)}.$$

The numerators $\mathbb{P}(H_{-k+1}|H_{-k}; \Phi)$ are known from the coalescent model, and the denominators $\mathbb{Q}(H_{-k}^{(i)}|H_{-k+1}^{(i)}; \Phi)$ are specified by the proposal distribution. The key point is that this reformulation of the importance sampling procedure can be viewed as exploring the state space of latent histories *backward* from the sample H_0 to the MRCA H_{-m} . One advantage of this is immediate: when the only data collection time is $t_0 = 0$, then $\mathcal{D} = H_0$, and by choosing $H_0^{(i)} = H_0$ for each $i = 1, \dots, M$ we can ensure that $\mathbb{P}(\mathcal{D}|\mathcal{H}^{(i)}) = 1$ for each i , so that no simulation is wasted. Moreover, the Markov specification of the proposal distribution helps to simplify our task of actually designing it.

A key observation of Stephens and Donnelly (2000) was that the posterior, conditional distribution of genealogies can be expressed in terms of the CSD, which is defined as the probability $\pi[(i, j)|\mathbf{n}; \Phi]$ that a haplotype sampled from the population is of type $(i, j) \in \{0, 1\}^l$ given that we have already sampled the haplotypes with configuration $\mathbf{n}$. Although this distribution is as intractable as $\mathbb{P}(\mathcal{G}|\mathcal{D}; \Phi)$, it provides a sensible starting point in the design of an efficient proposal distribution, since we have reduced the problem to one of finding a good approximation of $\pi[(i, j)|\mathbf{n}; \Phi]$. We emphasize that even if we use a proposal distribution based on an *approximation* to $\pi[(i, j)|\mathbf{n}; \Phi]$, the reweighting in (4) still provides us with an unbiased and consistent estimator.

New proposal distribution

Our aim then is to design a sequential proposal distribution $\mathbb{Q}(H_{-k}|H_{-k+1}, \dots; \Phi)$, defining a Markov chain backward in time that approximates $\mathbb{P}(\mathcal{H}|\mathcal{D}; \Phi)$. Following the discussion of the previous section, to achieve this we use two steps: first, express the true backward transition rates $\mathbb{P}(H_{-k}|H_{-(k-1)}; \Phi)$ in terms of the CSDs $\pi[\cdot|\cdot; \Phi]$; second, substitute the CSD for a well-motivated approximation $\hat{\pi}[\cdot|\cdot; \Phi]$. We address each of these steps in turn.

Backward transition rates: For now, fix N_e and assume a single sample collection time $t_0 = 0$. Griffiths *et al.* (2008) obtained the backward transition rates for a two-locus, finite-alleles model with isochronous samples taken from a population at equilibrium. Their formulation is efficient because it corresponds to ancestral recombination graphs (ARGs) in which only lineages carrying genetic material an-

cestral to the sample are simulated; entirely nonancestral lineages are integrated out. This dramatically reduces the state space and improves efficiency. However, the simulation continues to assign allelic types to the remaining nonancestral *loci*. For example, if a lineage in an ARG is ancestral to a member of the sample at locus A but not at locus B , then it is still necessary to assign an allelic status at locus B . In our application these loci represent tens or hundreds of nucleotides, and assigning a haplotypic status to such loci throughout the simulation is cumbersome. We therefore marginalize over these nonancestral loci too. The resulting backward transition rates are given in Table A1, and the entries in Table A1 are derived in *Appendix A*.

Thus far we have ignored changes in population size and heterochronous sampling. To account for these factors we introduce an explicit time variable $\mathcal{T} = (T_{-m}, T_{-m+1}, \dots, T_0 = 0)$, which contains the times between events in the history $\mathcal{H}$ (we now also include the collection of additional samples at times $t_{-1}, t_{-2}, \dots$ as valid events). To calculate the likelihood, we now need to marginalize over both histories $\mathcal{H}$ and interevent times $\mathcal{T}$,

$$L(\Phi) = \int_{\mathcal{H}, \mathcal{T}} \mathbb{P}(\mathcal{D}|\mathcal{H}, \mathcal{T}) \mathbb{P}(\mathcal{H}, \mathcal{T}; \Phi) d\mathcal{H} d\mathcal{T},$$

with IS estimator

$$\hat{L}(\Phi) = \frac{1}{M} \sum_{i=1}^M \mathbb{P}(\mathcal{D}|\mathcal{H}^{(i)}, \mathcal{T}^{(i)}) \frac{\mathbb{P}(\mathcal{H}^{(i)}, \mathcal{T}^{(i)}; \Phi)}{\mathbb{Q}(\mathcal{H}^{(i)}, \mathcal{T}^{(i)}|\mathcal{D}; \Phi)}, \quad (5)$$

where the parameters of interest are $\Phi = (N_e, \bar{\mu}, \bar{r})$, and $(\mathcal{H}^{(i)}, \mathcal{T}^{(i)})$ are independent samples from $\mathbb{Q}(\mathcal{H}, \mathcal{T}|\mathcal{D}; \Phi)$.

Our sequential proposal $\mathbb{Q}(H_{-k-1}, T_{-k-1}|H_{-k}, T_{-k}; \Phi)$ is described below. Suppose our genealogical reconstruction has proceeded backward at time $T_{-k}^{\text{sum}} = T_0 + T_1 + \dots + T_{-k}$, past the sample collection time t_{-j} . We first sample an event time backward according to an exponential distribution with rate depending on the current sample size (a, b, c) and the current effective population size N_{-j} . In apprehension of the appearance of new samples at given times in the past, we measure time in chronological units rather than in generation units (Rodrigo and Felsenstein 1999). Then the waiting time T (in days) for the next event is exponentially distributed with density

$$f_{\text{event}}(t) = \frac{D}{2N_{-j}\tau} \exp\left[-\frac{D}{2N_{-j}\tau}t\right], \quad t > 0, \quad (6)$$

where $D = n(n-1) + (a+c)\theta_A + (b+c)\theta_B + \rho c$ is (twice) the total prior rate of events and τ is the generation time in days. (The parameters θ_A , θ_B , and ρ are redefined each time the effective population size changes; e.g., $\rho = 4dN_{-j}\bar{r}$.) We sample a time T from (6): if the corresponding event time is more recent than the next (pastward) sample collection time, i.e., $T_{-k}^{\text{sum}} + T < t_{-j-1}$, then we accept this time and define our proposal by

$$\begin{aligned} & \mathbb{Q}(H_{-k-1}, T_{-k-1} | H_{-k}, T_k; \Phi) \\ &= f_{\text{event}}(T_{-k-1}) \mathbb{Q}(H_{-k-1} | H_{-k}; \Phi), \end{aligned}$$

with $\mathbb{Q}(H_{-k-1} | H_{-k}; \Phi)$ defined as in Table A1—only we replace each instance of $\pi[\cdot; \Phi]$ with an approximation $\hat{\pi}[\cdot; \Phi]$ developed below. Otherwise, with probability

$$\mathbb{P}(\text{no event in } [T_{-k}^{\text{sum}}, t_{-j-1}]) = \exp\left[-\frac{D}{2N\tau}(t_{-j-1} - T_{-k}^{\text{sum}})\right]$$

the next event is the insertion of additional samples at the next collection time, and we set the next waiting time to be $T_{-k-1} = t_{-j-1} - T_{-k}^{\text{sum}}$. At collection time t_{-j-1} , we insert the samples $\mathcal{D}_{-j-1}$ collected at time t_{-j-1} into our current configuration H_{-k} . That is, if $\mathcal{D}_{-j-1} = \mathbf{m}$ and $H_{-k} = \mathbf{n}$, then we set $H_{-k-1} = \mathbf{m} + \mathbf{n}$. We then update the effective population size parameter to N_{-j-1} and continue this iterative procedure until we have both passed the oldest collection time and reached a MRCA for each locus. The overall procedure is illustrated in Figure 1.

Conditional sampling distribution: Our proposal distribution is obtained by substituting an approximate CSD $\hat{\pi}$ for π in Table A1. We use the HMM approach devised by Paul *et al.* (2011), yielding an algorithm linear in both the number of loci and the number of haplotypes (i.e., linear in sequence length and sample size, respectively). It is a relatively accurate approximation of the true CSD and practical to compute. Briefly, the new haplotype (i, j) is sampled given an existing configuration $\mathbf{n}$ by assuming that the genealogy for the latter is a simple, improper “trunk” ancestry: one extending infinitely far back into the past with no mutations, recombinations, or coalescences (see Figure 2 and Paul and Song 2010). The alleles of the new haplotype are determined by allowing its ancestral lineage to undergo mutation and recombination and to coalesce into this trunk ancestry at appropriate rates. This framework can be cast as an HMM across loci, whose emissions are the observed alleles at the newly sampled haplotype and whose hidden state is the pair $s_\ell = (\tau_\ell, i)$, where τ_ℓ is the absorption time of this locus into the trunk, $i \in \{0, 1\}^{l_\ell}$ is the type of the lineage into which absorption occurs (Figure 2), and l_ℓ is the length of this locus (in nucleotides).

Although this framework is obviously a strong simplification of the true underlying genealogical process, the resulting CSD $\hat{\pi}[(i, j) | \mathbf{n}; \Phi]$ has many sensible properties (Paul *et al.* 2011). However, this approach does not lend itself immediately to handling unspecified alleles. We therefore make several further modifications to the CSD of Paul *et al.* (2011) to construct a practicable IS algorithm, as we now describe.

Emission probabilities: For our model, where all nucleotides are diallelic and mutate independently, we can simplify the emission probabilities used in the HMM calculation—increasing the computational efficiency greatly. Given a hidden state $s_\ell = (\tau_\ell, i)$ of the Markov chain at locus ℓ , the emission probability mass function for observed state $j \in \{0, 1\}^{l_\ell}$ is given by

Figure 1 Illustration of our two-locus recombination model: a sampled history $\mathcal{H}$ and interevent times $\mathcal{T}$. The two loci of each haplotype are each represented by a circle. White circles represent an unspecified locus and colored circles indicate the allelic type at that locus. For example, H_0 consists of types (blue, red) and (green, *). There are two sampling times and the collected samples are represented by the leaves of the tree (marked by rectangles). Time is measured in chronological units and run backward from the most recent collection time, $t_0 = 0$, to the most recent common ancestor, t_{MRCA} . Ancestral lineages are represented by black lines. At a coalescence event, two lineages are joined together; the model allows coalescence between fully specified haplotypes (H_{-4}), between a fully specified and partially specified haplotype (H_{-6}), and between two partially specified haplotypes (H_{-2}). At a recombination event, two lineages are created and their haplotypes are partially specified: one of the two loci becomes nonancestral and its allele type is left unspecified (H_{-1}). At the next collection time t_{-1} , a new sample $\mathcal{D}_{-1}$ is added to the existing lineages $H_{-2} : H_{-3} = H_{-2} + \mathcal{D}_{-1}$; and the effective population size is allowed to change.

