

Efficient Maximum-Likelihood Inference For The Isolation-With-Initial-Migration Model With Potentially Asymmetric Gene Flow

Rui J. Costa and Hilde Wilkinson-Herbots

Department of Statistical Science, University College London

Gower Street, London WC1E 6BT, UK

email: rui.costa.11@ucl.ac.uk

Abstract

The isolation-with-migration (IM) model is commonly used to make inferences about gene flow during speciation, using polymorphism data. However, Becquet and Przeworski (2009) report that the parameter estimates obtained by fitting the IM model are very sensitive to the model's assumptions – including the assumption of constant gene flow until the present. This paper is concerned with the isolation-with-initial-migration (IIM) model of Wilkinson-Herbots (2012), which drops precisely this assumption. In the IIM model, one ancestral population divides into two descendant subpopulations, between which there is an initial period of gene flow and a subsequent period of isolation. We derive a very fast method of fitting an extended version of the IIM model, which also allows for asymmetric gene flow and unequal population sizes. This is a maximum-likelihood method, applicable to data on the number of segregating sites between pairs of DNA sequences from a large number of independent loci. In addition to obtaining parameter estimates, our method can also be used to distinguish between alternative models representing different evolutionary scenarios, by means of likelihood ratio tests. We illustrate the procedure on pairs of *Drosophila* sequences from approximately 30,000 loci. The computing time needed to fit the most complex version of the model to this data set is only a couple of minutes. The R code to fit the IIM model can be found in the supplementary files of this paper.

Keywords: speciation, coalescent, maximum-likelihood, gene flow, isolation

1 Introduction

The 2-deme isolation-with-migration (IM) model is a population genetic model in which, at some point in the past, an ancestral population divided into two subpopulations. After the division, these subpopulations exchanged migrants at a constant rate until the present. The IM model has become one of the most popular probabilistic models in use to study genetic diversity under gene flow and population structure. Although applicable to populations within species, many researchers are using it to detect gene flow between diverging populations and to investigate the role of gene flow in the process of speciation. A meta-analysis of published research articles that used the IM model in the context of speciation can be found in Pinho and Hey (2010).

Several authors have developed computational methods to fit IM models to real DNA data. Some of the most used programs are aimed at data sets consisting of a large number of sequences from a small number of loci. This is the case of *MDIV* (Nielsen and Wakeley, 2001), *IM* (Hey and Nielsen, 2004; Hey, 2005), *IMa* (Hey and Nielsen, 2007) and *IMa2* (Hey, 2010), which rely on Bayesian MCMC methods to estimate the model parameters and are computationally very intensive.

In the past decade, the availability of whole genome sequences has increased significantly. Very large data sets consisting of thousands of loci, usually from just a few individuals, became available for analysis, and such multilocus data sets are more informative than their single locus counterparts. In fact, as the sample size for a single locus increases, the probability that an extra sequence adds a deep (i.e. informative) branch to the coalescent tree quickly becomes negligible (see, e.g., Hein et al., 2005, p. 28-29). This new type of data set is also much more suitable for likelihood inference: if at each locus the observation consists only of a pair or a triplet of sequences, the coalescent process of these sequences is relatively simple and can more easily be used to derive the likelihood for the locus concerned. Furthermore, if the loci studied are distant from each other, observations at different loci can be considered independent and thus the likelihood of a whole set will consist of the product of the likelihoods for the individual loci.

For these reasons, some of the most recent theoretical papers and computer implementations of the IM model are aimed at data sets consisting of a small number of sequences at a large number of loci. Wang and Hey (2010) and Zhu and Yang (2012), for example, developed maximum-likelihood methods based on numerical integration. Computationally less intensive methods of fitting the IM model were studied by Lohse et al. (2011), who used moment generating functions, and Andersen et al. (2014), who

resorted to matrix exponentiation and decomposition.

Despite all the progress, several authors have recently pointed out some limitations of the IM model. Becquet and Przeworski (2009) report that the parameter estimates obtained by fitting the IM model are highly sensitive to the model's assumptions – including the assumption of constant gene flow until the present. Strasburg and Rieseberg (2011) and Sousa et al. (2011) note that the gene flow timing estimates reported in the literature, based on the IM model and obtained with the *IMA2* program, have extremely wide confidence intervals, as the time of gene flow is non-identifiable in the IM model.

As a step to overcome these limitations, Wilkinson-Herbots (2012) studied an extension of the IM model, the isolation-with-initial-migration (IIM) model, which is more realistic than the IM model in the context of speciation. Broadly speaking, the IIM model is an IM model in which gene flow ceased at some point in the past. Explicit formulae for the distribution of the coalescence time of a pair of sequences, and the distribution of the number of nucleotide differences between them, are derived in Wilkinson-Herbots (2012). These analytic results enable a very fast computation of the likelihood of a data set consisting of observations on pairs of sequences at a large number of independent loci. However, for mathematical reasons this work was limited to the case of symmetric migration and equal subpopulation sizes during the migration period.

In this paper, we study a more general IIM model which allows for asymmetric gene flow during the migration period. It also allows for unequal subpopulation sizes during gene flow, as well as during the isolation stage. Both this model and other simpler models studied in this paper assume haploid DNA sequences, which accumulate mutations according to the infinite sites assumption (Watterson, 1975). An extension to the Jukes-Cantor model of mutation is feasible but beyond the scope of this paper.

