



Ancestral inference from haplotypes and mutations

Robert C. Griffiths^{a,*}, Simon Tavaré^b

^a Department of Statistics, University of Oxford, 24–29 St Giles, Oxford OX1 3LB, UK

^b DAMTP, University of Cambridge, Centre for Mathematical Sciences, Wilberforce Road, Cambridge CB3 0WA, UK

ARTICLE INFO

Article history:

Available online 25 April 2018

We dedicate this paper to the memory of Paul Joyce, friend and collaborator

Keywords:

Ancestral inference
Coalescent inference
Ewens sampling formula
Ancestral lineages
Standing variation

ABSTRACT

We consider inference about the history of a sample of DNA sequences, conditional upon the haplotype counts and the number of segregating sites observed at the present time. After deriving some theoretical results in the coalescent setting, we implement rejection sampling and importance sampling schemes to perform the inference. The importance sampling scheme addresses an extension of the Ewens Sampling Formula for a configuration of haplotypes and the number of segregating sites in the sample. The implementations include both constant and variable population size models. The methods are illustrated by two human Y chromosome datasets.

© 2018 Published by Elsevier Inc.

1. Introduction

In this paper we study aspects of the ancestral history of a random sample of n DNA sequences, conditional on features of the haplotype configuration obtained at the present time, labelled 0 . Initially, we assume an infinitely-many-sites mutation model with constant population size where the sites are in a completely linked region of DNA where there is no recombination. We begin with some theory that describes the effects of mutations between time 0 and time t in the past. We describe the distribution of the quantities $\tilde{S}_n(t)$, $A_n(t)$, $\tilde{K}_n(t)$ and $A_n^\theta(t)$, where $\tilde{S}_n(t)$ is the number of mutations that have arisen in $(0, t)$, $A_n(t)$ is the number of ancestors at time t , $\tilde{K}_n(t)$ is the number of distinct haplotypes present in those ancestors that are formed by the mutations in $(0, t)$, and $A_n^\theta(t)$ is the number of ancestors at time t whose descendants have no further mutation in $(0, t)$.

Results such as these describe the effects of *new* variation, that arising from the present time 0 to time t in the past. We then describe the effects of standing variation by providing a simulation approach for studying the joint distribution of S_n and K_n , the number of segregating sites and haplotypes in the sample at time 0 resp., and $S_n(t)$, $K_n(t)$, the number of segregating sites and haplotypes, resp., in the $A_n(t)$ ancestors formed by standing variation arising after time t in the past. This approach allows for variable population size, as well as essentially arbitrary binary branching models such as the Yule process.

If the full haplotype configuration and the number of segregating sites are available, then there are many other questions that

may be asked about the ancestral history of the current sample. A sequential importance sampling algorithm for studying the ancestral history of a sample of genes conditional on the haplotype configuration and the number of segregating sites is developed. If the complete pattern of mutations on haplotypes were known, then a perfect phylogeny, a genetree, could be constructed and ancestral inference, such as the ages of mutations and time to the most recent common ancestor made conditional on the genetree topology. An example appears in the analysis in Griffiths and Tavaré (1994b).

Here we consider the case when just the haplotype frequencies and the number of mutations are known. The theory and computational algorithms are then much simpler. The implementation computes, *inter alia*, the probability of a sample configuration of haplotypes and segregating sites in a stationary population (this is an extension of the Ewens Sampling Formula), the average coalescence times, mutation times, allele loss times and allele ages back in time conditional on the current haplotype configuration and number of segregating sites, and the conditional average allele configuration and distribution of ancestors at time t in the past. In the infinitely-many-sites model, the haplotype configuration follows the infinitely-many-alleles model so the haplotype configuration in a sample is the same as the allele configuration and has a probability distribution of the Ewens Sampling Formula.

2. Ancestral distributions in the coalescent: theory

We begin by setting some notation. We let $A_n(t)$ denote the number of ancestors of a sample of size n time t ago, and $A_n^\theta(t)$ the number of those ancestors whose descendants in $(0, t)$ have no

* Corresponding author.

E-mail address: griffiths.bob@gmail.com (R.C. Griffiths).

further mutation. We let $\tilde{S}_n(t)$ be the number of segregating sites that arise as mutations in $(0, t)$, $\tilde{K}_n(t)$ be the number of haplotypes in these ancestors that arise from mutations in $(0, t)$, and $S_n(t)$ the number of mutations in the sample arising after time t .

Coalescence times between events are denoted by T_2, T_3, \dots . In the constant size coalescent model these are independent exponential random variables with rates $\binom{2}{2}, \binom{3}{2}, \dots$ (Kingman, 1982). When considering non-mutant lineages back in time, it is appropriate to have lineages lost by either mutation or coalescence. Times between these events are denoted by $T_2^\theta, T_3^\theta, \dots$. In the constant population size case, these are independent exponential random variables with rates $\binom{2}{2} + 2\theta, \binom{3}{2} + 3\theta, \dots, \theta$ being the scaled mutation rate. We denote by $f_{nk}^\theta(t)$ the density of $T_n^\theta + \dots + T_k^\theta$.

The reader is referred to Ewens (1972) for the distribution of haplotypes in a sample, the Ewens Sampling Formula; Griffiths (1980), eq. (9), for an original derivation of the distribution of $A_\infty^\theta(t)$ from looking back in time in a diffusion process; Tavaré (1984) for an original derivation of $A_n^\theta(t)$, connections with the Kingman coalescent and review of ancestral lineage distributions; Tavaré (1984, 2004) for an introduction to ancestral lineage distributions and ancestral inference in population genetics; and Griffiths (2006) for a review and new representations for ancestral distributions.

The distribution of $A_n^\theta(t)$ is given by

$$\begin{aligned} \mathbb{P}(A_n^\theta(t) = k) &:= q_{nk}^\theta(t) \\ &= \sum_{j=k}^n \rho_j^\theta(t) (-1)^{j-k} \frac{(2j + \theta - 1)(k + \theta)_{j-1}}{k!(j-k)!} \\ &\quad \cdot \frac{n_{[j]}}{(n + \theta)_{[j]}}, \end{aligned} \tag{1}$$

where $\rho_j^\theta(t) = e^{-\frac{1}{2}j(j+\theta-1)t}$; see Tavaré (1984), eq. (5.2). $\mathbb{P}(A_n(t) = k)$ is given by setting $\theta = 0$. The formula also holds for $n = \infty$, where the interpretation is that of the whole infinite population coalescent.

The falling factorial moments of $A_n^\theta(t)$ are

$$\mathbb{E}[A_n^\theta(t)_{[r]}] = \sum_{k=r}^n \rho_k^\theta(t) (2k + \theta - 1) \binom{k-1}{r-1} (\theta + k)_{(r-1)} \frac{n_{[k]}}{(n + \theta)_{(k)}}; \tag{2}$$

see Tavaré (1984), p. 13. A simple rate argument establishes the correspondence between the distribution of coalescent times and the number of ancestors, namely

$$f_{nl}^\theta(t) = \frac{2}{l(l + \theta - 1)} \mathbb{P}(A_n^\theta(t) = l). \tag{3}$$

In a stationary population, the probability generating function (pgf) of the number of segregating sites in a sample of n genes is

$$\begin{aligned} H_n(z) &= \prod_{j=1}^{n-1} \left(1 + \frac{\theta}{j}(1-z)\right)^{-1} \\ &= (n-1) \cdot \frac{\Gamma(n-1)\Gamma(\theta(1-z)+1)}{\Gamma(n+\theta(1-z))} \\ &= \int_0^1 x^{\theta(1-z)} (n-1)(1-x)^{n-2} dx. \end{aligned} \tag{4}$$

Therefore in a stationary population,

$$\mathbb{P}(S_n = k) = (n-1) \int_0^1 (1-x)^{n-2} e^{-(-\theta \log x)^k} \frac{(-\theta \log x)^k}{k!} dx, \tag{5}$$

which is a Poisson mixture with mean $-\theta \log X$, where X has density $(n-1)(1-x)^{n-2}$, $0 < x < 1$. A calculation in Tavaré (2004)

eq. (5.3.6) using the first line of (4) shows that

$$\mathbb{P}(S_n = k) = \frac{n-1}{\theta} \sum_{l=1}^{n-1} (-1)^{l-1} \binom{n-2}{l-1} \left(\frac{\theta}{l+\theta}\right)^{k+\theta}.$$

There is the simple recursion that

$$n(n + \theta - 1)\mathbb{P}(S_n = s) = \theta\mathbb{P}(S_n = s - 1) + n(n - 1)\mathbb{P}(S_{n-1} = s), \quad n \geq 2,$$

with $\mathbb{P}(S_1 = s) = \delta_{s0}$.