$$\xi_\theta(j | s_\ell) = e^{-(\theta/2)\tau_\ell} \sum_{m=0}^{\infty} \frac{1}{m!} \left(\frac{\theta}{2}\tau_\ell\right)^m [\mathbf{Q}^m]_{ij}, \quad (7)$$

where $\mathbf{Q}$ is the mutation transition matrix (1). Because sites within a locus mutate independently in our model, we can reformulate this as

$$\xi_\theta(j | s_\ell) = \prod_{k=1}^{l_\ell} e^{-(\bar{\theta}/2)\tau_\ell} \sum_{m=0}^{\infty} \frac{1}{m!} \left(\frac{\bar{\theta}}{2}\tau_\ell\right)^m [\mathbf{P}^m]_{i_k j_k},$$

where we recall that $\mathbf{P} = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$. We can simplify this even further as

$$\begin{aligned} \xi_\theta(j | s_\ell) &= \left[e^{-(\bar{\theta}/2)\tau_\ell} \sum_{s=0}^{\infty} \frac{(\bar{\theta}\tau_\ell/2)^{(2s+1)}}{(2s+1)!} \right]^{S_{ij}} \left[e^{-(\bar{\theta}/2)\tau_\ell} \sum_{s=0}^{\infty} \frac{(\bar{\theta}\tau_\ell/2)^{(2s)}}{(2s)!} \right]^{l_\ell - S_{ij}} \\ &= e^{-(\theta/2)\tau_\ell} [\sinh(\bar{\theta}\tau_\ell/2)]^{S_{ij}} [\cosh(\bar{\theta}\tau_\ell/2)]^{l_\ell - S_{ij}}, \end{aligned} \quad (8)$$

where S_{ij} denotes the number of nucleotide differences between i and j . Thus, we have successfully eliminated the infinite sum in the emission distribution (7). Because the hidden state τ_ℓ is continuous, we follow Paul *et al.* (2011) and employ

Figure 2 Illustration of the sequential interpretation for a realization of $\hat{\pi}[\mathbf{e} = (i, j) | \mathbf{n}; \Phi]$ for two loci. The dotted and solid lines, respectively, represent the marginal genealogies (S_1, S_2) at loci A and B. The hidden state at locus A is $s_1 = (\tau_1, h_1)$. Haplotype h_1 would carry a green allele at its first locus, but a mutation results in the observed blue allele. The hidden state at locus B is $s_2 = (\tau_2, h_2)$. h_2 carries a yellow allele at its second locus, and no mutation occurs on the marginal genealogy at this locus. If there is no recombination, $s_2 = s_1$, but here a recombination occurs before τ_2 and the absorption time for the second locus is $\tau_2 \neq \tau_1$. As in Figure 1, white circles represent loci with unspecified alleles.

Gaussian quadrature to construct a discrete HMM. We chose Laguerre–Gauss quadrature (Abramowitz and Stegun 1972, table 25.9) with 16 points. Applying the forward algorithm to this HMM yields the required CSD. The expression for the emission distribution in this discretized model can also be reduced to a closed-form formula (Appendix B).

Emission probability when the absorbing hidden state is unspecified: The method of Paul *et al.* (2011) does not deal with the case where the locus of interest at the absorbing hidden state is unspecified, *i.e.*, calculating emissions of the form $\xi_\theta(j | (\tau_\ell, *))$. In a two-locus model this occurs, for example, when the absorbing state is the haplotype $(i, *)$ for some i and we are interested in the emission distribution at locus B. In this case, we *condition* on choosing an absorbing haplotype with the allele at this locus specified (Figure 3B). When there are no haplotypes in the trunk ancestry for which the locus of interest has a specified allele, we sample the allele from the stationary distribution of $\mathbf{Q}$ (Equation 1); that is, we pick uniformly from all possible 2^L haplotypes at this locus,

$$\xi_\theta(j | (\tau_\ell, *)) = \begin{cases} \sum_{k: a_k + c_k > 0} \frac{a_k + c_k}{a + c} \xi_\theta(j | (\tau_\ell, k)) & \text{if } \ell = A \text{ and } a + c > 0, \\ \sum_{l: b_l + c_l > 0} \frac{b_l + c_l}{b + c} \xi_\theta(j | (\tau_\ell, l)) & \text{if } \ell = B \text{ and } b + c > 0, \\ 2^{-L} & \text{otherwise,} \end{cases} \quad (9)$$

where the trunk ancestry has configuration $\mathbf{n} = (\mathbf{a}, \mathbf{b}, \mathbf{c})$ and we have used the notation $c_k = \sum_l c_{kl}$, $c_l = \sum_k c_{kl}$.

Emission probability for partially observed haplotypes: If the observed haplotype has an unspecified locus (*i.e.*, we seek to calculate $\hat{\pi}[(i, *) | \mathbf{n}; \Phi]$ or $\hat{\pi}[(*, j) | \mathbf{n}; \Phi]$), then we transform this to a one-locus problem with a degenerate HMM summing over hidden states at a single locus (Figure 3A). If the absorbing hidden state is unspecified (Figure 3B),

Figure 3 Sampling $\hat{\pi}[(*, j) | \mathbf{n}; \Phi]$, with the observed allele j represented by a yellow circle. (A) Absorption at the second lineage of the trunk ancestry for which the second locus is specified (red allele). A mutation event is still allowed in this one-locus model, as illustrated here by a mutation from a red to a yellow allele. (B) Absorption at the first lineage of the trunk ancestry for which the second lineage is unspecified. In such cases we choose uniformly from the other informative lineages as the absorbing state.

then we choose uniformly from the informative trunk lineages as described above.

CSD for recombination event: To model recombination, we require $\pi[\{(i, *), (*, j)\} | \mathbf{n}; \Phi]$, the CSD for two partially observed haplotypes (see Table A1). This quantity satisfies the following decompositions and symmetries:

$$\pi[\{(i, *), (*, j)\} | \mathbf{n}; \Phi] = \sum_{i', j'} \pi[(i, j') | \mathbf{n}; \Phi] \pi[(i', j) | \mathbf{n} + \mathbf{e}_{ij'}^C; \Phi] \quad (10a)$$

$$= \sum_{i', j'} \pi[(i', j) | \mathbf{n}; \Phi] \pi[(i, j') | \mathbf{n} + \mathbf{e}_{ij'}^C; \Phi], \quad (10b)$$

and

$$\pi[\{(i, *), (*, j)\} | \mathbf{n}; \Phi] = \pi[(i, *) | \mathbf{n}; \Phi] \pi[(*, j) | \mathbf{n} + \mathbf{e}_i^A; \Phi] \quad (11a)$$

$$= \pi[(*, j) | \mathbf{n}; \Phi] \pi[(i, *) | \mathbf{n} + \mathbf{e}_j^B; \Phi]. \quad (11b)$$

To estimate $\hat{\pi}[\{(i, *), (*, j)\} | \mathbf{n}]$, Griffiths *et al.* (2008) use relationship (10) by substituting approximate CSDs $\hat{\pi}$ for each fully observed haplotype. In addition, they noted that $\hat{\pi}$ may not satisfy symmetries (10a) and (10b)—and so they take the average of the two expressions. However, this strategy in

averaging over all the alleles at the unspecified loci is very computationally intensive. Here, we instead use relationship (11) and substitute the approximate CSDs $\hat{\pi}[(i, *)|\mathbf{n}; \Phi]$ and $\hat{\pi}[(*, j)|\mathbf{n}; \Phi]$ from the previous section (*Emission probability for partially observed haplotypes*). We still take the average of the two expressions (11a) and (11b) to account for asymmetries.

Forward transitions: To summarize, solving the forward equation associated with the HMM described above allows us to compute an approximate CSD, $\hat{\pi}[(i, j)|\mathbf{n}; \Phi]$, for each $(i, j) \in \{0, 1\}^l$ as well as for the special cases $(i, *)$ and $(*, j)$. Plugging this into Table A1 defines an IS proposal distribution for sampling genealogical histories associated with the two loci. Equation 5 then provides an unbiased estimator of the likelihood. We have not yet specified how to compute the numerator in (5); to do this we also need the prior probabilities $\mathbb{P}(\mathcal{H}, \mathcal{T}; \Phi)$. These can be computed easily and sequentially, using the rightmost column of Table A1 (with minor modifications to account for the simulation of $\mathcal{T}$). Table A1 covers events involving coalescence, mutation, and recombination. In addition, we also need to calculate the probability of subsampling the observed haplotypes $\mathcal{D}_{-j}$ from among the lineages in the reconstructed genealogy, for each collection time t_{-j} . Thus, there is a contribution to the numerator of (5) given by the hypergeometric distribution (Beaumont 2003): if $H_{-k} = \mathbf{n}^*$ is the configuration at the collection time (including the additional samples) and $H_{-k+1} = \mathbf{n}$ is the configuration just before (more recently than) the addition of these samples, then

$$\mathbb{P}(H_{-k+1}|H_{-k}) = \frac{\prod_i \binom{a_i^*}{a_i} \prod_j \binom{b_j^*}{b_j} \prod_{i,j} \binom{c_{ij}^*}{c_{ij}}}{\binom{a^*}{a} \binom{b^*}{b} \binom{c^*}{c}}. \quad (12)$$

After decomposing $\mathbb{P}(\mathcal{H}, \mathcal{T}; \Phi)$ in this manner it remains to compute the final term $\mathbb{P}(H_{-m})$ for the probability of the type of the MRCA. Under our uniform mutation model, the stationary distribution of the mutation process and thus the type of the MRCA is uniformly distributed such that $\mathbb{P}(H_{-m}) = 2^{-l}$. Finally, the probability of the data given a history, $\mathbb{P}(\mathcal{D}|\mathcal{H}, \mathcal{T})$, is simply an indicator for whether the configurations at the leaves of the simulated history coincide with the observed data. By construction this is identically equal to 1.

Further algorithmic improvements

In this section, we present two modifications to our model for reducing the computational overhead, as well as two strategies for extending the two-locus model to process multilocus data.

Proposal distribution: The work of Stephens and Donnelly (2000), De Iorio and Griffiths (2004a), Griffiths *et al.* (2008), Paul and Song (2010), and Paul *et al.* (2011) suggests that

our approach to IS will have attractive statistical efficiency. However, generating proposals according to Table A1 requires the evaluation of all possible events for all haplotypes in the current configuration H_{-k} , which may be too costly for complex samples. Instead we therefore make a simple modification, similar to the one suggested by Fearnhead and Donnelly (2001), which is a further approximation to the proposal suggested by Table A1: consider the waiting time to the next event as the minimum of independent competing exponential times for the events involving each of the (i, j) , $(i, *)$, and $(*, j)$ haplotypes. Now sample a haplotype to be involved in the next event according to the *prior* probability of its involvement. In this calculation we exclude the possibility of a coalescence between dissimilar haplotypes, resulting in the probabilities given in Table A2. Next, given the chosen haplotype we choose the event it is involved in with probabilities proportional to the relevant rows in Table A1; now only those rows need to be computed.

Time machine: As reported by Griffiths and Tavaré (1994b), Nielsen (1997), and Jasra *et al.* (2011), as we simulate the tree backward and closer to the MRCA, the simulation times increase greatly. This long simulation run results in undesirably high variance of the likelihood estimate and extensive CPU time—a particular drawback of an IS approach. To circumvent this, Jasra *et al.* (2011) considered stopping the simulation before the MRCA is reached, in an approach termed the Time Machine. The bias from this action is then characterized. First, it depends on the underlying mixing of the evolutionary process: the closer to the root node of the tree, the process is able to “forget” its initial condition. Second, it depends on the specific distribution of the process at the stopping (exit) time. Using the IS algorithm of Stephens and Donnelly (2000), they investigated the bias-variance effect of stopping simulations at the first time back that the number of lineages had decreased to 1%, 2%, 5%, 10%, 25%, and 50% of the original sample size n . They concluded that stopping at 5% strikes the right balance between computational efficiency and numerical accuracy.

We conservatively adopt their approach: stopping at $T_{\text{MRCA}-5}$, the first time the backward simulation reaches five ancestors. At this exit time, the alleles of the remaining lineages are drawn from a uniform distribution. Indeed, we experience sizable computational saving and reduced variance, without noticeable effect on the quality of the estimates.

Estimation of local and global parameters: The method described above estimates the likelihood $L(\Phi)$ for heterochronous deep sequencing data at a pair of loci. Recall that each locus is a stretch of nucleotides within which recombination is *ignored*, but between which recombination is *explicitly modeled*. By searching for Φ that maximizes $L(\Phi)$, we obtain the *pairwise maximum-likelihood estimate* (pairwise MLE) for the mean mutation rate $\bar{\mu}$ and effective population size across the pair and the mean recombination rate $\bar{r}$ between the pair.