We first describe, for different versions of the IIM model, an efficient method to compute the likelihood of a set of observations on the number of different nucleotides between pairs of sequences (also termed the number of pairwise differences). Each pair of sequences comes from a different locus and we assume free recombination between loci and no recombination within loci. Secondly, we illustrate how to use this method to fit the IIM model to real data. The data set of *Drosophila* sequences from Wang and Hey (2010), containing over 30,000 observations (i.e. loci), is used for this purpose. Finally, we demonstrate, using this data set, how different models, representing different evolutionary scenarios, can be compared using likelihood ratio tests.

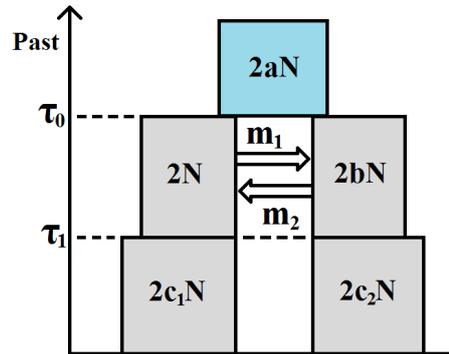


Figure 1: The IIM model. The left-hand side subpopulation is subpopulation 1; the right-hand side subpopulation is subpopulation 2.

2 Theory and methods

For the purposes of the present paper, and from a forward-in-time perspective, the isolation-with-migration (IM) model makes the following assumptions: a) until time τ_0 ago ($\tau_0 > 0$), a population of DNA sequences from a single locus followed a Wright-Fisher haploid model (Fisher, 1930; Wright, 1931); b) at time τ_0 ago, this ancestral population split into two Wright-Fisher subpopulations with constant gene flow between them. If we take an IM model and add the assumption that, at time τ_1 ago ($0 < \tau_1 < \tau_0$), gene flow ceased, we get an isolation-with-initial-migration (IIM) model. Figure 1 illustrates the fullest IIM model dealt with in this paper.

In the IIM model of Figure 1, the population sizes are given inside the boxes, in units of DNA sequences. All population sizes are assumed constant and strictly positive. The parameters a , b , c_1 and c_2 indicate the relative size of each population with respect to subpopulation 1 during the migration stage. For example, if $2N_{anc}$ is the number of sequences in the ancestral population, then $a = 2N_{anc}/2N$. Between times τ_0 and τ_1 ago (two time parameters in units of $2N$ generations) there is gene flow between the subpopulations: in each generation, a fraction m_i of subpopulation i are immigrants from subpopulation j ($i, j \in \{1, 2\}$ with $i \neq j$), i.e. m_i is the migration rate per generation from subpopulation i to subpopulation j backward in time. Within each subpopulation, reproduction follows the neutral Wright-Fisher model and, in each generation, restores the subpopulations to their original sizes, i.e. reproduction undoes any decrease or increase in size caused by gene flow (this assumption of constant population size is common and usually reflects restrictions on food and habitat size).

Under the IIM model, the genealogy of a sample of two DNA sequences from the present subpopulations can be described by successive Markov chains, working backward in time. We will define these in the simplest possible way, using the smallest state space necessary for the derivation of the coalescence time distribution. Hence, during the isolation stage (until time τ_1 into the past) and the migration stage (between τ_1 and τ_0), the process can only be in state 1 – both lineages in subpopulation 1 –, state 2 – both lineages in subpopulation 2 –, state 3 – one lineage in each subpopulation –, or state 4 – in which lineages have coalesced. After τ_0 , the lineages have either coalesced already – state 4 –, or have not – state 0. Only states 1, 2 and 3 can be initial states, according to whether we sample two sequences from subpopulation 1, two sequences from subpopulation 2, or one sequence from each subpopulation. When the genealogical process starts in state i (with $i \in \{1,2,3\}$), the time until the most recent common ancestor of the two sampled sequences is denoted $T^{(i)}$, whereas $S^{(i)}$ denotes the number of nucleotide differences between them.

If time is measured in units of $2N$ generations and N is large, the genealogical process is well approximated by a succession of three continuous-time Markov chains, one for each stage of the IIM model (Kingman, 1982a,b; Notohara, 1990). We refer to this stochastic process in continuous time as the *coalescent* under the IIM model. During the isolation stage, the approximation is by a Markov chain defined by the generator matrix

$$\mathbf{Q}_{\text{iso}}^{(i)} = \begin{matrix} & \begin{matrix} (i) & (4) \end{matrix} \\ \begin{matrix} (i) \\ (4) \end{matrix} & \begin{bmatrix} -\frac{1}{c_i} & \frac{1}{c_i} \\ 0 & 0 \end{bmatrix} \end{matrix}, \quad (1)$$

with $i \in \{1,2\}$ being the initial state (Kingman, 1982a,b). If 3 is the initial state, the lineages cannot coalesce before τ_1 . During the ancestral stage, the genealogical process is approximated by a Markov chain with generator matrix

$$\mathbf{Q}_{\text{anc}} = \begin{matrix} & \begin{matrix} (0) & (4) \end{matrix} \\ \begin{matrix} (0) \\ (4) \end{matrix} & \begin{bmatrix} -\frac{1}{a} & \frac{1}{a} \\ 0 & 0 \end{bmatrix} \end{matrix}. \quad (2)$$