The distribution of the number of alleles in a sample of n in a stationary population is

$$\mathbb{P}(K_n = k) = \theta^k |S_k^n| / \theta_{(n)}, \quad 1 \leq k \leq n,$$

where $\{S_k^n\}$ are Stirling numbers of the first kind.

2.1. The joint distribution of $\tilde{S}_n(t)$ and $A_n(t)$

Measuring time back from the present, $(\tilde{S}_n(t), A_n(t))$ is a Markov process beginning at $(0, n)$ at $t = 0$ such that

$$(s, l) \rightarrow \begin{cases} (s+1, l) & \text{at rate } \theta l/2 \\ (s, l-1) & \text{at rate } l(l-1)/2. \end{cases} \tag{6}$$

The pgf of the joint distribution of $A_n(t), \tilde{S}_n(t)$ distribution is derived in Griffiths (1981), eqs. (8), (9). The approach taken in the present paper is new.

The joint pgf of the number of mutations arising while there are n, \dots, k ancestors and the density of T_n, T_{n-1}, \dots, T_k in an n -coalescent tree is

$$\begin{aligned} &\prod_{j=k}^n e^{\theta(z-1)j t_j/2} \binom{j}{2} e^{-\binom{j}{2} t_j/2} \\ &= \frac{\prod_{j=k}^n \binom{j}{2}}{\prod_{j=k}^n \frac{1}{2}j(j + \theta(1-z))} \cdot \prod_{j=k}^n \frac{1}{2}j(j + \theta(1-z)) e^{-\frac{1}{2}j(j-1+\theta(1-z))t_j}. \end{aligned} \tag{7}$$

Integrating over $t_n + t_{n-1} + \dots + t_k \leq t$, the pgf of the number of mutations on edges of the coalescent tree during times T_n, \dots, T_k and the distribution function of $T_n + \dots + T_k$ is

$$\begin{aligned} &\frac{\prod_{j=k}^n \binom{j}{2}}{\prod_{j=k}^n \frac{1}{2}j(j + \theta(1-z))} \cdot \mathbb{P}(T_k^{\theta(1-z)} + \dots + T_n^{\theta(1-z)} \leq t) \\ &= \frac{\prod_{j=k}^n \binom{j}{2}}{\prod_{j=k}^n \frac{1}{2}j(j + \theta(1-z))} \mathbb{P}(A_n^{\theta(1-z)}(t) \leq k) \\ &= \prod_{j=k}^{n-1} \left(1 + \frac{\theta(1-z)}{j}\right)^{-1} \mathbb{P}(A_n^{\theta(1-z)}(t) \leq k). \end{aligned} \tag{8}$$

The joint probability that $A_n(t) = l$ and pgf of the number of mutations $S_n(t)$ arising in $(0, t)$, which will be denoted by $G_l(z; t)$, is the probability that $T_n + \dots + T_{l+1} = \tau < t$ (a coalescence necessarily occurs at τ), there is no coalescence in (τ, t) and the pgf of the numbers of mutations on T_n, \dots, T_{l+1} and the l lines from τ to t . Therefore, with notation $\theta_z = \theta(1-z)$,

$$\begin{aligned} G_l(z; t) &= \frac{\prod_{j=l+1}^n \binom{j}{2}}{\prod_{j=l+1}^n \frac{1}{2}j(j-1 + \theta(1-z))} \\ &\quad \cdot \int_0^t e^{-\binom{j}{2}(t-\tau) + \theta(z-1)l(t-\tau)} f_{n,l+1}^{\theta(1-z)}(\tau) d\tau \\ &= \frac{\prod_{j=l+1}^n \binom{j}{2}}{\prod_{j=l}^n \frac{1}{2}j(j-1 + \theta(1-z))} \cdot f_{n,l}^{\theta(1-z)}(t) \\ &= \frac{\prod_{j=l+1}^n \binom{j}{2}}{\prod_{j=l+1}^n \frac{1}{2}j(j-1 + \theta_z)} \cdot \mathbb{P}(A_n^{\theta_z}(t) = l) \end{aligned} \tag{9}$$

$$\begin{aligned}
 &= \prod_{j=l}^{n-1} \left(1 + \frac{\theta_z}{j}\right)^{-1} \mathbb{P}(A_n^{\theta_z}(t) = l) \\
 &= \prod_{j=l}^{n-1} \left(1 + \frac{\theta_z}{j}\right)^{-1} \\
 &\quad \times \sum_{j=l}^n \rho_j^{\theta_z}(t) (-1)^{j-l} \frac{(2j + \theta_z - 1)(l + \theta_z)_{(j-1)}}{l!(j-l)!} \\
 &\quad \cdot \frac{n_{[j]}}{(n + \theta_z)_{[j]}}. \tag{10}
 \end{aligned}$$

The identity (3) is used in the calculation. As a check of (10), if $z = 1$ the probability that there are l ancestors at t is $\mathbb{P}(A_n^0(t) = l)$, as it should be.

The marginal pgf of the number of mutations arising while there are at least l ancestors is given by

$$\prod_{j=l-1}^{n-1} \left(1 + \frac{\theta_z}{j}\right)^{-1}. \tag{11}$$

To see this, calculate the pgf of the total number of mutations accumulating while there are at least l ancestors as

$$\begin{aligned}
 \frac{l(l-1)}{2} \int_0^\infty G_l(z; t) dt &= \frac{l-1}{(l-1+\theta_z)} \cdot \prod_{j=l}^{n-1} \left(1 + \frac{\theta_z}{j}\right)^{-1} \\
 &= \prod_{j=l-1}^{n-1} \left(1 + \frac{\theta_z}{j}\right)^{-1}. \tag{12}
 \end{aligned}$$

In a calculation similar to (5),

$$\begin{aligned}
 &\prod_{j=l-1}^{n-1} \left(1 + \frac{\theta_z}{j}\right)^{-1} \\
 &= \frac{\Gamma(n)}{\Gamma(l-1)\Gamma(n-l+1)} \cdot \frac{\Gamma(l-1+\theta_z)\Gamma(n-l+1)}{\Gamma(n+\theta_z)}
 \end{aligned}$$

is the pgf of a Poisson mixture with a rate $-\theta \log X$, where X has a Beta($l-1, n-l+1$) distribution. Hence the probability of k mutations, from the pgf (11), is

$$B(l-1, n-l+1)^{-1} \int_0^1 x^{l-2} (1-x)^{n-l} e^{-(\theta \log x)} \frac{(-\theta \log x)^k}{k!} dx, \tag{13}$$

for $k = 0, 1, \dots$

The pgf of the number of mutations arising in $(0, t)$ (not counting mutations when there is one ancestor) is $\sum_{l=2}^n G_l(z; t)$. The joint pgf for $S_n(t), \tilde{S}_n(t)$ and the probability of l ancestors time t ago is $Q_l(r)G_l(z; t)$, where $Q_l(r)$ is the pgf of the number of segregating sites in a sample of l from the population at the initial time. In a stationary population $Q_l(r) = H_l(r)$, defined in (4).

The probability $\mathbb{P}(\tilde{S}_n(t) = k)$ does not have a simple form, but if one considers standing variation in a stationary population as well as mutations in $(0, t)$ then the distribution of the number of mutations is as in a stationary population at time t and the conditional distribution of the number of ancestors given the number of mutations is easier. The pgf/probability of the number of mutations and number of ancestors at time t in the past is then

$$G_l^*(z, t) = \prod_{j=2}^{n-1} \left(1 + \frac{\theta_z}{j}\right)^{-1} \mathbb{P}(A_n^{\theta_z}(t) = l). \tag{14}$$

Summing (14) over $l = 2, \dots, n$, we see that the marginal pgf of the number of mutations is

$$\prod_{j=2}^{n-1} \left(1 + \frac{\theta_z}{j}\right)^{-1}, \tag{15}$$

and the distribution of the number of ancestors at time t back, found by setting $z = 1$ in (14), is

$$\mathbb{P}(A_n^0(t) = l). \tag{16}$$

Eqs. (15) and (16) recover earlier results.