This two-locus IS algorithm can be used to analyze multilocus data. Given a region of length κ (typically $\sim 400\text{--}500$ nt), we partition it into loci of length δ (typically ~ 50 nt), defining a collection of $\omega = \kappa/\delta$ loci. We run the IS algorithm on pairs of nonadjacent loci $\{(i, j) : |j - i| > 1; i, j \in \{1, 2, \dots, \omega\}\}$. Adjacent pairs are excluded because recombination between loci separated by a single nucleotide is expected to be negligible. The resulting pairwise MLEs are then indicative of how the local population parameters are distributed within the region.

In practice, the size for the partitioned loci, δ , would be chosen based on the read length distribution of the data, so that a large proportion of the sample at any particular locus will be fully specified and thus can be included in our analysis. (Recall that if a sequence contains even a few missing nucleotides at a given locus, the rest of the information at this locus is discarded and we consider it as a block of missing data.) A balance should be struck: smaller locus length allows inference of finer recombination and mutation rate variation across the region, at the expense of higher computational complexity.

It is also of interest to have a single, global estimate for the population parameters that is representative of the whole region. An ideal situation is a short region (typically ~ 500 nt) where the evolutionary behavior is relatively homogeneous. In such cases, one can take the *median* or *mean* of pairwise MLEs as a simple global estimator; we found the median to be more robust. An alternative, more sophisticated global estimator is via a *pairwise composite likelihood* (reviewed in Larribe and Fearnhead 2011), which we include for comparison with the median of pairwise MLEs. In this strategy, the two-locus IS algorithm is run on each pair $i, j \in \{1, 2, \dots, \omega\}$ of loci such that $1 < |i - j| \leq \Delta$, for some threshold distance Δ . The global parameters are inferred by maximizing pairwise composite likelihood

$$L(\Phi) = \prod_{1 < |i - j| \leq \Delta} \mathbb{P}(\mathcal{D}_{ij}; \Phi), \quad (13)$$

where $\mathcal{D}_{ij}$ denotes the marginal data restricted to the pair of loci i and j . Equation 13 is similar to the composite likelihood of McVean *et al.* (2002), using pairs of multinucleotide loci rather than pairs of single-nucleotide polymorphisms (SNPs), as suggested by Jenkins and Griffiths (2011).

Introducing a threshold Δ is attractive for both statistical and computational reasons (Larribe and Fearnhead 2011), and there are several additional reasons why Δ should be small in our application. First, the short-read nature of the data means that if $|i - j| > \Delta$, then many samples will be missing at one or both of the two loci. Second, since our loci represent blocks of nucleotides rather than single SNPs, parameter variation along the sequence implies we should concentrate on proximate pairs of loci. Third, our modeling of complete blocks of nucleotides instead of isolated SNPs prohibits the precomputation of an exhaustive list of pairwise

likelihoods, as is possible in McVean *et al.* (2004), for example. For these reasons we focus on the composite-likelihood (13) in which $\Delta = 2$.

These procedures are illustrated in Figure 4.

Simulated data

Simulated heterochronous data sets were generated using the software package NetRecodon (Arenas and Posada 2010), using a Jukes–Cantor substitution model (Jukes and Cantor 1969) and with each nucleotide frequency value equal to 0.25. Population parameters were chosen to be comparable to those previously described for HIV: $\bar{\mu} = 2.5 \times 10^{-5}$ per site per generation, $\bar{\tau} = 10^{-6}$ per site per generation, $N_e = 10^3$, $\tau = 1.8$ days (e.g., Shriner *et al.* 2004; Neher and Leitner 2010; Batorsky *et al.* 2011). Each simulation was set to produce n_c sequences of length $\kappa = 500$ nt, at four collection times separated by 252 days, totaling $n = 4n_c$ sequences.

Under the above conditions, four data sets were generated under slightly different scenarios. The first pair of data sets, called *Const-120* and *Const-600*, were generated under a constant population size. Const-120 had $n_c = 30$ reads sampled at each collection time, totaling $n = 120$ reads. Const-600 had a fivefold increased sample size to $n_c = 120$ reads at each collection time, totaling $n = 600$ reads. The second pair of data sets, *Dynamic-120* and *Dynamic-600*, had the same population parameters and samples sizes as Const-120 and Const-600, respectively. However, they were generated under a fluctuating viral population such that N_e corresponded to 1000, 2000, 4000, and 500 at the four collection times (from the most recent to the earliest). This models a population undergoing both an expansion and a bottleneck through time.

Figure 4 Estimation of the local and global parameter estimates using the two-locus IS algorithm. A region of interest, typically $\sim 4\text{--}500$ nt, is shown as a blue segment. This region is partitioned into smaller loci, e.g., 50 nt. Sequence reads are shown as thinner horizontal bars. For non-adjacent pairs of loci, the two-locus engine computes pairwise MLEs as local estimates of the population parameters. Here, a pairwise comparison between loci 2 and 4 is illustrated (yellow shading). Reads that fully cover at least one of the two loci are highlighted in red and are used for the inference: three complete haplotypes (reads 1, 4, and 6) and two partial haplotypes (reads 2 and 3). These pairwise inferences can be combined to reach the global parameter for the whole region. Two approaches are described in the text: by taking the median of the pairwise MLEs or via a pairwise composite likelihood.

Real data

The subject studied here was enrolled in the control (no therapy) arm of the SPARTAC study. EDTA plasma samples were obtained at the start of the study (at an estimated 12 days from HIV-1 seroconversion) and at 28, 120, 176, 373, 429, and 695 days after the start of the study. Viral RNA was extracted, whole HIV-1 genomes were sequenced in a pool of 21 samples on a 1/2 PicoTiterPlate using a Genome Sequencer FLX TitaniumXL instrument (Roche/454Life Sciences), and a consensus sequence for each time point was derived as previously described (Gall *et al.* 2012). The seven consensus sequences were aligned using MAFFT v6.857 (Katoh *et al.* 2002), and a consensus sequence of this alignment was used as a reference sequence for mapping of all reads, with suffixes representing the time points, using Burrows-Wheeler Aligner v0.5.9 (Li and Durbin 2010). The resulting SAM file was converted into a FASTA file, using a custom Java script. The distribution of read lengths and positions across the genome are given in Supporting Information, Table S1 and Table S2, respectively.

HIV genome analysis: Nine regions were selected from the whole HIV genome alignment of mapped reads, each region ~600 nt in length and comprising samples from all time points. Furthermore, we ensured that these regions were in frame and nonoverlapping, spanning the four main open-reading frames (*i.e.*, gag, pol, env, and nef). Reads covering at least 95% of a region were retained in our analysis. The average number of sequences per time point was 34 (range 20–65). Further details about the sequence data set for each region can be found in Table S3.

BEAST analysis: To evaluate the performance of Coalescenator in estimating population parameters of interest, primarily the effective population size, N_e , and mutation rate, $\bar{\mu}$, we reanalyzed the sequence alignments for the nine regions described above, using the Bayesian MCMC coalescent approach implemented in BEAST (v1.8.0) (Drummond and Rambaut 2007; Drummond *et al.* 2012). Specifically, a codon-structured nucleotide substitution model (Shapiro *et al.* 2006), a constant-size population demographic model, and a strict molecular clock prior were employed. Two independent MCMC chains of 50 million steps were run to assess convergence and adequate mixing.

Accession numbers for sequencing data: The Roche/454 Life Sciences sequencing data obtained in this study are available from the EMBL/NCBI/DDBJ Sequence Read Archive sample accession nos. ERS661087–ERS661093.

Results

Simulated data

To evaluate the validity of our model implementation, Coalescenator, we simulated heterochronous data sets under

known population parameters. Employing this set, we assessed Coalescenator's performance with respect to the number of Monte Carlo iterations, the sample sizes of the data, and the data's underlying population size dynamics. For each study, we searched through a three-dimensional parameter space of 11 $\bar{\mu}$'s, 11 $\bar{r}$'s, and 6 N_e 's that best explain the observed diallelic sites in the data (see Table S4 for the summary of the sites classification and the numbers of sites used). We assessed whether parameter combinations with the highest likelihoods corresponded to the true parameters that generated the data. A total of $11 \times 11 \times 6 = 726$ parameter combinations of the following were used:

$$\begin{aligned} \bar{\mu} : & 1 \times 10^{-7}, 1 \times 10^{-6}, (1, 2.5, 5, 7.5) \times 10^{-5}, (1, 2.5, 5, 7.5) \\ & \times 10^{-4}, 1 \times 10^{-3}. \\ \bar{r} : & 1 \times 10^{-8}, (1, 2.5, 5, 7.5) \times 10^{-7}, (1, 2.5, 5, 7.5) \times 10^{-6}, \\ & 1 \times 10^{-5}, 1 \times 10^{-4}. \\ N_e : & 250, 750, 1000, 1250, 1750, 3000. \end{aligned}$$

Performance for data from constant population size: The behavior of Coalescenator when analyzing data generated under constant population size was assessed using data sets Const-120 and Const-600 (see *Materials and Methods*). Coalescenator was run with a constant population size model.

First, the performance of local inference using pairwise MLEs was assessed by analyzing a single pair of loci of lengths $l_A = l_B = 50$ nt at positions (1, 50) and (101, 150). Figure 5 shows the likelihood surfaces for Const-120, with 100 and 500 Monte Carlo iterations, respectively; a further increase to 1000 and 5000 iterations is shown in Figure S1. Here, across various choices of N_e , there is a clear clustering of higher likelihoods near the true parameter values. With only 100 iterations, the pairwise MLE recovers the correct mutation rate (among the 11 parameter values considered) and estimates the recombination rate at $\hat{r} = 10^{-7}$. The MLE for the effective population size is farther off target, at $\hat{N}_e = 3000$. By increasing the number of iterations to 500 (Figure 5B), there was a marked improvement in the likelihood surface; it is now much more peaked, and the recombination rate estimate improves to within a factor of 2 of its true value (although for this experiment the effective population size is still overestimated). Overall, parameter estimates are accurate for 500 iterations and certainly to within an order of magnitude. For this data set, we also investigated a more ambitious increase to 1000 and 5000 Monte Carlo iterations (Figure S1); the recombination rate estimate sees modest further improvement, but at a cost of infeasible computing requirements should we hope to scale this to real data sets. Indeed, it is worth emphasizing that these experiments use relatively few Monte Carlo iterations; typical IS approaches to coalescent inference might use several orders of magnitude more. Our aim here is to stretch our algorithm when only a few iterations are available, so that it can feasibly be scaled up to complex multilocus data sets. The results presented here suggest that the

Figure 5 Likelihood surfaces on a pair of loci at positions (1, 50), (101, 150), for dataset Const-120, using 100 and 500 Monte Carlo iterations. Cells correspond to the searched parameters, colored by log-likelihoods, with the top 10 estimates numbered. The true mutation, recombination, and population parameters are $\bar{\mu} = 2.5 \times 10^{-5}$, $\bar{r} = 10^{-6}$, $N_e = 10^3$. (A) One hundred twenty sequences (Seqs), 100 iterations. (B) One hundred twenty Seqs, 500 iterations.

Figure 6 As in Figure 5 but for data set Const-600. (A) Six hundred Seqs, 100 iterations. (B) Six hundred Seqs, 500 iterations.

sophistication of our algorithm adequately compensates for the dearth of CPU time, and comparisons with longer runs show that even as few as 100 iterations give decent estimates.

The nature of the data, with its high mutation rates and heterochronicity, may also assist in this exploration of tree space.