(Kingman, 1982a,b). In between, during the migration stage, the approximation is by a

Markov chain with generator matrix

$$\mathbf{Q}_{\text{mig}} = \begin{matrix} & \begin{matrix} (1) & (3) & (2) & (4) \end{matrix} \\ \begin{matrix} (1) \\ (3) \\ (2) \\ (4) \end{matrix} & \begin{bmatrix} -(1 + M_1) & M_1 & 0 & 1 \\ \frac{M_2}{2} & -\left(\frac{M_1 + M_2}{2}\right) & \frac{M_1}{2} & 0 \\ 0 & M_2 & -(1/b + M_2) & 1/b \\ 0 & 0 & 0 & 0 \end{bmatrix} \end{matrix} \quad (3)$$

(Notohara, 1990). In this matrix, $M_i/2 = 2Nm_i$ represents the rate of migration (in continuous time) of a single sequence when in subpopulation i . The rates of coalescence for two lineages in subpopulation 1 or 2 are 1 and $1/b$ respectively. Note that state 3 corresponds to the second row and column, and state 2 to the third row and column. This swap was dictated by mathematical convenience: the matrix \mathbf{Q}_{mig} should be as symmetric as possible because this facilitates a proof in the next section.

2.1 Distribution of the time until coalescence under bidirectional gene flow ($M_1 > 0$, $M_2 > 0$)

To find $f_T^{(i)}$, the density of the coalescence time $T^{(i)}$ of two lineages under the IIM model, given that the process starts in state i and there is gene flow in both directions, we consider separately the three Markov chains mentioned above. We let $T_{iso}^{(i)}$ ($i \in \{1, 2\}$), $T_{mig}^{(i)}$ ($i \in \{1, 2, 3\}$) and $T_{anc}^{(0)}$ denote the times until absorption of the time-homogeneous Markov chains defined by the generator matrices $\mathbf{Q}_{iso}^{(i)}$, \mathbf{Q}_{mig} and \mathbf{Q}_{anc} respectively. And we let the corresponding *pdf*'s (or *cdf*'s) be denoted by $f_{iso}^{(i)}$, $f_{mig}^{(i)}$ and $f_{anc}^{(0)}$ (or $F_{iso}^{(i)}$, $F_{mig}^{(i)}$ and $F_{anc}^{(0)}$). Then $f_T^{(i)}$ can be expressed in terms of the distribution functions just mentioned:

$$f_T^{(i)}(t) = \begin{cases} f_{iso}^{(i)}(t) & \text{for } 0 \leq t \leq \tau_1, \\ \left[1 - F_{iso}^{(i)}(\tau_1)\right] f_{mig}^{(i)}(t - \tau_1) & \text{for } \tau_1 < t \leq \tau_0, \\ \left[1 - F_{iso}^{(i)}(\tau_1)\right] \left[1 - F_{mig}^{(i)}(\tau_0 - \tau_1)\right] f_{anc}^{(0)}(t - \tau_0) & \text{for } \tau_0 < t < \infty, \\ 0 & \text{otherwise.} \end{cases} \quad (4)$$

for $i \in \{1, 2\}$. If 3 is the initial state,

$$f_T^{(3)}(t) = \begin{cases} f_{mig}^{(3)}(t - \tau_1) & \text{for } \tau_1 < t \leq \tau_0, \\ \left[1 - F_{mig}^{(3)}(\tau_0 - \tau_1)\right] f_{anc}^{(0)}(t - \tau_0) & \text{for } \tau_0 < t < \infty, \\ 0 & \text{otherwise.} \end{cases} \quad (5)$$

The important conclusion to draw from these considerations is that to find the distribution of the coalescence time under the IIM model, we only need to find the distributions of the absorption times under the simpler processes just defined.

A Markov process defined by the matrix \mathbf{Q}_{anc} , and starting in state 0, is simply Kingman's coalescent (Kingman, 1982a,b). For such a process, the distribution of the coalescence time is exponential, with rate equal to the inverse of the relative population size:

$$f_{anc}^{(0)}(t) = \frac{1}{a} e^{-\frac{1}{a}t}, \quad 0 \leq t < \infty. \quad (6)$$

A Markov process defined by $\mathbf{Q}_{iso}^{(i)}$, $i \in \{1, 2\}$, is again Kingman's coalescent, so

$$f_{iso}^{(i)}(t) = \frac{1}{c_i} e^{-\frac{1}{c_i}t}, \quad 0 \leq t < \infty. \quad (7)$$

Finally, with respect to the 'structured' coalescent process defined by the matrix \mathbf{Q}_{mig} , we prove in the Appendix that, for $i \in \{1, 2, 3\}$,

$$f_{mig}^{(i)}(t) = - \sum_{j=1}^3 V_{ij}^{-1} V_{j4} \lambda_j e^{-\lambda_j t}, \quad (8)$$

where V_{ij} is the (i, j) entry of the (non-singular) matrix \mathbf{V} , whose rows are the left eigenvectors of \mathbf{Q}_{mig} . The (i, j) entry of the matrix \mathbf{V}^{-1} is denoted by V_{ij}^{-1} . The λ_j ($j \in \{1, 2, 3\}$) are the absolute values of those eigenvalues of \mathbf{Q}_{mig} which are strictly negative (the remaining one is zero). Since the λ_j are real and strictly positive, the density function of $T_{mig}^{(i)}$ is a linear combination of exponential densities.