The joint pgf and expected number of ancestors time t ago is

$$\begin{aligned}
 &\prod_{j=2}^{n-1} \left(1 + \frac{\theta_z}{j}\right)^{-1} \mathbb{E}[A_n^{\theta_z}(t)] \\
 &= \prod_{j=2}^{n-1} \left(1 + \frac{\theta_z}{j}\right)^{-1} \sum_{k=1}^n \rho_k^{\theta_z}(t) (2k + \theta_z - 1) \frac{n_{[k]}}{(n + \theta_z)_{[k]}}. \tag{17}
 \end{aligned}$$

Inversion of (17) is straightforward, if a little messy. We calculate $\mathbb{E}[A_n(t) | S_n = r]$. Let $a(r, k)$ be the coefficient of z^r in

$$M(k, z) = \prod_{j=2}^{n-1} \left(1 + \frac{\theta_z}{j}\right)^{-1} (n + \theta_z)_{[k]}^{-1}.$$

Then by equating coefficients of z^r in

$$(n + \theta + k - 1 - \theta z)M(k, z) = M(k-1, z),$$

for $k = 1, 2, \dots, n, r = 0, 1, \dots$, we get

$$(n + \theta + k - 1)a(r, k) = \theta a(r-1, k) + a(r, k-1).$$

Note that $a(r, 0) = \mathbb{P}(S_n = r)$. Let

$$b(r, k) = (2k + \theta - 1)a(r, k) - \theta a(r, k-1)$$

and finally

$$c(r, k) = \sum_{m=0}^r e^{-k\theta t/2} \frac{(k\theta t/2)^m}{m!} b(r-m, k).$$

Then

$$\mathbb{E}[A_n(t) | S_n = r] = \frac{\sum_{k=1}^n \rho_k(t) c(r, k) n_{[k]}}{\mathbb{P}(S_n = r)}. \tag{18}$$

The pgfs (10) and (14) and the formula (18) are new.

In a later section, we give a rejection algorithm for simulating the distribution of $A_n(t)$ given $S_n = r$.

2.2. The joint distribution of $\tilde{S}_n(t), \tilde{K}_n(t), A_n(t), A_n^\theta(t)$

A method for computing the stationary joint distribution of (S_n, K_n) is derived in Griffiths (1982), exploiting a diffusion generator in a model with K alleles, then letting $K \rightarrow \infty$. It is found numerically that the joint distribution is strongly diagonal in that, approximately, $S_n = K_n - 1$. It is also shown that $S_n - K_n + 1$ has a proper limit distribution when $n \rightarrow \infty$, even though both S_n and K_n tend to infinity. It is always true that $\tilde{S}_n(t) - \tilde{K}_n(t) + 1 \geq 0$.

Here we consider the joint distribution of $(\tilde{S}_n(t), \tilde{K}_n(t), A_n(t), A_n^\theta(t))$ using a coalescent treatment different from that of Griffiths (1982). To obtain a Markov process, consider $(\tilde{S}_n(t), \tilde{K}_n(t), A_n(t), A_n^\theta(t)), t \geq 0$, beginning from $(0, 0, n, n)$. Let $B_n(t) = A_n(t) - A_n^\theta(t)$. Then

$$(s, k, b, a^\theta) \rightarrow \begin{cases} (s+1, k, b, a^\theta) & \text{at rate } \theta b/2 \\ (s+1, k+1, b+1, a^\theta-1) & \text{at rate } \theta a^\theta/2 \\ (s, k, b-1, a^\theta) & \text{at rate } (b(b-1) + 2ba^\theta)/2 \\ (s, k, b, a^\theta-1) & \text{at rate } a^\theta(a^\theta-1)/2 \end{cases} \tag{19}$$

The total coalescence rate is $a(a-1)/2 = (a^\theta + b)(a^\theta + b - 1)/2$. The total mutation rate is $a\theta/2 = \theta(a^\theta + b)/2$. This process counts mutations and alleles as they arrive back in time from t in a

sample of n and its ancestors, by considering the two groups of lines $a - a^\theta, a^\theta$ and in which groups mutations or coalescences occur.

The simpler process $(A_n(t), A_n^\theta(t))$ has rates

$$(a, a^\theta) \rightarrow \begin{cases} (a - 1, a^\theta) & \text{at rate } (b(b - 1) + 2ba^\theta)/2 \\ (a - 1, a^\theta - 1) & \text{at rate } a^\theta(a^\theta - 1)/2 \\ (a, a^\theta - 1) & \text{at rate } a^\theta \theta/2 \end{cases} \quad (20)$$

with total rate of $(a(a - 1) + a^\theta)/2$. The marginal transition functions are known explicitly from (1) when $\theta = 0$ and $\theta > 0$.

A slightly different approach, described in equation (2.9) of Griffiths (1981), is to consider sample paths that have i mutations and j alleles at t starting from fixed $m = a^\theta, r = a$ as being stationary distributions which satisfy the following recursive system:

$$\begin{aligned} & a(a - 1 + \theta)p(i, j; a^\theta, a) \\ &= (a - a^\theta)\theta p(i - 1, j; a^\theta, a) \\ &+ (a + a^\theta - 1)(a - a^\theta)p(i, j; a^\theta, a - 1) \\ &+ a^\theta \theta p(i - 1, j - 1; a^\theta - 1, a) \\ &+ a^\theta(a^\theta - 1)p(i, j; a^\theta - 1, a - 1), \end{aligned} \quad (21)$$

for $i = 0, 1, \dots, a^\theta; j = 1, 2, \dots$ and $a \geq a^\theta \geq 2$. The boundary probabilities satisfy

$$\begin{aligned} & p(i, j; 1, a) = 0, j > 1 \\ & (a + \theta - 1)p(i, 1; 1, a) = (a - 1)p(i, 1; 1, a - 1) + \theta p(i, 1; 1, a), \\ & \quad \quad \quad a = 2, 3, \dots \\ & p(i, 1; 1; 1) = \delta_{i0}. \end{aligned} \quad (22)$$

Then we want $p(i, j; n, n)$, which can be computed from (21). These equations can be argued from a probabilistic perspective, though a different approach is taken in Griffiths (1981), where the identification $m = a^\theta, r = a$ is not made. (21) is analogous, but more complex, to a similar recursion for the number of mutations which could be written as

$$a(a + \theta - 1)p(i, a) = \theta a p(i - 1; a) + (a - 1)p(i; a - 1). \quad (23)$$

for $i = 0, 1, \dots, a = 2, 3, \dots$

A pgf version of (21) for

$$q(a^\theta, a) = \sum_{i=0}^{\infty} \sum_{j=1}^{i+1} p(i, j; a^\theta, a) z^i w^j$$

appears in Griffiths (1982), eq. (2.8).

3. Simulation-based approaches

In this section, we develop theory that will be used in the rejection sampling and importance sampling schemes. We begin by recalling a simulation method that generates stationary samples of haplotype counts together with the number of mutations in the tree. This works for constant population size coalescent models.

3.1. Growing a tree

A useful way to simulate the ancestral history of haplotype configurations in age order together with the mutations is to use a condensation of an algorithm for growing a gene tree whose nodes are mutations. The algorithm for the tree, described in Ethier and Griffiths (1987), Theorem 5.4 and Griffiths (1989), p. 7, is the following:

Algorithm 1

1. Start with a tree of two leaves (individuals) as two edges joined at a node.

2. When there are m leaves, select one of the m leaves at random and duplicate from the same immediate mutation node with probability $(m - 1)/(\theta + m - 1)$; or add a mutation node on the chosen edge with probability $\theta/(\theta + m - 1)$.
3. To get a sample of n , stop when there are first $n + 1$ leaves and select the configuration just before the last leaf appeared.

Times between events can be added as exponential random variables $T_2^\theta, \dots, T_n^\theta$.