Figure 7 Population parameter estimates for simulated data sets generated under a constant population size model. The effect of using different combinations of sample size and number of Monte Carlo iterations is compared. Circles correspond to the pairwise MLEs between neighboring pairs of loci (those separated by 50 nt). The horizontal lines correspond to the median of the pairwise MLEs. Crosses indicate the pairwise composite-likelihood estimates. (A) Mutation and recombination rate estimates, per site per generation (1.8 days). (B) Effective population size estimates.

Figure 6 shows the corresponding likelihood surfaces for Const-600. The likelihood hypersurface concentrates in a similar manner with respect to the true parameter values, but this time there is a shift toward the correct effective population size. This suggests that Coalescentor performs favorably with increasing sample size.

It is also worth commenting on the general shape of the likelihood hypersurfaces in Figure 5 and Figure 6. One might expect to see a ridge corresponding to a negative correlation in the estimation of mutation and recombination rates, as alternative ways of resolving any incompatibilities in the data. In fact the main ridge, evident from several panels in Figure 5 and Figure 6, is parallel to the recombination axis. This reflects the simple fact that recombination is inherently more difficult than mutation to estimate to a given accuracy. A second, weaker ridge is evident in the contours along lines of constant $N_e\mu$, as might be expected. That no other “ridges” seem to be discernible from these likelihood hypersurfaces is encouraging, as it is consistent with the fact that the data contain sufficient information for mutation and recombina-

tion parameters to be estimated separately (and that our method can reconstruct the likelihood surface to sufficient accuracy).

The above analysis shows that by using the information from a single pair of loci of 50 nt each, Coalescentor is able to arrive at a sensible indication of the population parameters for the whole 500-nt region. This raises two questions:

1. For effective inference, how far apart should the two loci be?
2. A pairwise MLE provides only partial information since it utilizes a subset of the whole region. How, then, should one come up with the global estimate for the whole region?

To answer the first question, inference quality was assessed with respect to the separation d of a pair of loci. The 500-nt sequence region was divided into 10 neighboring loci of 50 nt, and for each of the 36 pairs of nonadjacent loci, pairwise MLEs for the population parameters were calculated. Figure S2 shows how the MLEs vary with d .

Figure 8 Population parameter estimates for simulated data sets generated under a dynamic population size model. Coalescentor was run under a constant population model, and this analysis shows its robustness to unmodeled changes in population size. Circles correspond to the pairwise MLEs between neighboring pairs of loci (those separated by 50 nt). The horizontal lines correspond to the median of the pairwise MLEs. Crosses indicate the pairwise composite-likelihood estimates. (A) Mutation and recombination rate estimates, per site per generation (1.8 days). (B) Estimates for effective population size.

We observed that the mutation rate estimate is stable across separation distance. However, this is not the case with the recombination estimates. Specifically, (i) variance in MLEs for fixed d is greater than for mutation, and (ii) with increasing d , recombination rate estimates (per site) seem to be lower.

These two observations could have been anticipated to some extent; the signal for recombination in genetic data is generally weaker than that of mutation. Furthermore, with increasing separation between two loci the signal for recombination can become “saturated,” so that recombination events are undetected and recombination rates are underestimated. More precisely, the curvature of the likelihood

curve for ρ is flatter when ρ is greater (Chan *et al.* 2012). Ultimately, this provides the basis in deciding the appropriate constraint on locus pair separation for reliable inference. This is ~ 50 – 100 nt, and from here onward we proceeded by computing pairwise MLEs for pairs of loci separated by 50 nt.

To address the second challenge of obtaining global parameter estimates, we assessed both of the two approaches discussed in the section *Estimation of local and global parameters*. The first approach is to take the median of the pairwise MLEs. Figure 7 shows pairwise MLEs and the corresponding median across the data sets. We see that the median can effectively capture the true estimates, albeit with a slight overestimation of N_e . There is also greater variance in

estimates for the recombination rates, which in turn reflect the confidence we should have in these estimates.

The second approach is via a pairwise composite likelihood (Equation 13), in which likelihood surfaces for pairs of loci separated by 50 nt are multiplicatively combined. Figure S3 shows the composite-likelihood surfaces for data set Const-120. The MLEs using 100 and 500 Monte Carlo iterations are both of high accuracy and differ only slightly in the recombination rate estimates; this further illustrates the fast convergence of Coalescentor. With 500 iterations, the true mutation and recombination rates are captured, although the effective population size is off target at $\hat{N}_e = 3000$. Figure S4 shows the result for Const-600, where the larger sample size also provides accurate estimates of the effective population size. For example, with 500 iterations the joint MLE is $(\hat{\mu}, \hat{r}, \hat{N}_e) = (5 \times 10^{-5}, 10^{-6}, 1250)$, only two cells away from the true values (on this grid). These estimates are also included in Figure 7 for direct visual comparison with the first approach.

Performance for data from dynamic population sizes: The behavior of Coalescentor for analyzing data generated under dynamic population sizes was assessed using data sets Dynamic-120 and Dynamic-600. Figure 8 shows the pairwise MLEs and the global estimates based on the median and pairwise composite likelihoods, for both data sets. Each estimate was based on 100 Monte Carlo iterations.

With the incorrect assumption of a constant population size, Coalescentor slightly overestimates the mutation rates. The pairwise MLEs for recombination rates increase in accuracy for the larger sample size, but there is greater variability compared to the analysis in the previous section on Const-120 and Const-600. Effective population size estimates are still of very reasonable accuracy and stability. This suggests that Coalescentor is reasonably robust to unmodeled changes in population size and can be reliably used in this way. Of course, it is straightforward to extend our algorithm to impose population size changes between sample collection times, perhaps guided by viral load information, and even to *infer* changes in effective population size. However, in the latter case we found that our chosen ranges of sample size and Monte Carlo iterations were insufficient to compensate for the significant increase in model complexity.

As all the cases here and in the previous section illustrate (Figure 7 and Figure 8), the two methods for providing global parameter estimates—taking the median of pairwise MLEs and constructing a composite-likelihood estimate—have generally shown good agreement, which corroborates each other’s accuracy. The former seems to give similar or better global measures than the latter, which seems slightly more susceptible to outlier estimates, so we focus on the median of pairwise MLEs to provide our global summaries from here onward. Clearly, strategies more sophisticated than a simple median might be able to aggregate pairwise estimates more efficiently, and we leave this for future work.

Table 1 Assessment of the variation of the importance weights, under different simulated data sets

Data set	CV threshold				
	<0.15	<0.1	<0.05	<0.01	<0.001
Const-120	1	1	0.92	0.83	0.83
Const-600	1	0.86	0.86	0.86	0.86
Dynamic-120	1	1	0.99	0.95	0.95
Dynamic-600	1	0.91	0.9	0.88	0.88

For each data set, a CV of the importance weights was calculated for each run of 100 Monte Carlo iterations on a particular parameter combination and a pair of loci separated by 50 nt. This results in 5808 CVs being computed in total for each data set. Each cell in the table shows the proportion of these CVs that is below the indicated threshold.

Variance of the importance weights: For each of the four simulated data sets, we assessed the stability of Coalescentor’s likelihood estimates by computing a coefficient of variation $CV(= \sqrt{\text{Variance}/\text{Mean}})$ of the importance weights [*i.e.*, the summands in (5)]. We used 100 Monte Carlo iterations for each particular parameter combination, analyzing each pair of loci separated by 50 nt. There are $11 \times 11 \times 6 = 726$ parameter combinations and eight pairs of loci, resulting in 5808 CVs being computed for each data set.

Table 1 summarizes the distribution of the coefficient of variation for the four simulated data sets. The variance appears to be well constrained, with the majority of CVs <0.001 . The larger CVs originate from parameter combinations that deviate farthest from the ground truth parameters. While a small CV (equivalently, a large effective sample size) is not a guarantee of good performance, it is consistent with the idea that the proposal distribution is not too skewed in its exploration of the space of ARGs.

Running time: Coalescentor scales approximately linearly with respect to the sample size and the number of Monte Carlo iterations. For data set Const-120, a likelihood calculation for a single parameter combination using 100 iterations takes on average 10 sec when run in parallel on a 16-core computer with a 2.0-GHz processor. For data set Const-600, the corresponding calculation increases fivefold to about 50 sec.

Real data analysis

Coalescentor was run on nine HIV regions from HIV genome alignments collected at seven time points from a patient infected with HIV over a period of 2 years (see *Materials and Methods*). The sequence data are summarized in Table S3.

Each gene region is 600 nt long, but some reads had missing data, chiefly up to the first 30 nt and the last 30 nt of the region. Although Coalescentor can handle missing data, for the purpose of this analysis we concentrate on the central positions (51, 550) of each gene region. Sites containing ambiguous nucleotides, indels, and nondiallelic sites were filtered out. Table S5 summarizes the sites classification, as well as the total and breakdown of the number of sites used for each gene region.

Figure 9 Population parameter estimates for nine HIV gene regions. Data are from HIV genome alignments collected at seven time points over the period of 2 years. Circles correspond to the pairwise MLEs between neighboring pairs of loci separated by 50 nt. The horizontal lines correspond to the median of the pairwise MLEs. Crosses indicate the pairwise composite-likelihood estimates. (A) Mutation and recombination rate estimates, per site per generation (1.8 days). (B) Effective population size estimates.

The parameter combinations extended those used in the simulation study. Additional grid points $(2.5, 5, 7.5) \times 10^{-5}$ for the recombination rate were considered, to account for the observation in a preliminary run that several MLEs for recombination rates across the regions seem to be between 1×10^{-5} and 1×10^{-4} . Following the simulation study, each gene region was partitioned into ten 50-nt loci, and likelihood surfaces were generated for pairs of loci separated by 50 nt, using 100 iterations.

Figure 9 shows the evolutionary parameter estimates across each region. The median of the pairwise MLEs and the pairwise composite-likelihood estimates agree quite closely for the mutation and effective population size estimates, and all lie in the

range $[2.5 \times 10^{-5}, 1 \times 10^{-4}]$ for mutation rate estimates in each region. There was greater variability across the genome for effective population size estimates: for example, all median-based estimates were in the range 875–1000 in regions encompassing env, while the median-based estimate for the region pol 2836–3436 had $\hat{N}_e = 1500$. There was some disagreement between the two estimates of recombination rate. The pairwise composite-likelihood estimates seem to be less stable; the estimate in the region env 6415–7015 is 2–10 times lower than in any other region and 1/10th of the estimate obtained using the median of pairwise MLEs. (As discussed above, this behavior was also evident in the simulation study.) We prefer therefore to use the median of the pairwise MLEs as our regional

Figure 10 Likelihood surfaces for the env 6415–7015 region, using 100 Monte Carlo iterations. Cells correspond to the searched parameters, colored by log-likelihoods, with the top 10 estimates numbered. (A) Likelihood surface for a single pair of loci within the region. (B) Pairwise composite likelihood aggregating all valid pairs of loci within the region.

summaries, with the uncertainty expressed by the variability in individual MLEs.

It is worth noting that some pairwise MLEs are suggestive of a higher mutation rate in parts of the env region (described previously by Alizon and Fraser 2013). There is also some variability in the median-based recombination rate estimates across regions, but this is consistent with Monte Carlo sampling error. A single pairwise likelihood surface between loci (1, 50) and (101, 151) and a pairwise composite-likelihood surface is shown in Figure 10 for the env 6415–7015 region. This serves as a sanity check that we indeed have convergence to a plausible likelihood surface, since each grid point is simulated independently.

Finally, we compared our estimates against the output of the corresponding analysis using BEAST (Drummond and Rambaut 2007; Drummond *et al.* 2012). (We do not compare recombination rates—recall that BEAST is unable to infer these. Our recombination rate estimates are summarized separately in Table 2.) The comparison is illustrated in Figure 11. The raw values for all the estimates are reported in detail in Table S6. The mutation μ and recombination $\bar{\rho}$ rate estimates from Coalescentor were converted into the number of substitutions and recombinations per site per year, respectively. N_e denotes effective population size, and the time to the most recent common ancestor, T_{MRCA} , or the time until we reach five ancestors, T_{MRCA-5} , is given in years. Estimates of mutation rates and T_{MRCA} 's show very good agreement, although our method seems to be more cautious in estimating variation in N_e across regions. Across the mutation rate estimates, the only notable differences are our higher estimates for the second gag region and the first pol region (each greater by a factor of ≈ 4). This discrepancy could reflect that Coalescentor can be prone to slight overestimation of mutation rates, as suggested by the simulation results on Dynamic-120 and Dynamic-600.