Substituting the *pdf*'s from equations (6), (7) and (8) into the equations (4) and (5),

and denoting by \mathbf{A} the three by three matrix with entries $A_{ij} = -V_{ij}^{-1}V_{j4}$, we obtain

$$f_T^{(i)}(t) = \begin{cases} \frac{1}{c_i} e^{-\frac{1}{c_i}t} & \text{for } 0 \leq t \leq \tau_1, \\ e^{-\frac{1}{c_i}\tau_1} \sum_{j=1}^3 A_{ij} \lambda_j e^{-\lambda_j(t-\tau_1)} & \text{for } \tau_1 < t \leq \tau_0, \\ e^{-\frac{1}{c_i}\tau_1} \sum_{j=1}^3 A_{ij} e^{-\lambda_j(\tau_0-\tau_1)} \frac{1}{a} e^{-\frac{1}{a}(t-\tau_0)} & \text{for } \tau_0 < t < \infty, \\ 0 & \text{otherwise,} \end{cases} \quad (9)$$

for $i \in \{1, 2\}$, and

$$f_T^{(3)}(t) = \begin{cases} \sum_{j=1}^3 A_{3j} \lambda_j e^{-\lambda_j(t-\tau_1)} & \text{for } \tau_1 < t \leq \tau_0, \\ \sum_{j=1}^3 A_{3j} e^{-\lambda_j(\tau_0-\tau_1)} \frac{1}{a} e^{-\frac{1}{a}(t-\tau_0)} & \text{for } \tau_0 < t < \infty, \\ 0 & \text{otherwise.} \end{cases} \quad (10)$$

If $M_1 = M_2$ and $b = 1$ (i.e. in the case of symmetric gene flow and equal subpopulation sizes during the gene flow period), results (9) and (10) above simplify to the corresponding results in Wilkinson-Herbots (2012) – in this case, the coefficient A_{i3} in the linear combination is zero for $i \in \{1, 2, 3\}$.

2.2 Distribution of the time until coalescence under unidirectional gene flow, and in the absence of gene flow

If either M_1 or M_2 is equal to zero, or if both are equal to zero, the above derivation of $f_{mig}^{(i)}$ is no longer applicable, as the similarity transformation in part (ii) of the proof is no longer defined (see the denominators in some entries of the matrix \mathbf{D}). In this section, we derive $f_{mig}^{(i)}$, the density of the absorption time of the Markov chain defined by the matrix \mathbf{Q}_{mig} given in equation (3), starting from state i , when one or both the migration rates are zero. Again, this is all we need to fill in equations (4) and (5) and obtain the distribution of the coalescence time of a pair of DNA sequences under the IIM model. Having gene flow in just one direction considerably simplifies the coalescent. For this reason, we resort to moment generating functions, instead of eigendecomposition, and derive fully explicit *pdf*'s.

2.2.1 Migration from subpopulation 2 to subpopulation 1 backward in time

$$(M_1 = 0, M_2 > 0)$$

Let $T_{mig}^{(i)}$ again be defined as the absorption time of the Markov chain generated by \mathbf{Q}_{mig} , now with $M_1 = 0$ and $M_2 > 0$, given that the initial state is $i \in \{1, 2, 3\}$. A first-step argument gives the following system of equations for the *mgf* of $T_{mig}^{(i)}$:

$$\begin{aligned} \mathbb{E} \left[\exp \left(-sT_{mig}^{(1)} \right) \right] &= \left(\frac{1}{1+s} \right) \\ \mathbb{E} \left[\exp \left(-sT_{mig}^{(2)} \right) \right] &= \left(\frac{M_2}{1/b+M_2+s} \right) \mathbb{E} \left[\exp \left(-sT_{mig}^{(3)} \right) \right] + \left(\frac{1/b}{1/b+M_2+s} \right) \\ \mathbb{E} \left[\exp \left(-sT_{mig}^{(3)} \right) \right] &= \left(\frac{M_2}{M_2+2s} \right) \mathbb{E} \left[\exp \left(-sT_{mig}^{(1)} \right) \right] . \end{aligned}$$

Solving this system of equations and applying a partial fraction decomposition, the distributions of $T_{mig}^{(1)}$, $T_{mig}^{(2)}$ and $T_{mig}^{(3)}$ can be expressed as linear combinations of exponential distributions:

$$\begin{aligned} \mathbb{E} \left[\exp \left(-sT_{mig}^{(1)} \right) \right] &= \left(\frac{1}{1+s} \right) \\ \mathbb{E} \left[\exp \left(-sT_{mig}^{(2)} \right) \right] &= \left(\frac{M_2}{1/b+M_2+s} \right) \left(\frac{M_2}{M_2+2s} \right) \left(\frac{1}{1+s} \right) + \left(\frac{1/b}{1/b+M_2+s} \right) \\ &= \left(\frac{bM_2^2}{(M_2-2)(1-b+bM_2)} \right) \left(\frac{1}{1+s} \right) + \left(\frac{4bM_2}{(2-M_2)(2+bM_2)} \right) \left(\frac{M_2}{M_2+2s} \right) \\ &\quad + \left(\frac{1/b}{1/b+M_2} + \frac{b^2M_2^2}{(2+bM_2)(1-b+bM_2)(1/b+M_2)} \right) \left(\frac{1/b+M_2}{1/b+M_2+s} \right) \\ \mathbb{E} \left[\exp \left(-sT_{mig}^{(3)} \right) \right] &= \left(\frac{M_2}{M_2+2s} \right) \left(\frac{1}{1+s} \right) \\ &= \left(\frac{M_2}{M_2-2} \right) \left(\frac{1}{1+s} \right) + \left(\frac{2}{2-M_2} \right) \left(\frac{M_2}{M_2+2s} \right) . \end{aligned}$$