If we just look at haplotype frequencies in age order from the oldest and keep track of s , the number of accumulated mutations, the state space is $(m_1, \dots, m_k; s)$ and transitions are Markovian. The condensed algorithm is the following:

Algorithm 2

1. Start with a configuration $(n_1 = 2; 0)$ of two identical oldest haplotypes and no mutations.
2. When there are m individuals make a transition $(m_1, \dots, m_k; s) \rightarrow (m_1, \dots, m_j + 1, \dots, m_k; s)$ with probability $(m_j/m)(m - 1)/(m + \theta - 1)$; or $(m_1, \dots, m_k; s) \rightarrow (m_1, \dots, m_i - 1, \dots, m_k, 1; s + 1)$ with probability $(m_i/m)\theta/(m + \theta - 1)$.
3. To get a sample of n , stop when there are first $n + 1$ leaves and select the configuration just before the last leaf appeared.

This algorithm is useful for simulation of an ancestral path forward in time which contains full information of haplotype count configurations and mutations.

Let $p^\circ(\mathbf{m})/\prod_{j=1}^m \alpha_j!$ be the probability of a non-age labelled configuration \mathbf{m} (labelled in an arbitrary order), with α_j the number of allele frequencies equal to j , under the Markov chain without the stopping rule. Then

$$\begin{aligned} p^\circ(\mathbf{m}; s) &= \frac{m - 2}{m + \theta - 2} \sum_{m_j > 1} \frac{m_j - 1}{m - 1} p^\circ(m_1, \dots, m_j - 1, \dots, m_k; s) \\ &+ \frac{\theta}{m + \theta - 1} \sum_{i, l: m_i = 1} \frac{m_l + 1 - \delta_{li}}{m} p^\circ(\mathbf{m} + \mathbf{e}_l - \mathbf{e}_i; s - 1). \end{aligned} \quad (24)$$

The sample probability $p(\mathbf{n}; s)/\prod_{j=1}^n \alpha_j!$ is such that

$$p(\mathbf{n}; s) = \frac{n - 1}{n + \theta - 1} p^\circ(\mathbf{n}; s).$$

A recursion is therefore

$$\begin{aligned} p(\mathbf{n}; s) &= \frac{n - 1}{n + \theta - 1} \sum_{n_j > 1} \frac{n_j - 1}{n - 1} p(\mathbf{n} - \mathbf{e}_j; s) \\ &+ \frac{\theta}{n + \theta - 1} \sum_{i, l: n_i = 1} \frac{n_l + 1 - \delta_{li}}{n} p(\mathbf{n} - \mathbf{e}_i + \mathbf{e}_l; s - 1). \end{aligned} \quad (25)$$

The recursion in (25) is used in calculating the importance sampling weights in Section 4.

Note that the last sum includes the case when $i = l$, and it can be written as

$$\sum_{i \neq l: n_i = 1} \frac{n_l + 1}{n} p(\mathbf{n} - \mathbf{e}_i - \mathbf{e}_l; s - 1) + \alpha_1 \frac{1}{n} p(\mathbf{n}; s - 1).$$

If a, s are the number of haplotypes and mutations, respectively, with \mathbf{n} then $s - a + 1 \geq 0$ so $p(\mathbf{n}; s) = 0$ if $s - a + 1 < 0$ in the recursive equations. The Ewens Sampling Formula Ewens (1972) satisfies a similar equation to (25), summing over s , with

$$p(\mathbf{n}; \cdot) = \frac{n!}{n_1 \cdots n_k} \cdot \frac{\theta^k}{\theta_{(n)}}.$$

The probability of the sample configuration is

$$p(\mathbf{n}; \cdot) / \prod_{j=1}^n \alpha_j! = \frac{n!}{\prod_{j=1}^k n_j \alpha_j!} \cdot \frac{\theta^k}{\theta_{(n)}} = \frac{n!}{\prod_{j=1}^k j^{\alpha_j} \alpha_j!} \cdot \frac{\theta^k}{\theta_{(n)}}.$$

The joint distribution of $(\mathbf{n}; s)$ is studied in Innan et al. (2005), where their Eq. (1) is essentially the same as (25). The authors then use these recursive equations to derive exact probabilities for small sample sizes. In Section 4, we use an importance sampling technique based on (25) which allows calculation for much larger sample sizes than are possible with exact calculation and we also carry out ancestral inference back in time.

3.2. The number of ancestors and mutations

The simplest ancestral question is to ask about the distribution of $A_n(t)$ conditional on $S_n = s$ segregating sites in the sample of size n . Rejection algorithms work well for problems like this, as was illustrated by Tavaré et al. (1997). They took a general Bayesian approach in which θ was given a prior distribution and the times T_n, \dots, T_2 come from a variable population size coalescent model. The combinatorics for a general binary coalescent tree are studied in Griffiths and Tavaré (1998), where the sample frequency spectrum and the mean age of a mutation known to have b descendants in a sample of n genes are calculated.

We illustrate by describing how to generate observations from the conditional distribution of $(\theta, A_n(t))$ conditional on $S_n = s$. To this end, let $W_1 = T_n, W_2 = T_n + T_{n-1}, \dots, W_j = T_n + \dots + T_{n-j+1}, \dots, W_{n-1} = T_n + \dots + T_2, W_{n-1}$ is the height of the coalescent tree. Define the total length of the tree as $L_n = nT_n + \dots + 2T_2$. Form a set of bins as follows:

$$B_1 = (0, W_1], \dots, B_{n-1} = (W_{n-2}, W_{n-1}], B_n = (W_{n-1}, \infty).$$

Let J be the bin that covers t . Then the number of ancestors at time t is

$$A_n(t) = n - J + 1, \tag{26}$$

and the length of the coalescent tree from 0 back to t is

$$\tilde{L}_n(t) = \begin{cases} nt & \text{if } J = 1 \\ L_n & \text{if } J = n \\ \sum_{l=n-J+2}^n lT_l + (n - J + 1)(t - W_{J-1}) & \text{if } 2 \leq J \leq n - 1. \end{cases} \tag{27}$$

Finally, define the length of the coalescent tree from time t to the most recent common ancestor as $L_n(t) = L_n - \tilde{L}_n(t)$. Let $\tilde{S}_n(t)$ be the number of mutations arising in $(0, t)$ and let $S_n(t)$ the number of mutations from t to the most recent common ancestor. Then, conditional on T_n, \dots, T_2 and J ,

$$\tilde{S}_n(t) \sim \text{Po}(\theta \tilde{L}_n(t)/2), \quad S_n(t) \sim \text{Po}(\theta L_n(t)/2),$$

where $\text{Po}(\lambda)$ denotes the Poisson distribution with parameter λ , such that

$$\text{Po}(\lambda)\{s\} = \frac{e^{-\lambda} \lambda^s}{s!}, \quad s = 0, 1, \dots,$$

and $\tilde{S}_n(t)$ and $S_n(t)$ are conditionally independent. The total number of segregating sites in the sample at time 0 is $S_n = \tilde{S}_n(t) + S_n(t)$. The simplest algorithm gives the probability distribution of J , and therefore the distribution of $A_n(t)$, conditional on $S_n = s$. This gives the following:

Table 1

Result of 10,000 runs for the constant population size coalescent model with a sample of size $n = 1544$ and $s = 9$ segregating sites. Table shows average value of $A_n(t), S_n(t)$ given $S_n = 9$. Righthand columns give SE of the mean.

t	$S_n(t)$	SE	$A_n(t)$	SE
0	9.0		1544	
0.1	3.18	0.031	19.47	0.051
0.5	1.11	0.023	3.72	0.023
1.0	0.38	0.015	1.78	0.015
1.5	0.12	0.009	1.25	0.009

Algorithm 3

1. Simulate θ from the prior, $\pi(\cdot)$
2. Simulate T_n, \dots, T_2 from an appropriate coalescent model
3. Compute $J, A = A_n(t) = n - J + 1, L_n$
4. Accept (θ, A) as an observation from the posterior with probability

$$h = \frac{\text{Po}(\theta L_n/2)\{s\}}{\text{Po}(s)\{s\}}$$

We also note that the same rejection approach may be used to approximate conditional distributions for many other ancestral variables. For example, to study the distribution of the number of mutations $S_n(t)$ present in the ancestors at time t , which is a measure of the standing variation at time t , we can use the following.