Discussion

We have presented a method capable of handling many of the challenges associated with modeling evolution of rapidly evolving viral populations measured using high-throughput sequencing. It is based on recent developments in the approximation of CSDs (Stephens and Donnelly 2000; De Iorio and Griffiths 2004a; Griffiths *et al.* 2008; Paul and Song 2010; Paul *et al.* 2011), and the adoption of this work into a practical IS algorithm should be a useful contribution in itself. The model we derive handles recombination by allowing for partially specified haplotypes and thus avoids the inflation of the model state space caused by alternative imputation procedures. It also provides a natural way of dealing with missing data. One limitation is that data are lost when we simply consider an entire locus as unobserved if any nucleotides are missing; another possibility for future research would be to impute short stretches of missing nucleotides within a locus.

High-throughput sequence data from viral genomes present many challenges. We have focused in particular

Table 2 Coalescentor's recombination rate $\bar{\rho}$ estimates, given in the number of recombinations per site per year

Region	$\bar{\rho}$
gag 311–940	3.04×10^{-3}
gag 960–1560	5.07×10^{-3}
pol 2005–2605	1.27×10^{-2}
pol 2836–3436	5.07×10^{-3}
pol 3796–4396	5.07×10^{-3}
env 5812–6412	2.03×10^{-3}
env 6415–7015	5.07×10^{-3}
env 7357–7957	5.07×10^{-3}
nef 8376–9011	3.55×10^{-3}

on handling high mutation rates, recombination, heterochronous sampling, and missing data. Each of these presents its own hurdles; dealing with them simultaneously raises numerous further complications. In this article we have therefore also explored several statistical and algorithmic techniques for reducing the large computational overhead. In particular, employing the Time Machine strategy developed by Jasra *et al.* (2011) proves to be effective in saving computational time, while controlling the bias and variance. Nevertheless, the overhead remains large, and it is likely that further implementation strategies will have to be explored in the future as sample sizes continue to grow. For example, one advantage of an IS approach is that it is especially suited for parallelization using the rapidly growing graphics processing unit (GPU) (Lee *et al.* 2010). Lee *et al.* (2010) focused on MCMC and IS as case studies and concluded that significant performance boost can be achieved for IS using GPU, whereas BEAST's MCMC is inherently less suitable for exploitation using GPU's massive parallelization.

In spite of these challenges, it is encouraging that our new method performs comparably to BEAST in the analysis of a deep sequencing data set generated from a serially sampled acute HIV infection. Like previous studies based on coalescent inference, BEAST and Coalescentor also indicate that the effective population size is of the order $\sim 10^3$ (Brown 1997). While there has been great interest in measuring the effective population size of within-host HIV infections, especially in regard to addressing whether evolution is largely stochastic or deterministic (Rouzine *et al.*, 2014), it is important to note that estimates of effective population size from methods based on the coalescent framework do not typically consider natural selection or other factors that may lead to variation in offspring distribution. Consequently, in populations where variation in reproductive success is likely to be present, such as in within-host HIV infections, interpretation of the effective population size is confounded. Instead, a more accurate interpretation of effective population size is that it represents a measure of population turnover or relative genetic diversity. We further note that in this article we have focused our analyses on inference of a single effective population size parameter, while a priority for future work will be also to make

Figure 11 Comparison of BEAST's parameter estimates against Coalescentor's. An identity line $y = x$ is plotted in each case, to visually assess the agreement of the estimates by the two programs. The following abbreviations are used for the nine HIV gene regions: gag-1 = gag 311–940, gag-2 = gag 960–1560, pol-1 = pol 2005–2605, pol-2 = pol 2836–3436, pol-3 = pol 3796–4396, env-1 = env 5812–6412, env-2 = env 6415–7015, env-3 = env 7357–7957, nef = nef 8376–9011. (A) Comparison of mutation rate $\bar{\mu}$ estimates, converted to the number of substitutions per site per year. (B) Comparison of effective population size N_e estimates. (C) Comparison of BEAST's time to the most recent common ancestor T_{MRCA} estimates, against Coalescentor's time until five ancestors is reached, T_{MRCA-5} estimates, given in years.

inference of temporal changes in this parameter computationally tractable.

The estimates of recombination rate from Coalescensor strongly agree with non-coalescent-based estimates obtained by Neher and Leitner (2010) and Batorsky *et al.* (2011). This is encouraging as the sequence data analyzed in this study have been sampled from an early phase of infection, in contrast to previous studies that utilized data from patients with a much longer follow-up (Neher and Leitner 2010; Batorsky *et al.* 2011). Consequently, this demonstrates the applicability of our method even in spite of limited accumulation of diversity in the viral population.

Despite the complicated evolutionary model considered in this article, we have omitted other mechanisms that are likely to be important to within-host HIV evolution. Our simplified model of population size changes provides a restrictive model for the phases of exponential expansion and contraction HIV and other viral populations are known to go through (Nowak and May 2000). A natural solution to this problem would be to fit exponential functions between consecutive time points. More importantly, however, selective responses to the host immune system are known to be an important evolutionary factor (Edwards *et al.* 2006; Lemey *et al.* 2006), having an effect on intrahost genealogies that is not fully captured merely by adjusting the effective population size. Other commonly used methods such as BEAST (Drummond *et al.* 2002, 2005, 2012; Drummond and Rambaut 2007; Minin *et al.* 2008) also ignore the effects of selection on genealogies (as well as ignoring recombination); thus, developing methods that can robustly account for both selection and recombination (and indeed other factors) remains a challenging task. Another important consideration is population substructure: HIV is known to undergo *compartmentalization*, forming distinct subpopulations at different anatomical sites and in different cell types (Ewing *et al.* 2004). Incorporating substructure and migration into IS algorithms is in principle straightforward and has been achieved in other applications (Bahlo and Griffiths 2000; De Iorio and Griffiths 2004b; Griffiths *et al.* 2008), although it obviously introduces yet more computational burden.

Exploring more complicated mutation models would also be interesting, although the diallelic model we consider here allows for a simple analytical solution, which in turn allows for more efficient computation. However, it should be straightforward to allow for time-varying mutation rates as a proxy for the effects of selection. Inclusion of indels into the mutation model would also be desirable; inspection of the HIV data (*cf.* Table S5) shows that these are highly prevalent. The model also considers only crossover recombination; however, McVean *et al.* (2002) suggested that a gene-conversion model may be appropriate for bacteria and viruses. Incorporating this type of recombination would also be a worthwhile extension.

Finally, we remark that in our analysis we tuned our algorithm to analyze data generated under the 454 platform. However, it should be straightforward to adapt our work for

other platforms such as Illumina paired-end sequencing, in which a haplotype would now comprise a pair of reads. Since the number of reads generated by this method is typically severalfold larger, this raises further computational and statistical challenges that form the basis of ongoing work.

Acknowledgments

This research was initiated during the 2013 Oxford Summer School in Computational Biology. We gratefully acknowledge Sarah Fidler, Steve Kaye, Jonathan Weber, Myra McClure, and the SPARTAC Trial participants and investigators. We thank James Anderson, Farah Colchester, Pierre Haas, Thomas Leitner, István Miklós, and Yee Whye Teh for helpful comments on this work. We acknowledge the use of the University of Oxford Advanced Research Computing facility in carrying out this work (<http://dx.doi.org/10.5281/zenodo.22558>). This work was supported by the Wellcome Trust and the National Institute for Health Research (NIHR) Biomedical Research Centre funding scheme at Imperial College Healthcare National Health Service (NHS) Trust and University College London Hospitals NHS Foundation Trust. We also acknowledge funding from the Novo Nordisk Foundation (to J.A.S. and L.M.) and the Engineering and Physical Sciences Research Council (EPSRC) (to P.A.J., grant EP/L018497/1). J.R. is supported by the Oxford Martin School. O.G.P. received funding from the European Research Council (ERC) under the European Union's Seventh Framework Programme (FP7/2007-2013)/ERC grant agreement 614725-PATHPHYLODYN.

Literature Cited

- Abramowitz, M., and I. Stegun (Editors), 1972 *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables* (National Bureau of Standards Applied Mathematics Series, Ed. 10, Vol. 55). Dover Publications, New York, USA.
- Alizon, S., and C. Fraser, 2013 Within-host and between-host evolutionary rates across the HIV-1 genome. *Retrovirology* 10(1): 49.
- Anderson, E. C., 2005 An efficient Monte Carlo method for estimating N_e from temporally spaced samples using a coalescent-based likelihood. *Genetics* 170: 955–967.
- Archer, J., J. W. Pinney, J. Fan, E. Simon-Loriere, E. J. Arts *et al.*, 2008 Identifying the important HIV-1 recombination breakpoints. *PLoS Comput. Biol.* 4: e10000178.
- Arenas, M., and D. Posada, 2010 Coalescent simulation of intra-codon recombination. *Genetics* 184: 429–437.
- Bahlo, M., and R. Griffiths, 2000 Inference from gene trees in a subdivided population. *Theor. Popul. Biol.* 57: 79–95.
- Batorsky, R., M. F. Kearney, S. E. Palmer, F. Maldarelli, I. M. Rouzine *et al.*, 2011 Estimate of effective recombination rate and average selection coefficient for HIV in chronic infection. *Proc. Natl. Acad. Sci. USA* 108: 5661–5666.
- Beaumont, M. A., 1999 Detecting population expansion and decline using microsatellites. *Genetics* 153: 2013–2029.
- Beaumont, M. A., 2003 Estimation of population growth or decline in genetically monitored populations. *Genetics* 164: 1139–1160.