Thus we obtain the following *pdf*'s:

$$\begin{aligned} f_{mig}^{(1)}(t) &= e^{-t} \\ f_{mig}^{(2)}(t) &= \left(\frac{bM_2^2}{(M_2-2)(1-b+bM_2)} \right) e^{-t} + \left(\frac{4bM_2}{(2-M_2)(2+bM_2)} \right) \frac{M_2}{2} e^{-\frac{M_2}{2}t} \\ &\quad + \left(\frac{1}{1+bM_2} + \frac{b^2M_2^2}{(2+bM_2)(1-b+bM_2)(1/b+M_2)} \right) \left(\frac{1}{b} + M_2 \right) e^{-(1/b+M_2)t} \\ f_{mig}^{(3)}(t) &= \left(\frac{M_2}{M_2-2} \right) e^{-t} + \left(\frac{2}{2-M_2} \right) \frac{M_2}{2} e^{-\frac{M_2}{2}t} \end{aligned}$$

for $t > 0$.

The *pdf* of the coalescence time of a pair of DNA sequences under an IIM model with $M_1 = 0$ and $M_2 > 0$ can now be easily derived by comparing the above expressions with equation (8): $f_T^{(i)}(t)$ is given by equations (9) and (10) above, but now with

$$\boldsymbol{\lambda} = \begin{bmatrix} 1 & \frac{M_2}{2} & \frac{1}{b} + M_2 \end{bmatrix} ,$$

and

$$\mathbf{A} = \begin{bmatrix} 1 & 0 & 0 \\ \frac{bM_2^2}{(M_2-2)(1-b+bM_2)} & \frac{4bM_2}{(2-M_2)(2+bM_2)} & \frac{1}{1+bM_2} + \frac{b^2M_2^2}{(2+bM_2)(1-b+bM_2)(1/b+M_2)} \\ \frac{M_2}{M_2-2} & \frac{2}{2-M_2} & 0 \end{bmatrix} .$$

2.2.2 Migration from subpopulation 1 to subpopulation 2 backward in time ($M_1 > 0, M_2 = 0$)

In the opposite case of unidirectional migration, and using the same derivation procedure, we find that:

$$f_{mig}^{(1)}(t) = \left(\frac{b^2M_1^2}{(bM_1-2)(b-1+bM_1)} \right) \frac{1}{b} e^{-\frac{1}{b}t} + \left(\frac{4M_1}{(2-bM_1)(2+M_1)} \right) \frac{M_1}{2} e^{-\frac{M_1}{2}t} + \left(\frac{1}{(1+M_1)} + \frac{M_1^2}{(2+M_1)(b-1+bM_1)(1+M_1)} \right) (1+M_1) e^{-(1+M_1)t}$$

$$f_{mig}^{(2)}(t) = \frac{1}{b} e^{-\frac{1}{b}t}$$

$$f_{mig}^{(3)}(t) = \left(\frac{bM_1}{bM_1-2} \right) \frac{1}{b} e^{-\frac{1}{b}t} + \left(\frac{2}{2-bM_1} \right) \frac{M_1}{2} e^{-\frac{M_1}{2}t} .$$

As a result, the *pdf* of the coalescence time of a pair of sequences under the IIM model, $f_T^{(i)}(t)$, is again given by equations (9) and (10), now with

$$\boldsymbol{\lambda} = \begin{bmatrix} \frac{1}{b} & \frac{M_1}{2} & 1 + M_1 \end{bmatrix}$$

and

$$\mathbf{A} = \begin{bmatrix} \frac{b^2 M_1^2}{(bM_1-2)(b-1+bM_1)} & \frac{4M_1}{(2-bM_1)(2+M_1)} & \frac{1}{1+M_1} + \frac{M_1^2}{(2+M_1)(b-1+bM_1)(1+M_1)} \\ 1 & 0 & 0 \\ \frac{bM_1}{bM_1-2} & \frac{2}{2-bM_1} & 0 \end{bmatrix}.$$