Algorithm 4

1. Simulate θ from the prior, $\pi(\cdot)$.
2. Simulate T_n, \dots, T_2 from an appropriate coalescent model.
3. Compute $J, A = A_n(t) = n - J + 1, L_n(t)$ and $\tilde{L}_n(t)$, defined in (27).
4. Accept (θ, A) with probability

$$h = \frac{\text{Po}(\theta L_n/2)\{s\}}{\text{Po}(s)\{s\}}$$

and else return to 1.

5. Simulate S from a Binomial($s, L_n(t)/L_n$) distribution, and return (θ, A, S) as an observation from the posterior of $(\theta, A_n(t), S_n(t))$ given $S_n = s$.

To address the problems studied in the first section of the paper, we may treat θ as fixed in this approach (that is, as having a degenerate prior).

Blum and Rosenberg (2007) construct a rejection algorithm for maximum likelihood estimation of the number of ancestral lineages time t ago, based on the sample frequency spectrum. The algorithm is constructed differently from Algorithms 3 and 4.

3.3. Hammer et al. example

As an illustration we consider the Y chromosome data of 1544 sequences from Hammer et al. (1998). In that paper a perfect phylogeny was constructed from the sequence data and the program GENETREE was used to find the TMRCA and ages of mutations in the ancestral tree shown in Figure 7 of Hammer et al. (1998) with $\theta = 2.5$. There were 9 segregating sites and 10 haplotypes observed in the data. The unconditional expected height of the coalescent tree is 2 time units, and the expected height conditional on $s = 9$ segregating sites is 1.21 units, this latter found from the method in Tavaré et al. (1997).

We generated 10,000 repetitions of the previous algorithm for a series of times t , obtaining the information in Table 1 for the conditional expectations of $S_n(t)$ and $A_n(t)$.

The simple rejection schemes illustrated here are not as useful for considering more complicated summaries of the data. In the

next section, we show how to exploit an importance sampling approach to derive conditional distributions given the haplotype frequency distribution and the number of mutations.

4. Importance sampling

Sequential importance sampling for ancestral inference in population genetics has a long history, illustrated by Griffiths and Tavaré (1994a, b), Griffiths and Tavaré (1999), Felsenstein et al. (1999), Stephens and Donnelly (2000), Griffiths (2002) and De Iorio and Griffiths (2004). The technique can be described as constructing a proposal distribution for events back in time, simulating back in time, then correcting for the approximate proposal distribution by calculating the exact probability of the path forward in time and taking the ratio of the probability of the forward path divided by the approximate probability of the backward path as the importance weight. If there are r simulation runs then an empirical ancestral history is returned as $(\hat{p}_1, \mathcal{H}_1), \dots, (\hat{p}_r, \mathcal{H}_r)$ where $\{\hat{p}_j\}$ are the importance weights scaled to add to unity and $\{\mathcal{H}_j\}$ are the histories. A general reference to sequential importance sampling is Liu (2001).

Choosing a proposal distribution is an art. We use the principal of choosing a lineage which can be involved in a transition back in time uniformly. This has a theoretical justification, described in Stephens and Donnelly (2000) and De Iorio and Griffiths (2004). Sequential importance sampling for a haplotype configuration can be regarded as a simplification of the technique used for a complete genetree, constructed as a perfect phylogeny from the pattern of mutations on DNA sequences. The simpler scheme counts different haplotypes of the sequences, the extra information being the number of mutations back to the most recent common ancestor. Time information, such as coalescence times, ages of mutations and time to the most recent ancestor, can be included. Details of how to include time in an importance sampling algorithm are given in Griffiths and Tavaré (1997). Here we develop a new importance sampling approach for the Kingman coalescent models conditional on an observed configuration \mathbf{n}, s .

The proposal distribution \hat{p} for reverse transitions in a haplotype history is detailed in the following equations. Suppose a current configuration is $\mathbf{n} = (n_1, \dots, n_k)$, the number of mutations to the most recent common ancestor is s , and the number of singletons is q .

For $k > 2$, if $s - k + 1 > 0$,

$$\hat{p}(\mathbf{n} - \mathbf{e}_i; s | \mathbf{n}; s) = \frac{n_i}{n} \text{ if } n_i > 1 \tag{28}$$

$$\hat{p}(\mathbf{n} - \mathbf{e}_i + \mathbf{e}_l; s - 1 | \mathbf{n}; s) = \frac{n_l}{n} \cdot \frac{1}{n}, \quad n_i = 1, l \neq i$$

$$\hat{p}(\mathbf{n}; s - 1 | \mathbf{n}; s) = \frac{q}{n} \cdot \frac{1}{n}, \tag{29}$$

or if $s - k + 1 = 0$ then

$$\hat{p}(\mathbf{n} - \mathbf{e}_i; s | \mathbf{n}; s) = \frac{n_i}{n} \text{ if } n_i > 1$$

$$\hat{p}(\mathbf{n} - \mathbf{e}_i + \mathbf{e}_l; s - 1 | \mathbf{n}; s) = \frac{n_l}{n - 1} \cdot \frac{1}{n}, \quad n_i = 1, l \neq i. \tag{30}$$

The first factors in (29) involve a choice of either mutations that define allele types and those which appear on lineages between defined alleles. Importance weights for transitions back in time are therefore

$$(\mathbf{n}, s) \rightarrow (\mathbf{n} - \mathbf{e}_i, s), n_j > 1: \frac{n_i - 1}{n + \theta - 1} \cdot \frac{1}{\hat{p}(\mathbf{n} - \mathbf{e}_i, s | \mathbf{n}, s)}$$

$$(\mathbf{n}, s) \rightarrow (\mathbf{n} - \mathbf{e}_i + \mathbf{e}_l, s - 1), n_k = 1, k \neq l:$$

$$\frac{\theta}{n + \theta - 1} \cdot \frac{n_l + 1}{n} \cdot \frac{1}{\hat{p}(\mathbf{n} - \mathbf{e}_i + \mathbf{e}_l, s - 1 | \mathbf{n}, s)}$$

$$(\mathbf{n}, s) \rightarrow (\mathbf{n}, s - 1): \frac{\theta}{n + \theta - 1} \cdot \frac{1}{n} \cdot \frac{1}{\hat{p}(\mathbf{n}, s - 1 | \mathbf{n}, s)}$$

When $k = 2$ we have to consider the following possible cases:

(a) If $n_1 > 1, n_2 > 1$

$$\hat{p}(\mathbf{n} - \mathbf{e}_i; s | \mathbf{n}; s) = \frac{n_i}{n}, \quad i = 1, 2,$$

(b) if $n_1 > 1, n_2 = 1, s > 1$,

$$\hat{p}(\mathbf{n} - \mathbf{e}_1; s | \mathbf{n}; s) = \frac{n_1}{n},$$

$$\hat{p}(\mathbf{n}; s - 1 | \mathbf{n}; s) = \frac{1}{n},$$

(c) if $n_1 > 1, n_2 = 1, s = 1$,

$$\hat{p}(\mathbf{n} - \mathbf{e}_1; 1 | \mathbf{n}; 1) = \frac{n_1}{n},$$

$$\hat{p}(\mathbf{n} + \mathbf{e}_1 - \mathbf{e}_2, 0 | \mathbf{n}; 1) = \frac{1}{n},$$