- Brown, A. J. L., 1997 Analysis of HIV-1 *env* gene sequences reveals evidence for a low effective number in the viral population. *Proc. Natl. Acad. Sci. USA* 94: 1862–1865.
- Chan, A. H., P. A. Jenkins, and Y. S. Song, 2012 Genome-wide fine-scale recombination rate variation in *Drosophila melanogaster*. *PLoS Genet.* 8: e1003090.
- De Iorio, M., and R. C. Griffiths, 2004a Importance sampling on coalescent histories I. *Adv. Appl. Probab.* 36: 417–433.
- De Iorio, M., and R. C. Griffiths, 2004b Importance sampling on coalescent histories II. *Adv. Appl. Probab.* 36: 434–454.
- Drummond, A. J., and A. Rambaut, 2007 BEAST: Bayesian evolutionary analysis by sampling trees. *BMC Evol. Biol.* 7: 214.
- Drummond, A. J., G. K. Nicholls, A. G. Rodrigo, and W. Solomon, 2002 Estimating mutation parameters, population history and genealogy simultaneously from temporally spaced sequence data. *Genetics* 161: 1307–1320.
- Drummond, A. J., O. G. Pybus, A. Rambaut, R. Forsberg, and A. G. Rodrigo, 2003 Measurably evolving populations. *Trends Ecol. Evol.* 18: 481–488.
- Drummond, A. J., A. Rambaut, B. Shapiro, and O. G. Pybus, 2005 Bayesian coalescent inference of past population dynamics from molecular sequences. *Mol. Biol. Evol.* 22: 1185–1192.
- Drummond, A. J., M. A. Suchard, D. Xie, and A. Rambaut, 2012 Bayesian phylogenetics with BEAUti and the BEAST 1.7. *Mol. Biol. Evol.* 29: 1969–1973.
- Edwards, C. T. T., E. C. Holmes, O. G. Pybus, D. J. Wilson, R. P. Viscidi *et al.*, 2006 Evolution of the human immunodeficiency virus envelope gene is dominated by purifying selection. *Genetics* 174: 1441–1453.
- Ethier, S. N., and R. C. Griffiths, 1990 On the two-locus sampling distribution. *J. Math. Biol.* 29: 131–159.
- Ewing, G., G. Nicholls, and A. Rodrigo, 2004 Using temporally spaced sequences to simultaneously estimate migration rates, mutation rate and population sizes in measurably evolving populations. *Genetics* 168: 2407–2420.
- Fan, J., M. Negroni, and D. L. Robertson, 2007 The distribution of HIV-1 recombination breakpoints. *Infect. Genet. Evol.* 7: 717–723.
- Fearnhead, P., 2008 Computational methods for complex stochastic systems: a review of some alternatives to MCMC. *Stat. Comput.* 18: 151–171.
- Fearnhead, P., and P. Donnelly, 2001 Estimating recombination rates from population genetic data. *Genetics* 159: 1299–1318.
- Gall, A., B. Ferns, C. Morris, S. Watson, M. Cotten *et al.*, 2012 Universal amplification, next-generation sequencing, and assembly of HIV-1 genomes. *J. Clin. Microbiol.* 50: 3838–3844.
- Gall, A., S. Kaye, S. Hué, D. Bonsall, R. Rance *et al.*, SPARTAC Trial Investigators, 2013 Restriction of V3 region sequence divergence in the HIV-1 envelope gene during antiretroviral treatment in a cohort of recent seroconverters. *Retrovirology* 10: 8.
- Grenfell, B. T., O. G. Pybus, J. R. Gog, J. L. N. Wood, J. M. Daly *et al.*, 2004 Unifying the epidemiological and evolutionary dynamics of pathogens. *Science* 303: 327–332.
- Griffiths, R. C., and P. Marjoram, 1996 Ancestral inference from samples of DNA sequences with recombination. *J. Comput. Biol.* 3: 479–502.
- Griffiths, R. C., and S. Tavaré, 1994a Sampling theory for neutral alleles in a varying environment. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 344: 403–410.
- Griffiths, R. C., and S. Tavaré, 1994b Simulating probability distributions in the coalescent. *Theor. Popul. Biol.* 46: 131–159.
- Griffiths, R. C., P. A. Jenkins, and Y. S. Song, 2008 Importance sampling and the two-locus model with subdivided population structure. *Adv. Appl. Probab.* 40: 473–500.
- Hein, J., M. H. Schierup, and C. Wiuf, 2005 *Gene Genealogies, Variation and Evolution*. Oxford University Press, London/New York/Oxford.
- Henn, M. R., C. L. Boutwell, P. Charlebois, N. J. Lennon, K. A. Power *et al.*, 2012 Whole genome deep sequencing of HIV-1 reveals the impact of early minor variants upon immune recognition during acute infection. *PLoS Pathog.* 8: e1002529.
- Jasra, A., M. De Iorio, and M. Chadeau-Hyam, 2011 The time machine: a simulation approach for stochastic trees. *Proc. R. Soc. Math. Phys. Eng. Sci.* 467: 2350–2368.
- Jenkins, P. A., and R. C. Griffiths, 2011 Inference from samples of DNA sequences using a two-locus model. *J. Comput. Biol.* 18: 109–127.
- Jenkins, P. A., and Y. S. Song, 2009 Closed-form two-locus sampling distributions: accuracy and universality. *Genetics* 183: 1087–1103.
- Jukes, T. H., and C. R. Cantor, 1969 Evolution of protein molecules, pp. 21–132 in *Mammalian Protein Metabolism*, Vol. III H. N. Munro (editor). Academic Press, New York.
- Katoh, K., K. K. Misawa, K. Kuma, and T. Miyata, 2002 MAFFT: a novel method for rapid multiple sequence alignment based on fast Fourier transform. *Nucleic Acids Res.* 30: 3059–3066.
- Kellam, P., and B. A. Larder, 1995 Retroviral recombination can lead to linkage of reverse transcriptase mutations that confer increased zidovudine resistance. *J. Virol.* 69: 669–674.
- Kuhner, M. K., J. Yamato, and J. Felsenstein, 2000 Maximum likelihood estimation of recombination rates from population data. *Genetics* 156: 1393–1401.
- Larribe, F., and P. Fearnhead, 2011 On composite likelihoods in statistical genetics. *Stat. Sin.* 21: 43–69.
- Leblois, R., P. Pudlo, J. Néron, F. Bertaux, C. R. Beeravolu *et al.*, 2014 Maximum likelihood inference of population size contractions from microsatellite data. *Mol. Biol. Evol.* 31: 2805–2823.
- Lee, A., C. Yau, M. B. Giles, A. Doucet, and C. C. Holmes, 2010 On the utility of graphics cards to perform massively parallel simulation of advanced Monte Carlo methods. *J. Comput. Graph. Stat.* 19: 769–789.
- Lemey, P., A. Rambaut, and O. G. Pybus, 2006 HIV evolutionary dynamics within and among hosts. *AIDS Rev.* 8: 125–140.
- Lemey, P., S. L. K. Pond, A. J. Drummond, O. G. Pybus, B. Shapiro *et al.*, 2007 Synonymous substitution rates predict HIV disease progression as a result of underlying replication dynamics. *PLoS Comput. Biol.* 3: e29.
- Li, H., and R. Durbin, 2010 Fast and accurate long-read alignment with Burrows-Wheeler transform. *Bioinformatics* 26: 589–595.
- McVean, G., P. Awadalla, and P. Fearnhead, 2002 A coalescent-based method for detecting and estimating recombination from gene sequences. *Genetics* 160: 1231–1241.
- McVean, G. A. T., S. R. Myers, S. Hunt, P. Deloukas, D. R. Bentley *et al.*, 2004 The fine-scale structure of recombination rate variation in the human genome. *Science* 304: 581–584.
- Minin, V. N., E. W. Bloomquist, and M. A. Suchard, 2008 Smooth skyline through a rough skyline: Bayesian coalescent-based inference of population dynamics. *Mol. Biol. Evol.* 25: 1459–1471.
- Neher, R. A., and T. Leitner, 2010 Recombination rate and selection strength in HIV intra-patient evolution. *PLoS Comput. Biol.* 6: e10000660.
- Nielsen, R., 1997 A likelihood approach to populations samples of microsatellite alleles. *Genetics* 146: 711–716.
- Nowak, M., and R. M. May, 2000 *Virus Dynamics: Mathematical Principles of Immunology and Virology*. Oxford University Press, London/New York/Oxford.
- Paul, J. S., and Y. S. Song, 2010 A principled approach to deriving approximate conditional sampling distributions in population genetics models with recombination. *Genetics* 186: 321–338.

- Paul, J. S., M. Steinrücken, and Y. S. Song, 2011 An accurate sequentially Markov conditional sampling distribution for the coalescent with recombination. *Genetics* 187: 1115–1128.
- Pennings, P. S., S. Kryazhimskiy, and J. Wakeley, 2014 Loss and recovery of genetic diversity in adapting populations of HIV. *PLoS Genet.* 10: e1004000.
- Poon, A. F. Y., L. C. Swenson, E. M. Bunnik, D. Edo-Matas, H. Schuitemaker *et al.*, 2012 Reconstructing the dynamics of HIV evolution within hosts from serial deep sequence data. *PLoS Comput. Biol.* 8: e1002753.
- Pybus, O. G., and A. Rambaut, 2009 Evolutionary analysis of the dynamics of viral infectious disease. *Nat. Rev. Genet.* 10: 540–550.
- Pybus, O. G., A. Rambaut, and P. H. Harvey, 2000 An integrated framework for the inference of viral population history from reconstructed genealogies. *Genetics* 155: 1429–1437.
- Rasmussen, M. D., M. J. Hubisz, I. Gronau, and A. Siepel, 2014 Genome-wide inference of ancestral recombination graphs. *PLoS Genet.* 10: e1004342.
- Rodrigo, A. G., and J. Felsenstein, 1999 Coalescent approaches to HIV population genetics, pp. 233–272 in *The Evolution of HIV*, edited by K. A. Crandall. Johns Hopkins University Press, Baltimore.
- Ross, H. A., and A. G. Rodrigo, 2002 Immune-mediated positive selection drives human immunodeficiency virus type 1 molecular variation and predicts disease duration. *J. Virol.* 76: 11715–11720.
- Rouzine, I., and J. Coffin, 1999 Linkage disequilibrium test implies a large effective population number for HIV *in vivo*. *Proc. Natl. Acad. Sci. USA* 96: 10758–10763.
- Rouzine, I. M., J. M. Coffin, and L. S. Weinberger, 2014 Fifteen years later: hard and soft selection sweeps confirm a large population number for HIV *in vivo*. *PLoS Genet.* 10: e1004179.
- Shankarappa, R., J. B. Margolick, S. J. Gange, A. G. Rodrigo, D. Upchurch *et al.*, 1999 Consistent viral evolutionary changes associated with the progression of human immunodeficiency virus type 1 infection. *J. Virol.* 73: 10489–10502.
- Shapiro, B., A. Rambaut, and A. J. Drummond, 2006 Choosing appropriate substitution models for the phylogenetic analysis of protein-coding sequences. *Mol. Biol. Evol.* 23: 7–9.
- Sheehan, S., K. Harris, and Y. S. Song, 2013 Estimating variable effective population sizes from multiple genomes: a sequentially Markov conditional sampling distribution approach. *Genetics* 194: 647–662.
- Shriner, D., A. G. Rodrigo, D. C. Nickle, and J. I. Mullins, 2004 Pervasive genomic recombination of HIV-1 *in vivo*. *Genetics* 167: 1573–1583.
- Stephens, M., and P. Donnelly, 2000 Inference in molecular population genetics. *J. R. Stat. Soc. B* 62: 605–655.
- Wang, Y., and B. Rannala, 2008 Bayesian inference of fine-scale recombination rates using population genomic data. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 363: 3921–3930.
- Williamson, S., 2003 Adaptation in the *env* gene of HIV-1 and evolutionary theories of disease progression. *Mol. Biol. Evol.* 20: 1318–1325.
- Wilson, I. J., M. E. Weale, and D. J. Balding, 2003 Inferences from DNA data: population histories, evolutionary processes and forensic match probabilities. *J. R. Stat. Soc. Ser. A* 166: 155–201.

Communicating editor: J. D. Wall

Appendix

Appendix A: Backward Transition Rates in the ARG

The distribution on the most recent event back in time in a two-locus ARG, *conditioned* on observing the sample configuration $H_{-k} = \mathbf{n}$, is given in Table A1. This distribution is expressed in terms of the (unknown) CSD $\pi[(i,j)|\mathbf{n};\Phi]$ and transitions are marginalized over nonancestral genetic material; thus, for example, a recombination event changes an (i,j) haplotype to an $(i,*)$ and an $(*,j)$, with “*” denoting unspecified alleles (Such a marginalization approach was proposed by Jenkins and Griffiths 2011, but was based instead on an infinite-sites model.) The corresponding forward transition probabilities are also given in the rightmost column of Table A1. Here, $\mathbf{e}_i$ denotes a unit vector with i th entry 1 and the rest 0, with $\mathbf{e}_{ij}$ denoting a unit matrix whose only nonzero entry is a 1 at position (i,j) . The notation $\mathbf{n} - \mathbf{e}_{ij}^C$ is shorthand for $(\mathbf{a}, \mathbf{b}, \mathbf{c} - \mathbf{e}_{ij})$ and so on, and we write $c_i = \sum_j c_{ij}$ and $c_j = \sum_i c_{ij}$ for the marginal counts of fully specified haplotypes. To avoid computing each row in Table A1 at each step, a simplified proposal distribution first picks a haplotype according to Table A2, as described in the main text.