2.2.3 Distribution of the time until coalescence under an IIM model with $M_1 = M_2 = 0$

In this case, the IIM model reduces to a complete isolation model where both descendant populations may change size at time τ_1 into the past. The distribution of the absorption time $T_{mig}^{(i)}$ corresponding to \mathbf{Q}_{mig} will now be either exponential, if both sampled sequences are from the same subpopulation (i.e. for $i \in \{1, 2\}$), or coalescence will not be possible at all until the ancestral population is reached, if we take a sequence from each subpopulation (i.e. if $i = 3$). It follows that the *pdf* of the coalescence time of a pair of sequences in the IIM model is given by equations (9) and (10) with

$$\boldsymbol{\lambda} = \begin{bmatrix} 1 & \frac{1}{b} & 0 \end{bmatrix}$$

and

$$\mathbf{A} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}.$$

2.3 The distribution of the number S of segregating sites

Let $S^{(i)}$ denote the number of segregating sites in a random sample of two sequences from a given locus, when the ancestral process of these sequences follows the coalescent under the IIM model and the initial state is state i ($i \in \{1, 2, 3\}$). Recall the infinite sites assumption and assume that the distribution of the number of mutations hitting one sequence in a single generation is Poisson with mean μ . As before, time is measured in units of $2N$ generations and we use the coalescent approximation. Given the coalescence time $T^{(i)}$ of two sequences, $S^{(i)}$ is Poisson distributed with mean $\theta T^{(i)}$, where $\theta = 4N\mu$ denotes the scaled mutation rate. Since the *pdf* of $T^{(i)}$, $f_T^{(i)}$, is known, the likelihood $L^{(i)}$ of an observation from a single locus corresponding to the initial state i can be derived

by integrating out $T^{(i)}$:

$$L^{(i)}(\gamma, \theta; s) = P(S^{(i)} = s; \gamma, \theta) = \int_0^{\infty} P(S^{(i)} = s | T^{(i)} = t) f_T^{(i)}(t) dt,$$

where γ is the vector of parameters of the coalescent under the IIM model, that is, $\gamma = (a, b, c_1, c_2, \tau_1, \tau_0, M_1, M_2)$. There is no need to compute this integral numerically: because $f_T^{(i)}$ has been expressed in terms of a piecewise linear combination of exponential or shifted exponential densities, we can use standard results for a Poisson process superimposed onto an exponential or shifted exponential distribution.

The equations (18) and (29) of Wilkinson-Herbots (2012) use this superimposition of processes to derive the distribution of S under a mathematically much simpler IIM model with symmetric migration and equal subpopulation sizes during the period of migration. These equations can now be adapted to obtain the *pmf* of S under each of the migration scenarios dealt with in this paper. The changes accommodate the fact that the density of the coalescence time during the migration stage of the model is now given by a different linear combination of exponential densities, where the coefficients in the linear combination, as well as the parameters of the exponential densities, are no longer the same. The *pmf* of S has the following general form:

$$\begin{aligned} P(S^{(i)} = s) &= \frac{(c_i \theta)^s}{(1 + c_i \theta)^{s+1}} \left(1 - e^{-\tau_1 (\frac{1}{c_i} + \theta)} \sum_{l=0}^s \frac{(\frac{1}{c_i} + \theta)^l \tau_1^l}{l!} \right) \\ &+ e^{-\frac{1}{c_i} \tau_1} \sum_{j=1}^3 A_{ij} \frac{\lambda_j \theta^s}{(\lambda_j + \theta)^{s+1}} \left(e^{-\theta \tau_1} \sum_{l=0}^s \frac{(\lambda_j + \theta)^l \tau_1^l}{l!} \right. \\ &\quad \left. - e^{-\lambda_j (\tau_0 - \tau_1) - \theta \tau_0} \sum_{l=0}^s \frac{(\lambda_j + \theta)^l \tau_0^l}{l!} \right) \\ &+ \frac{e^{-\frac{1}{c_i} \tau_1 - \theta \tau_0} (a \theta)^s}{(1 + a \theta)^{s+1}} \left(\sum_{l=0}^s \frac{(\frac{1}{a} + \theta)^l \tau_0^l}{l!} \right) \sum_{j=1}^3 A_{ij} e^{-\lambda_j (\tau_0 - \tau_1)} \end{aligned} \quad (11)$$

for $i \in \{1, 2\}$ and

$$\begin{aligned}
 P(S^{(3)} = s) &= \sum_{j=1}^3 A_{3j} \frac{\lambda_j \theta^s}{(\lambda_j + \theta)^{s+1}} \left(e^{-\theta \tau_1} \sum_{l=0}^s \frac{(\lambda_j + \theta)^l \tau_1^l}{l!} \right. \\
 &\quad \left. - e^{-\lambda_j(\tau_0 - \tau_1) - \theta \tau_0} \sum_{l=0}^s \frac{(\lambda_j + \theta)^l \tau_0^l}{l!} \right) \\
 &\quad + \frac{e^{-\theta \tau_0} (a\theta)^s}{(1+a\theta)^{s+1}} \left(\sum_{l=0}^s \frac{(\frac{1}{a} + \theta)^l \tau_0^l}{l!} \right) \sum_{j=1}^3 A_{3j} e^{-\lambda_j(\tau_0 - \tau_1)}
 \end{aligned} \tag{12}$$

for $s \in \{0, 1, 2, 3, \dots\}$. As defined in Section 2.1, under bidirectional migration $\boldsymbol{\lambda} = (\lambda_1, \lambda_2, \lambda_3)$ is the vector of the absolute values of the strictly negative eigenvalues of \mathbf{Q}_{mig} and $A_{ij} = -V_{ij}^{-1}V_{j4}$. If migration occurs in one direction only, with $M_1 = 0$ and $M_2 > 0$, the matrix \mathbf{A} and the vector $\boldsymbol{\lambda}$ are those given in Section 2.2.1. In the remaining cases, when $M_1 > 0$ and $M_2 = 0$ or when there is no gene flow, \mathbf{A} and $\boldsymbol{\lambda}$ are given in Sections 2.2.2 and 2.2.3 respectively. In the special case of $M_1 = M_2$ and $b = 1$, equations (11) and (12) reduce to the results of Wilkinson-Herbots (2012).