(d) similarly when $n_1 = 1, n_2 > 1$,

(e) if $n_1 = 1, n_2 = 1, s > 1$,

$$\hat{p}(\mathbf{n}; s - 1 | \mathbf{n}; s) = 1,$$

(f) if $n_1 = 1, n_2 = 1, s = 1$

$$\hat{p}(\mathbf{n} + \mathbf{e}_1 - \mathbf{e}_2; 0 | \mathbf{n}; 1) = 1.$$

Importance weights are

$$(a) (\mathbf{n}, s) \rightarrow (\mathbf{n} - \mathbf{e}_i, s): \frac{n_i - 1}{n + \theta - 1} \cdot \frac{1}{\hat{p}(\mathbf{n} - \mathbf{e}_i, s | \mathbf{n}, s)}, \quad n_1, n_2 > 1$$

$$(b) (\mathbf{n}, s) \rightarrow (\mathbf{n} - \mathbf{e}_1, s): \frac{n_1 - 1}{n + \theta - 1} \cdot \frac{1}{\hat{p}(\mathbf{n} - \mathbf{e}_1, s | \mathbf{n}, s)},$$

$$n_1 > 1, n_2 = 1, s > 1$$

$$(\mathbf{n}, s) \rightarrow (\mathbf{n}, s - 1): \frac{\theta}{n + \theta - 1} \cdot \frac{1}{n} \cdot \frac{1}{\hat{p}(\mathbf{n}, s - 1 | \mathbf{n}, s)}$$

$$(c) (\mathbf{n}, 1) \rightarrow (\mathbf{n} - \mathbf{e}_1, 1): \frac{n_1 - 1}{n + \theta - 1} \cdot \frac{1}{\hat{p}(\mathbf{n} - \mathbf{e}_1, 1 | \mathbf{n}, 1)},$$

$$n_1 > 1, n_2 = 1, s = 1$$

$$(\mathbf{n}, 1) \rightarrow (\mathbf{n} + \mathbf{e}_1 - \mathbf{e}_2, 0): \frac{\theta}{n + \theta - 1} \cdot \frac{1}{\hat{p}(\mathbf{n} + \mathbf{e}_1 - \mathbf{e}_2, 1 | \mathbf{n}, 1)}$$

$$(e) (\mathbf{n}, s) \rightarrow (\mathbf{n}, s - 1): \frac{\theta}{1 + \theta} \cdot \frac{1}{\hat{p}(\mathbf{n}, s - 1 | \mathbf{n}, s)},$$

$$n_1 = 1, n_2 = 1, s > 1$$

$$(f) (\mathbf{n}, 1) \rightarrow (\mathbf{n} + \mathbf{e}_1 - \mathbf{e}_2, 0): \frac{2\theta}{1 + \theta} \cdot \frac{1}{\hat{p}(\mathbf{n} + \mathbf{e}_1 - \mathbf{e}_2, 0 | \mathbf{n}, 1)},$$

$$n_1 = 1, n_2 = 1, s = 1.$$

4.1. Implementation

Our implementation provides

- The probability of a sample configuration of haplotypes and number of segregating sites. This is an extension of the Ewens Sampling Formula, which is the probability of the configuration of haplotypes.

The next calculations are conditional on the configuration of haplotypes and segregating sites at time 0.

- The average coalescence times in the past.
- The average mutation times in the past.
- The average times when alleles are lost in the past. (The time of loss of the last haplotype is truncated at the TMRCA if not lost by mutation.)

- The average allele ages in the past.
- The distribution of ancestor lines and the average allele configuration at a given time in the past.

The program also implements a variable population size option with exponential growth. Coalescent times then have a distribution that depends on the time when they occur. We do not go into detail here, but refer the reader to Griffiths and Tavaré (1994c). The analogue of the Ewens Sampling Formula in this case is derived in Griffiths and Lessard (2005).

Accuracy of the implementation was checked by ensuring that for smaller sample sizes the equation

$$p(\mathbf{n}) = \sum_{s=k-1}^B p(\mathbf{n}; s)$$

was approximately satisfied, where k is the number of alleles in \mathbf{n} and B is a suitable upper bound. The simulation variance, starting with different seeds, was observed to be small for sample sizes such as in Section 4.2 that follows.

4.2. Hammer et al. example, continued

We continue with the example started in Section 3.3. The Y haplotype data of $n = 1,544$ sequences from Hammer et al. (1998), had 10 haplotypes and $s = 9$ segregating sites. We continue to use their value of $\theta = 2.5$ for illustration. The 10 haplotype frequencies are

21 23 853 188 75 1 68 31 67 217

in the lineage order shown in Figure 7 of Hammer et al. (1998). The average values in the tables below are conditional on the configuration and number of segregating sites, thereby extending the results of Section 3.3.

The Appendix describes the input for the implementation of the method. A command line of

```
esf_st1 HammerHap.dat 10 9 2.5 1000000 93849 -a
```

in which the input file HammerHap.dat contains the haplotype frequencies in the order above, produces the output described below; the average coalescence times are not shown.

Two runs with different seeds gave identical output to three significant places, showing some confidence in the output. The probability of obtaining the sample configuration and $s = 9$ segregating sites was 1.4785×10^{-19} . As a comparison, the probability of the sample configuration, calculated from the Ewens Sampling Formula, was 1.1722×10^{-18} . The mean TMRCA, conditional on the data, in coalescent units was 1.15, which may be compared to the value of 1.21 obtained in Section 3.3.

4.2.1. Stationary properties

Here we record some information about the sample at time 0.

The conditional expected mutation times in increasing time order are

0.003 0.022 0.039 0.062 0.094 0.142 0.219 0.360 0.675

while the conditional expected haplotype loss times in increasing time order are

0.003 0.022 0.039 0.062 0.094 0.142 0.219 0.360 0.761.

Most of the tree structure has developed by an average time of less than 1.00 coalescent time unit.

The conditional expected haplotype ages, in the order they are listed above, are

0.051 0.092 0.995 0.406 0.216 0.007 0.201 0.114 0.200 0.446.

These are monotonic in the number of copies of the haplotype in the sample, confirming the intuition that common haplotypes tend to be older.

Table 2

Result of 1000,000 runs for the constant population size coalescent model with a sample of size $n = 1544$ and 9 segregating sites, and haplotype frequencies given above. Table shows average values of $K_n(t)$, $A_n(t)$, $S_n(t)$ conditional on the haplotype frequencies and $S_n = 9$.

t	$K_n(t)$	$S_n(t)$	$A_n(t)$
0	10	9	1544
0.1	5.09	4.09	19.9
0.5	1.84	0.85	3.94
1.0	0.74	0.17	1.75
1.5	0.21	0.03	1.19

Table 3

Relative errors for estimates from Tables 1 and 2.

t	$S_n(t)$	$A_n(t)$
0.1	0.25	0.02
0.5	0.75	0.06
1.0	0.75	0.02
1.5	1.2	0.05

4.2.2. Time-varying properties

At a given time t in the past, the distribution of the configuration, number of mutations, and number of ancestral lineages conditional on the current configuration and number of mutations can be calculated by the importance sampling program. We illustrate this by considering time points $t = 0.1, 0.5, 1.0, 1.5$ and taking averages at those times.

We begin by comparing the conditional distribution of $A_n(t)$ and $S_n(t)$ with the analogous results in Table 1. Additionally, Table 2 shows the average of the number of haplotypes, $K_n(t)$, present at time t .

The results in the second and third columns should be compared with those in Table 1; they show qualitatively the same results. Table 3 shows the relative errors for estimates of $S_n(t)$ and $A_n(t)$. For example, letting superscripts 1 and 2 denote estimates from Tables 1 and 2, the relative error for $S_n(t)$ is $|S_n^1(t) - S_n^2(t)| / \frac{1}{2}(S_n^1(t) + S_n^2(t))$.

Table 4 gives the number of haplotype counts at different times t in the past. Haplotypes decrease because of coalescence and types are eventually lost when their defining mutation takes place. In Table 5 we give the distribution of the number of ancestral lines at different times t in the past, conditional on the current haplotype configuration and the number of segregating sites. The most interesting time configurations are when $t \leq 0.1$; afterwards the number of lineages and number of haplotypes decrease rapidly.

4.3. A 1000 Genomes Y chromosome dataset

A larger Y chromosome dataset comes from the 1000 Genomes Project. An analysis of these data is made in Poznik et al. (2016), where a phylogeny is constructed. The paper concludes that the data show evidence of expansion. As an example of our approach, we consider a subset of this dataset consisting of sequences in the A, B and E haplotype groups. These are the three oldest groups in the phylogeny, and are composed of 334 sequences. There may be explanations other than expansion for the data configuration, such as a non-random choice of individuals, however we will assume a random sample for this example analysis. For illustration we focus on the TBL1Y gene, composed of some 180,000bp, and containing 278 biallelic SNPs. The haplotype configuration, with α_j equal to the number of alleles of multiplicity j , is given in Table 6.

Watterson's estimate of θ (Watterson, 1975) based on the number of segregating sites s in a constant-sized population is

$$\hat{\theta}_W = \frac{s}{\sum_{j=1}^{n-1} 1/j} = 44$$

Table 4
Extant haplotype counts for the Hammer data at different times in the past.