Table A1 Possible backward transitions from state $H_{-k} = \mathbf{n}$ to state H_{-k-1} in a two-locus ARG in which loci are permitted to remain unspecified

Event	H_{-k-1}	$\mathbb{P}(H_{-k-1} \mathbf{n};\Phi)$	$\mathbb{P}(H_{-k} \mathbf{n};\Phi)$
Coalescence of two (i,j) s	$\mathbf{n} - \mathbf{e}_{ij}^C$	$\frac{c_{ij}(c_{ij}-1)}{D} \cdot \frac{1}{\pi[(i,j) \mathbf{n}-\mathbf{e}_{ij}^C;\Phi]}$	$\frac{c(c_{ij}-1)}{D}$
Coalescence of $(i,*)$ with $(i,*)$ or (i,j)	$\mathbf{n} - \mathbf{e}_i^A$	$\frac{a_i(a_i+2c_i-1)}{D} \cdot \frac{1}{\pi[(i,*) \mathbf{n}-\mathbf{e}_i^A;\Phi]}$	$\frac{a(a_i+2c_i-1)}{D}$
Coalescence of $(*,j)$ with $(*,j)$ or (i,j)	$\mathbf{n} - \mathbf{e}_j^B$	$\frac{b_j(b_j+2c_j-1)}{D} \cdot \frac{1}{\pi[(*,j) \mathbf{n}-\mathbf{e}_j^B;\Phi]}$	$\frac{b(b_j+2c_j-1)}{D}$
Coalescence of $(i,*)$ and $(*,j)$	$\mathbf{n} - \mathbf{e}_i^A - \mathbf{e}_j^B + \mathbf{e}_{ij}^C$	$\frac{2a_i b_j}{D} \cdot \frac{\pi[(i,j) \mathbf{n}-\mathbf{e}_i^A - \mathbf{e}_j^B + \mathbf{e}_{ij}^C;\Phi]}{\pi[(i,*) \mathbf{n}-\mathbf{e}_i^A;\Phi] \pi[(*,j) \mathbf{n}-\mathbf{e}_j^B;\Phi]}$	$\frac{2ab(c_{ij}+1)}{(c+1)D}$
Mutation of (i,j) to (k,j)	$\mathbf{n} - \mathbf{e}_{ij}^C + \mathbf{e}_{kj}^C$	$\frac{\theta_A P_{ij}^A c_{ij}}{D} \cdot \frac{\pi[(k,j) \mathbf{n}-\mathbf{e}_{ij}^C + \mathbf{e}_{kj}^C;\Phi]}{\pi[(i,j) \mathbf{n}-\mathbf{e}_{ij}^C;\Phi]}$	$\frac{\theta_A P_{ij}^A (c_{ij}+1)}{D}$
Mutation of (i,j) to (i,l)	$\mathbf{n} - \mathbf{e}_{ij}^C + \mathbf{e}_{il}^C$	$\frac{\theta_B P_{ij}^B c_{ij}}{D} \cdot \frac{\pi[(i,l) \mathbf{n}-\mathbf{e}_{ij}^C + \mathbf{e}_{il}^C;\Phi]}{\pi[(i,j) \mathbf{n}-\mathbf{e}_{ij}^C;\Phi]}$	$\frac{\theta_B P_{ij}^B (c_{ij}+1)}{D}$
Mutation of $(i,*)$ to $(k,*)$	$\mathbf{n} - \mathbf{e}_i^A + \mathbf{e}_k^A$	$\frac{\theta_A P_{ij}^A a_i}{D} \cdot \frac{\pi[(k,*) \mathbf{n}-\mathbf{e}_i^A + \mathbf{e}_k^A;\Phi]}{\pi[(i,*) \mathbf{n}-\mathbf{e}_i^A;\Phi]}$	$\frac{\theta_A P_{ij}^A (a_i+1)}{D}$
Mutation of $(*,j)$ to $(*,l)$	$\mathbf{n} - \mathbf{e}_j^B + \mathbf{e}_l^B$	$\frac{\theta_B P_{ij}^B b_j}{D} \cdot \frac{\pi[(*,l) \mathbf{n}-\mathbf{e}_j^B + \mathbf{e}_l^B;\Phi]}{\pi[(*,j) \mathbf{n}-\mathbf{e}_j^B;\Phi]}$	$\frac{\theta_B P_{ij}^B (b_j+1)}{D}$
Recombination of (i,j)	$\mathbf{n} - \mathbf{e}_{ij}^C + \mathbf{e}_i^A + \mathbf{e}_j^B$	$\frac{\rho c_{ij}}{D} \cdot \frac{\pi[(i,*) \mathbf{n}-\mathbf{e}_{ij}^C + \mathbf{e}_i^A;\Phi] \pi[(*,j) \mathbf{n}-\mathbf{e}_{ij}^C + \mathbf{e}_j^B;\Phi]}{\pi[(i,j) \mathbf{n}-\mathbf{e}_{ij}^C;\Phi]}$	$\frac{\rho c(a_i+1)(b_j+1)}{(a+1)(b+1)D}$

Backward rates are expressed in terms of the CSD $\pi[\cdot|\cdot;\Phi]$. $\mathbf{e}_i$ denotes a unit vector with i th entry 1 and the rest 0, with $\mathbf{e}_{ij}$ denoting a unit matrix whose only nonzero entry is a 1 at position (i,j) . The notation $\mathbf{n} - \mathbf{e}_{ij}^C$ is shorthand for $(\mathbf{a}, \mathbf{b}, \mathbf{c} - \mathbf{e}_{ij})$ and so on, and we write $c_i = \sum_j c_{ij}$ and $c_j = \sum_i c_{ij}$ for the marginal counts of fully specified haplotypes. The normalizing constant D denotes (twice) the total prior rate of events, $D = n(n-1) + (a+c)\theta_A + (b+c)\theta_B + \rho c$.

Table A2 Proposal probability of choosing a particular (i,j) , $(i,*)$, or $(*,j)$, to be involved in the next event back in time, for each i,j

Haplotype	Probabilities up to a normalizing constant
(i,j)	$c_{ij}(c_{ij}-1 + \theta + \rho + a_i + b_j)$
$(i,*)$	$a_i(a_i-1 + b + c_i + \theta_A)$
$(*,j)$	$b_j(b_j-1 + a + c_j + \theta_B)$

These probabilities are given up to a normalizing constant found by summing all these events over i and j .

To derive the entries in Table A1 we use sampling exchangeability and Bayes’ theorem (see Stephens and Donnelly 2000; De Iorio and Griffiths 2004a). For example, to obtain the first row of Table A1, with $H_{-k-1} = \mathbf{n} - \mathbf{e}_{ij}^C$, note that sampling exchangeability implies

$$\mathbb{P}(\mathbf{n} - \mathbf{e}_{ij}^C) \pi[(i,j)|\mathbf{n} - \mathbf{e}_{ij}^C] = \mathbb{P}(\mathbf{n}) \frac{c_{ij}}{c}, \quad (\text{A1})$$

since both sides are equal to the probability that we sample the configuration $\mathbf{n}$ and then subsample a type (i,j) , on the left by selecting the fully specified haplotype we sampled most recently and on the right by subsampling from $\mathbf{c}$. [Our convention for “unordered” is that $\mathbb{P}(\mathbf{n}) = \sigma(\mathbf{a}, \mathbf{b}, \mathbf{c}) \cdot \mathbb{P}(\mathcal{A}_n)$, where $\sigma(\mathbf{a}, \mathbf{b}, \mathbf{c}) = a!b!c! / (\prod_i a_i! \prod_j b_j! \prod_{i,j} c_{ij}!)$ and $\mathcal{A}_n$ denotes any *particular* ordering of the sample configuration $\mathbf{n}$. In this convention, whether or not a haplotype is fully specified is not considered to be “random”; that is, the missingness or otherwise of either allele need not be assigned a probability. Instead, we can think of this as having conditioned on the missingness status of each locus.]

Now, by Bayes' theorem,

$$\mathbb{P}(H_{-k-1}|H_{-k}) = \mathbb{P}(H_{-k}|H_{-k-1}) \frac{\mathbb{P}(H_{-k-1})}{\mathbb{P}(H_{-k})}, \quad (\text{A2})$$

with

$$\mathbb{P}(H_{-k}|H_{-k-1}) = \frac{c(c_{ij} - 1)}{n(n-1) + (a+c)\theta_A + (b+c)\theta_B + \rho c} \quad (\text{A3})$$

known from the coalescent prior. Probabilities such as (A3) can be obtained, for example, from the recursive system of equations studied by Ethier and Griffiths (1990) and Jenkins and Song (2009) [for example, multiply equation 5 of Jenkins and Song 2009 by $\sigma(\mathbf{a}, \mathbf{b}, \mathbf{c})$, rearrange to express the system in terms of $\mathbb{P}(\mathbf{n})$, and then read off the coefficients of the resulting system].

The entry for $\mathbb{P}(H_{-k-1}|H_{-k})$ in the first row of Table A1 follows by combining (A1), (A2), and (A3). The remaining entries follow similarly. At the bottom row, $\pi[\{(i, *), (*, j)\}|\mathbf{n} - \mathbf{e}_{ij}^C; \Phi]$ denotes

$$\pi[\{(i, *), (*, j)\}|\mathbf{n} - \mathbf{e}_{ij}^C; \Phi] = \pi[(i, *)|\mathbf{n} - \mathbf{e}_{ij}^C; \Phi] \pi[(*, j)|\mathbf{n} - \mathbf{e}_{ij}^C + \mathbf{e}_i^A; \Phi] = \pi[(*, j)|\mathbf{n} - \mathbf{e}_{ij}^C; \Phi] \pi[(i, *)|\mathbf{n} - \mathbf{e}_{ij}^C + \mathbf{e}_j^B; \Phi].$$

Appendix B: Emission Probabilities for Gaussian Quadrature

After applying Gaussian quadrature it is necessary to compute discretized emission probabilities of the form

$$\xi_\theta(j|([x, y], i)) = \frac{1}{w^{(x, y)}} \int_x^y e^{-\tau} \xi_\theta(j|s_\ell) d\tau, \quad (\text{B1})$$

where $w^{(x, y)} = \int_x^y e^{-\tau} d\tau$ (Paul *et al.* 2011, equation 17). In light of (8), we are also able to reduce this to a closed-form formula as follows. Substituting (8) into (B1), we find

$$\begin{aligned} \xi_\theta(j|([x, y], i)) &= \frac{1}{w^{(x, y)}} \int_x^y e^{-\tau} \left[e^{-\bar{\theta}\tau/2} \sinh\left(\frac{\bar{\theta}}{2}\tau\right) \right]^{S_{ij}} \left[e^{-\bar{\theta}\tau/2} \cosh\left(\frac{\bar{\theta}}{2}\tau\right) \right]^{l_\ell - S_{ij}} d\tau \\ &= \frac{1}{w^{(x, y)}} \int_x^y \frac{e^{-\tau}}{2^{l_\ell}} \left[1 - e^{-\bar{\theta}\tau} \right]^{S_{ij}} \left[1 + e^{-\bar{\theta}\tau} \right]^{l_\ell - S_{ij}} d\tau \\ &= \frac{1}{w^{(x, y)}} \int_x^y \frac{e^{-\tau}}{2^{l_\ell}} \sum_{k=0}^{S_{ij}} \binom{S_{ij}}{k} e^{-\bar{\theta}\tau k} (-1)^k \sum_{m=0}^{l_\ell - S_{ij}} \binom{l_\ell - S_{ij}}{m} e^{-\bar{\theta}\tau m} d\tau \\ &= \frac{2^{-l_\ell}}{w^{(x, y)}} \sum_{k=0}^{S_{ij}} \binom{S_{ij}}{k} (-1)^k \sum_{m=0}^{l_\ell - S_{ij}} \binom{l_\ell - S_{ij}}{m} \int_x^y e^{-[(\bar{\theta}(k+m)+1)\tau]} d\tau \\ &= \frac{2^{-l_\ell}}{w^{(x, y)}} \sum_{k=0}^{S_{ij}} \binom{S_{ij}}{k} (-1)^k \sum_{m=0}^{l_\ell - S_{ij}} \binom{l_\ell - S_{ij}}{m} \frac{e^{-[\bar{\theta}(k+m)+1]x} - e^{-[\bar{\theta}(k+m)+1]y}}{\bar{\theta}(k+m) + 1}. \end{aligned}$$