2.4 The likelihood of a multilocus data set

Recall that, for our purposes, an observation consists of the number of nucleotide differences between a pair of DNA sequences from the same locus. To jointly estimate all the parameters of the IIM model, our method requires a large set of observations on each of the three initial states (i.e. on pairs of sequences from subpopulation 1, from subpopulation 2, and from both subpopulations). To compute the likelihood of such a data set, we use the assumption that observations are independent, so we should have no more than one observation or pair of sequences per locus and there should be free recombination between loci, i.e. loci should be sufficiently far apart.

Let each locus for the initial state i be assigned a label $j_i \in \{1_i, 2_i, 3_i, \dots, J_i\}$, where J_i is the total number of loci associated with initial state i . Denote by $\theta_{j_i} = 4N\mu_{j_i}$ the scaled mutation rate at locus j_i , where μ_{j_i} is the mutation rate per sequence per generation at that locus. Let θ denote the average scaled mutation rate over all loci and denote by $r_{j_i} = \frac{\theta_{j_i}}{\theta}$ the relative mutation rate of locus j_i . Then $\theta_{j_i} = r_{j_i}\theta$. If the relative mutation rates are known, we can represent the likelihood of the observation at locus j_i simply by $L(\boldsymbol{\gamma}, \theta; s_{j_i})$. By independence, the likelihood of the data set is then given by

$$L(\boldsymbol{\gamma}, \theta; \mathbf{s}) = \prod_{i=1}^3 \prod_{j_i=1}^{J_i} L(\boldsymbol{\gamma}, \theta; s_{j_i}) \quad . \tag{13}$$

In our likelihood method, the r_{j_i} are treated as known constants. In practice, however, the relative mutation rates at the different loci are usually estimated using outgroup sequences (Yang, 2002; Wang and Hey, 2010).

3 Results

3.1 Simulated data

We generated two batches of data sets by simulation, each batch having one hundred data sets. Each data set consists of thousands of independent observations, where each observation represents the number of nucleotide differences between two DNA sequences belonging to the same locus, when the genealogy of these sequences follows an IIM model. In batch 1 each data set has 40,000 observations: 10,000 observations for initial state 1 (two sequences drawn from subpopulation 1), 10,000 for initial state 2 (two sequences drawn from subpopulation 2), and 20,000 for initial state 3 (one sequence from each of the two subpopulations). In batch 2 each data set has 800,000 observations: 200,000 for initial state 1, 200,000 for initial state 2, and 400,000 for initial state 3.

The data sets shown in this section were generated using the following parameter values: $a = 0.75$, $\theta = 2$, $b = 1.25$, $c_1 = 1.5$, $c_2 = 2$, $\tau_0 = 2$ and $\tau_1 = 1$, $M_1 = 0.5$ and $M_2 = 0.75$. Each observation in a data set refers to a different genetic locus j , and hence was generated using a different scaled mutation rate θ_j for that locus. For batch 1, we first fixed the average mutation rate over all sites to be $\theta = 2$. Then, a vector of 40,000 relative size scalars r_j was randomly generated using a Gamma(15,15) distribution. The scaled mutation rate at locus j was then defined to be $\theta_j = r_j\theta$, where r_j denotes the relative mutation rate at locus j , that is, the relative size of θ_j with respect to the average mutation rate, θ . All data sets in batch 1 were generated using the same vector of relative mutation rates. The generation of the 800,000 mutation rates θ_j used in batch 2 was carried out following the same procedure.

When fitting the IIM model to data sets generated in this manner, the relative mutation rates r_j are included as known constants in the log-likelihood function to be maximised. So, as far as mutation rates are concerned, only the average over all loci is estimated (i.e. the parameter θ). To increase the robustness and performance of the fitting procedure (see also Wilkinson-Herbots (2015) and the references therein), we found the maximum-likelihood estimates for a reparameterised model with parameters θ , $\theta_a = \theta a$, $\theta_b = \theta b$, $\theta_{c_1} = \theta c_1$, $\theta_{c_2} = \theta c_2$, $V = \theta(\tau_0 - \tau_1)$, $T_1 = \theta\tau_1$, M_1 and M_2 .