Time t	Haplotype frequency									
0	21	23	853	188	75	1	68	31	67	217
0.1	0.227	0.250	11.7	2.30	0.862	0.010	0.777	0.340	0.765	2.69
0.5	0.033	0.036	2.63	0.372	0.130	0.002	0.117	0.050	0.115	0.441
1.0	0.015	0.011	0.837	0.140	0.050	0.001	0.045	0.020	0.044	0.165
1.5	0.006	0.003	0.206	0.043	0.016	0.000	0.014	0.006	0.014	0.051

Table 5
Average number of ancestral lineages in the Hammer data at time t in the past.

t	12	13	14	15	16	17	18	19	20
0.1	0.001	0.003	0.009	0.023	0.048	0.083	0.121	0.149	0.159
	21	22	23	24	25	26	27	28	
	0.142	0.111	0.074	0.044	0.022	0.010	0.004	0.001	
0.5	1	2	3	4	5	6	7	8	
	0.011	0.088	0.259	0.337	0.216	0.073	0.014	0.002	
1.0	1	2	3	4	5				
	0.426	0.414	0.143	0.016	0.001				
1.5	1	2	3						
	0.826	0.163	0.011						

The maximum likelihood estimate $\hat{\theta}_E$ based on the Ewens' sampling formula uses $k = \sum_{j=1}^n \alpha_j$, a sufficient statistic for θ . $\hat{\theta}_E$ satisfies

$$k = 1 + \sum_{j=1}^{n-1} \frac{\hat{\theta}_E}{\hat{\theta}_E + j}.$$

In the TBL1Y dataset $\hat{\theta}_E = 82$. The large number of singletons α_1 in the data suggests exponential growth in the population. Growth produces a *star shaped* coalescent tree, which leads to a greater number of singleton sequences. The mean number of singletons in the constant size population setting is

$$\mathbb{E}[\alpha_1] = \frac{n\theta}{n + \theta - 1}.$$

If $\theta = 82$, then $\mathbb{E}[\alpha_1] = 66$, which is much less than the observed $\alpha_1 = 107$. Tajima's D (Tajima, 1983) is given by

$$D = \frac{\pi - \hat{\theta}_W}{\sqrt{\widehat{\text{var}}(\pi - \hat{\theta}_W)}}$$

where π is the average number of pairwise differences, an unbiased estimate of θ . This may be used to test for population growth or other departures from the coalescent model with no growth. Large negative values of D indicate population growth. In our data $\pi = 6.49$ and $D = -2.6$, consistent with expansion. In the Appendix we describe another statistic for testing the no-growth model based on the frequency spectrum, and particularly on the number of singletons, when θ is large. Achaz (2009) develops neutrality tests based on the frequency spectrum which generalize tests based on the number of segregating sites and Tajima's D .

In this large dataset it is difficult to obtain a very precise estimate of θ and growth rate β because there is a large amount of variation in the importance sampling scheme due to the number and length of the sequences. This is not so much an importance sampling issue, but due to the size and structure of the data. Random subsets of the data could be chosen, but the main feature of the data is the large number of singletons. Ancestral inference of the number of ancestral lines at t back would also be difficult to interpret for subsamples. We try a large value $\theta = 100$ with different growth rates β . Growth decreases the variation in the sample, but increases the proportion of singletons, because the coalescent lengths are shortened and the tree is star shaped. Increasing θ with growth keeps the variation as well as increasing the number of

Table 6
Allele multiplicities observed in the TBL1Y dataset.

α_1	α_2	α_3	α_4	α_5	α_6	α_7	α_{14}	α_{32}	α_{50}	α_{61}
107	12	6	1	1	2	1	1	1	1	1

Table 7
Likelihoods with expansion in the TBL1Y data.

β	Replicate 1	Replicate 2	Average
0	2.0934e-61	3.4592e-62	1.2197e-61
0.5	1.2203e-60	1.0782e-61	6.6405e-61
1.0	4.1350e-60	1.1497e-60	2.6424e-60
1.5	9.2120e-61	2.8962e-61	6.0541e-61
2.0	1.5297e-61	5.9120e-61	3.7209e-61
2.5	4.9886e-62	3.7215e-62	4.3551e-62

Table 8
Allele age within groups.

Haplotype groups					
α_1	α_2	α_3	α_4	α_5	α_6
0.0123	0.0086	0.0222	0.0249	0.0421	0.0390
α_7	α_{14}	α_{32}	α_{50}	α_{61}	
0.0290	0.1338	0.1331	0.1259	0.1304	

singletons. The likelihood of the allele configuration and number of segregating sites was calculated for $\theta = 100$ with several values of the growth rates β . Two different replicates each with 10 million runs gave the results in Table 7.

A plausible maximum likelihood estimate when $\theta = 100$ is $\hat{\beta} = 1.0$. If θ and β are increased together, it is possible that the likelihood estimates fall on a ridge. In the first replicate with these values of θ and β the TMRCA was 1.461 and the average of ages within haplotype groups is shown in Table 8.

With the large value of $\theta = 100$, times where mutation creates an allele are close to the leaves of the tree.

The scale of modern molecular datasets points out the difficulty of exact inference techniques, and highlights the need for alternative approaches. Among these are Approximate Bayesian Computation (ABC) and other advanced simulation methods, to which Paul Joyce made several contributions. Two examples are estimating evolutionary rates of trait evolution by ABC (Slater et al., 2012), and a perfect simulation method for simulation from a non-neutral high dimensional distribution of allele frequencies (Joyce et al., 2012).

Acknowledgements

We thank Chris Tyler-Smith, Yali Xue and two referees for their helpful comments about the 1000 Genomes Y chromosome data.

Code for the rejection and importance sampling methods may be obtained from the authors.

Appendix

Importance sampling code

The input for the importance sampling method illustrated in Section 4.2 is

```
esf_stl configfile k [#alleles] m [#mutations] theta replicates seed
```

Options

```
-g beta [exponential growth]
-a [age information]
-t time [Configuration at time]
```

Bob Griffiths 4 May 2017, Version 1.7

Poisson approximation for large θ

Motivated by the discussion in Section 4.3, we discuss the behaviour of the Ewens sampling Formula for large values of θ and n . The Ewens Sampling Formula gives the distribution of the number of haplotypes and their frequencies in a sample taken from a constant-size population. Writing α_j for the number of haplotypes with frequency j , the distribution is

$$p_E(\alpha_1, \alpha_2, \dots, \alpha_n) = \frac{n!}{\theta_{(n)}} \prod_{j=1}^n \binom{\theta}{j}^{\alpha_j} \frac{1}{\alpha_j!}, \quad (31)$$

where $\theta_{(n)} := \theta(\theta+1)\cdots(\theta+n-1)$, and $\alpha_1 + 2\alpha_2 + \cdots + n\alpha_n = n$. The formula (31) shows that, were it not for the condition that $\alpha_1 + 2\alpha_2 + \cdots + n\alpha_n = n$, the α_j would be independent Poisson random variables with mean θ/j . Indeed, for fixed θ it is known that for any $b = o(n)$ as $n \rightarrow \infty$, the total variation distance between the distribution of $(\alpha_1, \dots, \alpha_b)$ and that of (Z_1, \dots, Z_b) , for independent Poisson random variables with $\mathbb{E}[Z_j] = \theta/j$, is $O(b/n)$ as $n \rightarrow \infty$. See Arratia et al. (2003), Theorem 5.2.

Here we consider the case in which $\theta \rightarrow \infty$ with n , and we show that for fixed b , $(\alpha_1, \dots, \alpha_b)$ has asymptotically the distribution of $(Z_1^\theta, \dots, Z_b^\theta)$, where

$$\mathbb{E}Z_j^\theta = \frac{\theta}{j} \left(\frac{n}{n+\theta} \right)^j, \quad j = 1, \dots, b.$$

To see this, consider the joint falling factorial moments of $\alpha_1, \dots, \alpha_b$, given by Watterson (1974) as

$$\begin{aligned} \mathbb{E} \prod_{j=1}^b (\alpha_j)_{[r_j]} &= \mathbb{1}(m \leq n) \frac{n! \Gamma(\theta + n - m)}{(n - m)! \Gamma(\theta + n)} \prod_{j=1}^b \left(\frac{\theta}{j} \right)^{r_j}, \\ &= \mathbb{1}(m \leq n) \frac{n(n-1)\cdots(n-m)}{(\theta+n-1)\cdots(\theta+n-m)} \prod_{j=1}^b \left(\frac{\theta}{j} \right)^{r_j} \\ &\sim \left(\frac{\theta}{\theta+n} \right)^m \prod_{j=1}^b \left(\frac{\theta}{j} \right)^{r_j} \\ &= \prod_{j=1}^b \left(\frac{\theta}{j} \left(\frac{n}{\theta+n} \right) \right)^{r_j} \end{aligned}$$

where $m := r_1 + 2r_2 + \cdots + br_b$. The term on the right gives the falling factorial moments of $(Z_1^\theta, \dots, Z_b^\theta)$, and the result follows from the method of moments.