GENETICS

Supporting Information

www.genetics.org/lookup/suppl/doi:10.1534/genetics.115.177931/-/DC1

Coalescent Inference Using Serially Sampled, High-Throughput Sequencing Data from Intrahost HIV Infection

**Kevin Dialdestoro, Jonas Andreas Sibbesen, Lasse Maretty, Jayna Raghwani, Astrid Gall,
Paul Kellam, Oliver G. Pybus, Jotun Hein, and Paul A. Jenkins**

SUPPORTING INFORMATION:

COALESCENT INFERENCE USING SERIALY SAMPLED, HIGH-THROUGHPUT SEQUENCING DATA FROM INTRA-HOST HIV INFECTION

Kevin Dialdestoro^a, Jonas Andreas Sibbesen^{b,*}, Lasse Maretty^{b,*}, Jayna Raghvani^c,
Astrid Gall^d, Paul Kellam^{d,e}, Oliver G. Pybus^c, Jotun Hein^a, Paul A. Jenkins^{f,g}

^a Department of Statistics, University of Oxford, Oxford, United Kingdom

^b The Bioinformatics Centre, Department of Biology, University of Copenhagen, Copenhagen,
Denmark

^c Department of Zoology, University of Oxford, Oxford, United Kingdom

^d Wellcome Trust Sanger Institute, Wellcome Trust Genome Campus, Hinxton, Cambridge,
United Kingdom

^e UCL/MRC Centre for Medical Molecular Virology, Division of Infection and Immunity,
University College London, United Kingdom

^f Department of Statistics, University of Warwick, Coventry, United Kingdom

^g Department of Computer Science, University of Warwick, Coventry, United Kingdom

* These authors contributed equally to the work.

(A) 120 Seqs, 1000 iterations

(B) 120 Seqs, 5000 iterations

FIGURE S1: As in Figure 5 but with Monte Carlo iterations increased by a factor of 10.

FIGURE S2: Variation of maximum likelihood estimates with respect to distance d between pairs of loci. Each point corresponds to the maximum likelihood estimate (y -axis) for a pair of loci separated by a particular distance (x -axis). The size of each point corresponds to the number of estimates with these values.

(A) 120 Seqs, 100 iterations

(B) 120 Seqs, 500 iterations

FIGURE S3: Pairwise composite likelihood surfaces for dataset *Const-120*, using 100 and 500 Monte Carlo iterations. Cells correspond to the searched parameters, colored by log-likelihoods, with the top 10 estimates numbered. The true mutation, recombination, and population parameters, are: $\bar{\mu} = 2.5 \times 10^{-5}$, $\bar{r} = 10^{-6}$, $N_e = 10^3$.

(A) 600 Seqs, 100 iterations

(B) 600 Seqs, 500 iterations

FIGURE S4: As in Figure S3 but for dataset *Const-600*.

TABLE S1: Distribution of read lengths for the HIV-1 dataset.

Read length	Number of reads
0–100	47
100–200	880
200–300	6255
300–400	12840
400–500	21806
500–600	30676
600–700	14480
700–800	1573
800–900	20

TABLE S2: Distribution of reads across the genome with 95% coverage (i.e. $< 5\%$ gaps in sequences) for the HIV-1 dataset.

Midpoint position of window	Number of reads with 95% coverage	Midpoint position of window	Number of reads with 95% coverage
50	841	4650	1280
150	13881	4750	1355
250	14143	4850	1244
350	11495	4950	1044
450	10044	5050	1320
550	5234	5150	1558
650	4599	5250	1626
750	4887	5350	1780
850	4581	5450	1989
950	3848	5550	2385
1050	4897	5650	6732
1150	11110	5750	5990
1250	9043	5850	6587
1350	5890	5950	6631
1450	7639	6050	5574
1550	11577	6150	3311
1650	12547	6250	4602
1750	12558	6350	441
1850	11779	6450	3682
1950	844	6550	3531
2050	1497	6650	3861
2150	1545	6750	4187
2250	1763	6850	3756
2350	1815	6950	3653
2450	1884	7050	3484
2550	2005	7150	6
2650	1904	7250	2955
2750	1759	7350	3026
2850	1599	7450	2088
2950	1574	7550	2164
3050	1482	7650	1963
3150	1455	7750	2518
3250	1364	7850	2366
3350	1332	7950	2069
3450	1192	8050	1955
3550	1218	8150	1797
3650	1784	8250	1799
3750	2756	8350	1825
3850	3293	8450	1728
3950	3192	8550	2334
4050	2890	8650	2462
4150	3173	8750	2311
4250	3469	8850	2137
4350	2934	8950	2088
4450	3771	9050	2067
4550	1057		

TABLE S3: Regions of the HIV-1 genome chosen for study and their sample sizes. Samples are collected at seven time points over the period of two years: 0, 267, 323, 520, 576, 667, 695 (days towards the past)

Region	Nucleotide position relative to consensus genome start	Nucleotide position relative to HXB2 genome start	Number of sequences at collection times	Total sequences
gag	311–911	790–1413	38, 34, 33, 33, 37, 37, 44	256
	960–1560	1432–2033	47, 16, 15, 27, 45, 36, 32	218
pol	2005–2605	2482–3072	27, 16, 17, 24, 23, 37, 58	202
	2836–3436	3312–3910	20, 15, 23, 12, 29, 19, 56	175
	3796–4396	4272–4870	55, 29, 49, 39, 53, 46, 159	430
env	5812–6412	6288–6840	21, 9, 6, 27, 27, 17, 58	165
	6415–7015	6843–7420	52, 37, 51, 60, 66, 42, 81	389
	7357–7957	7779–8379	34, 12, 24, 22, 23, 16, 43	174
nef	8376–9011	8797–9417	29, 12, 20, 11, 18, 6, 57	153

TABLE S4: Classification of sites for the simulated datasets. For each gene region, proportions of site types out of the total number of sites, are shown in brackets for proportions that are above 0.05. Only diallelic sites are considered. The last two column of the table shows the number of sites that are analysed, and their breakdown across the partition of ten 50 nt loci.

Dataset	Sites by types				Breakdown of diallelic sites across loci
	Total	Diallelic	3-allelic	4-allelic	
<i>Const-120</i>	500	202 (0.4)	72 (0.14)	3	18, 17, 22, 26, 15, 18, 19, 20, 24, 23
<i>Const-600</i>	500	226 (0.45)	117 (0.23)	29 (0.06)	21, 24, 28, 21, 23, 16, 21, 21, 26, 25
<i>Dynamic-120</i>	500	122 (0.24)	11	2	12, 10, 14, 14, 13, 13, 14, 10, 13, 9
<i>Dynamic-600</i>	500	194 (0.39)	55 (0.11)	10	24, 18, 21, 20, 28, 20, 18, 16, 18, 11

TABLE S5: Classification of sites for the HIV-1 dataset. For each gene region, proportions of site types out of the total number of sites, are shown in brackets for proportions that are above 0.05. Only diallelic sites and those in the central positions (51, 550) of each gene region are considered. The last two columns of the table show the number of sites that are analysed, and their breakdown across the partition of ten 50 nt loci.

Region	Sites by types							Diallelic sites at central position	
	Total	Ambiguous	Indels	Missing	Diallelic	3-allelic	4-allelic	Total	Breakdown across loci
gag 311–940	630	5	36 (0.06)	1	159 (0.25)	16	1	131	17, 14, 10, 9, 14, 21, 16, 8, 12, 10
gag 960–1560	602	4	76 (0.13)	2	112 (0.19)	12	1	89	6, 13, 9, 7, 7, 8, 6, 9, 14, 10
pol 2005–2605	599	2	53 (0.09)	0	102 (0.17)	10	0	86	8, 8, 6, 9, 5, 6, 7, 9, 16, 12
pol 2836–3436	599	2	80 (0.13)	0	116 (0.19)	15	0	85	19, 9, 8, 9, 4, 5, 6, 7, 6, 12
pol 3796–4396	599	3	81 (0.14)	0	130 (0.22)	30 (0.05)	4	104	9, 5, 11, 9, 11, 13, 9, 14, 16, 7
env 5812–6412	601	1	79 (0.13)	0	107 (0.18)	12	0	87	7, 9, 10, 3, 8, 7, 15, 10, 9, 9
env 6415–7015	601	9	148 (0.25)	0	164 (0.27)	27	7	141	15, 9, 20, 20, 9, 14, 14, 11, 16, 13
env 7357–7957	601	16	74 (0.12)	0	97 (0.16)	15	0	79	10, 10, 4, 7, 7, 11, 8, 9, 7, 6
nef 8376–9011	636	16	32 (0.05)	0	106 (0.17)	6	2	82	9, 12, 7, 5, 8, 8, 4, 6, 13, 10

TABLE S6: Comparison of BEAST’s results against Coalescenator. Mutation $\bar{\mu}$ and recombination $\bar{\rho}$ rates are converted to the number of substitutions and recombinations per site per year, respectively. N_e represents effective population size. The time to the most recent common ancestor T_{MRCA} , or the time until 5 ancestors is reached T_{MRCA-5} , are given in years. Shown in brackets are BEAST’s 95% highest posterior densities.

Region	BEAST			Coalescenator: median of pairwise MLEs			
	$\bar{\mu}$	N_e	T_{MRCA}	$\bar{\mu}$	$\bar{\rho}$	N_e	T_{MRCA-5}
gag 311–940	1.33e-02 [9.39e-03, 1.75e-02]	3790 [2820, 4770]	3.94 [3.27, 4.69]	1.27e-02	3.04e-03	1250	4.59
gag 960–1560	4.05e-03 [3.04e-03, 5.12e-03]	2550 [1710, 3430]	4.38 [3.3, 5.59]	1.52e-02	5.07e-03	1000	4.31
pol 2005–2605	2.3e-03 [1.45e-03, 3.25e-03]	3660 [1900, 5750]	7.44 [4.59, 10.76]	1.01e-02	1.27e-02	1125	4.52
pol 2836–3436	6.73e-03 [5.14e-03, 8.59e-03]	2450 [1730, 3310]	5.13 [3.74, 6.54]	7.61e-03	5.07e-03	1500	6.34
pol 3796–4396	1.16e-02 [9.39e-03, 1.38e-02]	1720 [1370, 2070]	3.31 [2.77, 3.9]	1.01e-02	5.07e-03	750	3.81
env 5812–6412	1.52e-02 [1.14e-02, 1.9e-02]	455 [304, 602]	2.45 [2.21, 2.72]	1.52e-02	2.03e-03	1000	4.2
env 6415–7015	2.76e-02 [1.14e-02, 3.2e-02]	2090 [1710, 2500]	3.15 [2.69, 3.68]	1.52e-02	5.07e-03	875	3.43
env 7357–7957	1.06e-02 [8.18e-03, 1.33e-02]	1220 [883, 1580]	3.44 [2.82, 4.13]	1.78e-02	5.07e-03	875	3.81
nef 8376–9011	1.99e-02 [1.51e-02, 2.53e-02]	1430 [1040, 1860]	3.19 [2.67, 3.77]	1.52e-02	3.55e-03	750	3.94