In practice, different limit laws are obtained depending on the way θ varies with n . For example, if $\theta \sim \eta n$, then α_1 has approximately a Poisson distribution with mean $\theta/(1+\eta)$. For the data in Section 4.3, with $\theta = 82$, the number of singleton haplotypes has mean $82 \cdot (334/416) \approx 65.84$. Since the probability of observing 107 or more singletons is then $\approx 1.92 \times 10^{-6}$, we conclude that the constant-size model does not provide an adequate fit. In a similar spirit, $\alpha_1 + \alpha_2$ has approximately a Poisson distribution with mean 92.27. We observed $\alpha_1 + \alpha_2 = 119$, the probability of a larger value being ≈ 0.0043 ; once more, this suggests that the constant-size model is not a good fit.

References

- Achaz, G., 2009. Frequency spectrum neutrality tests: one for all and all for one. *Genetics* 183, 249–258.
- Arratia, R., Barbour, A.D., Tavaré, S., 2003. Logarithmic Combinatorial Structures: A Probabilistic Approach. In: Monographs in Mathematics, European Mathematical Society.
- Blum, M.G., Rosenberg, N.A., 2007. Estimating the number of ancestral lineages using a maximum-likelihood method based on rejection sampling. *Genetics* 176, 1741–1757.
- De Iorio, M., Griffiths, R.C., 2004. Importance sampling on coalescent histories. I. *Adv. Appl. Probab.* 36, 417–433.
- Ethier, S.N., Griffiths, R.C., 1987. The infinitely-many-sites model as a measure-valued diffusion. *Ann. Probab.* 15, 515–545.
- Felsenstein, J., Kuhner, M., Yamato, J., Beerli, P., 1999. Likelihoods on coalescents: a Monte Carlo sampling approach to inferring parameters from population samples of molecular data. *IMS Lect. Notes Monogr. Ser.* 33, 163–185.
- Griffiths, R.C., 1980. Lines of descent in the diffusion approximation of neutral Wright-Fisher models. *Theor. Popul. Biol.* 17, 37–50.
- Griffiths, R.C., 1981. Transient distribution of the number of segregating sites in a neutral infinite-sites model with no recombination. *J. Appl. Probab.* 18, 42–51.
- Griffiths, R.C., 1982. The number of alleles and segregating sites in a sample from the infinite-alleles model. *Adv. Appl. Probab.* 14, 225–239.
- Griffiths, R.C., 1989. Genealogical-tree probabilities in the infinitely-many-sites model. *J. Math. Biol.* 27, 667–680.
- Griffiths, R.C., 2002. Ancestral inference from gene trees. In: Veuille, M., Slatkin, M. (Eds.), *Modern Developments in Theoretical Population Genetics: The Legacy of Gustave Malécot*. Oxford University Press, New York, pp. 94–117.
- Griffiths, R.C., 2006. Coalescent lineage distributions. *Adv. Appl. Probab.* 38, 405–429.
- Griffiths, R.C., Lessard, S., 2005. Ewens' sampling formula and related formulae: Combinatorial proofs, extensions to variable population size and applications to ages of alleles. *Theor. Popul. Biol.* 68, 167–177.
- Griffiths, R.C., Tavaré, S., 1994a. Simulating probability distributions in the coalescent. *Theor. Popul. Biol.* 46, 131–159.
- Griffiths, R.C., Tavaré, S., 1994b. Ancestral inference in population genetics. *Stat. Sci.* 9, 307–319.
- Griffiths, R.C., Tavaré, S., 1994c. Sampling theory for neutral alleles in a varying environment. *Phil. Trans. R. Soc. Lond. B* 344, 403–410.
- Griffiths, R.C., Tavaré, S., 1997. Computational methods for the coalescent. In: Donnelly, P., Tavaré, S. (Eds.), *Progress in Population Genetics and Human Evolution*. In: IMA Volumes in Mathematics and its Applications, vol. 87, Springer Verlag, Berlin, pp. 165–182.
- Griffiths, R.C., Tavaré, S., 1998. The age of a mutation in a general coalescent tree. *Stoch. Models* 14, 273–295.
- Griffiths, R.C., Tavaré, S., 1999. The ages of mutations in gene trees. *Ann. Appl. Probab.* 9, 567–590.
- Hammer, M.F., Karafet, T., Rasanayagam, A., Wood, E.T., Altheide, T.K., Jenkins, T., Griffiths, R.C., Templeton, A.R., Zegura, S.L., 1998. Out of Africa and back again: Nested clastic analysis of human Y chromosome variation. *Mol. Biol. Evol.* 15, 427–441.
- Innan, H., Zhang, K., Marjoram, P., Tavaré, S., Rosenberg, N.A., 2005. Statistical tests of the coalescent model based on the haplotype frequency distribution and the number of segregating sites. *Genetics* 169, 1763–1777.
- Joyce, P., Genz, A., Buzbas, E.O., 2012. Efficient simulation and likelihood methods for non-neutral multi-allele models. *J. Comput. Biol.* 19, 650–661.
- Kingman, J.F.C., 1982. On the genealogy of large populations. *J. Appl. Probab.* 19A, 27–43.
- Liu, J.S., 2001. *Monte Carlo Strategies in Scientific Computing*. Springer, New York.

- Poznik, G.D., Xue, Y., Mendez, F.L., Willems, T.F., Massaia, A., Sayres, M.A.W., Ayub, Q., McCarthy, S.A., Narechania, A., Kashin, S., Chen, Y., Banerjee, R., Rodriguez-Flores, J.L., Cerezo, M., Shao, H., Gymrek, M., Malhotra, A., Louzada, S., Desalle, R., Ritchie, G.R.S., Cerveira, E., Fitzgerald, T.W., Garrison, E., Marcketta, A., Mittelman, D., Romanovitch, M., Zhang, C., Zheng-Bradley, X., Abecasis, G.R., McCarroll, S.A., Flicek, P., Underhill, P.A., Coin, L., Zerbino, D.R., Yang, F., Lee, C., Clarke, L., Auton, A., Erlich, Y., Handsaker, R.E., The 1000 Genomes Project Consortium, Bustamante, C.D., Tyler-Smith, C., 2016. Punctuated bursts in human male demography inferred from 1,244 worldwide Y-chromosome sequences. *Nature Genet* 48, 593–599.
- Slater, G.J., Harmon, L.J., Joyce, P., Revell, L.J., Alfaro, M.E., 2012. Fitting models of continuous trait evolution to incompletely sampled comparative data using Approximate Bayesian Computation. *Evolution* 66, 752–762.
- Stephens, M., Donnelly, P., 2000. Inference in molecular population genetics. *J. Roy. Statist. Soc. B* 62, 605–655.
- Tajima, F., 1983. Evolutionary relationship of DNA sequences in finite populations. *Genetics* 105, 437–460.
- Tavaré, S., 1984. Line-of-descent and genealogical processes, and their application in population genetics models. *Theor. Popul. Biol.* 26, 119–164.
- Tavaré, S., 2004. Ancestral Inference in Population Genetics. In: *Lectures on Probability Theory and Statistics*, vol. 1837, Springer, Berlin Heidelberg, pp. 1–188.
- Tavaré, S., Balding, D., Griffiths, R.C., Donnelly, P., 1997. Inferring coalescence times from DNA sequence data. *Genetics* 145, 505–518.
- Watterson, G.A., 1974. The sampling theory of selectively neutral alleles. *Adv. Appl. Probab.* 6, 463–488.
- Watterson, G.A., 1975. On the number of segregating sites in genetical models without recombination. *Theor. Popul. Biol.* 7, 256–276.
- Ewens, W.J., 1972. The sampling theory of selectively neutral alleles. *Theor. Popul. Biol.* 3, 87–112.