The boxplots of the maximum-likelihood estimates obtained for both batches of sim-

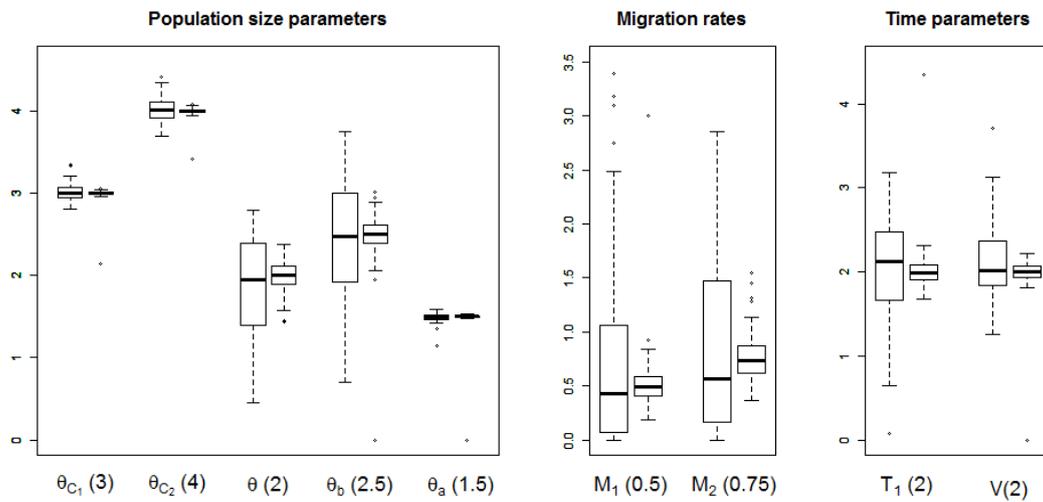


Figure 2: Estimates of population size parameters (left), migration rates (centre) and time parameters (right), for simulated data. For each parameter, the estimates shown in the left and right boxplots are based on sample sizes of 40,000 and 800,000 loci respectively. The values stated in parentheses are the true parameter values used to generate the data.

ulated data are shown in Figure 2. For each parameter, the boxplot on the left refers to batch 1 and the one on the right to batch 2. From the boxplots of time and population size parameters, it is seen that the estimates are centred around the true parameter values. Estimates for the migration rates are skewed to the right for batch 1, possibly because the true parameter values for these rates are closer to the boundary (zero) than the ones for population sizes and splitting times. For all types of parameters, increasing the sample size will decrease the variance of the maximum-likelihood estimator, as would be expected from using the correct expressions for the likelihood. In the case of the migration rate parameters, increasing the sample size eliminates most of the skewness.

Figure 3 shows a pair of normal Q-Q plots for each of three parameters: θ_{c_1} (a size parameter), T_1 (a time parameter) and M_1 (a migration parameter). The estimates in the top and bottom Q-Q plots are based on batch 1 and batch 2 respectively. It is clear from the figure that the distributions of the maximum-likelihood estimators of θ_{c_1} , T_1 and M_1 become increasingly Gaussian as we increase the number of observations (compare top and bottom Q-Q plots). This is also true for the estimators of the remaining parameters (results not shown here). In addition, we can also observe that the distributions of the estimators of time and population size parameters have already a reasonably Gaussian shape for a sample size of 40,000 observations (batch 1, top plots). Again, this is true

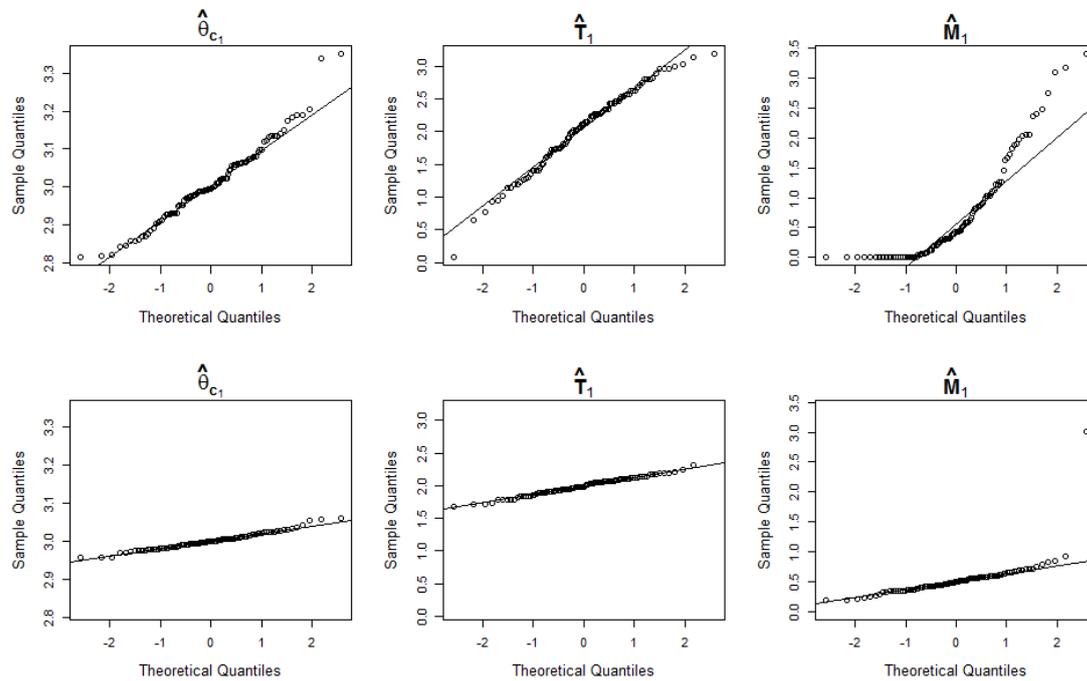


Figure 3: Normal Q-Q plots of maximum-likelihood estimates based on simulated data, for the parameters θ_{c_1} , T_1 and M_1 . The estimates shown in the top and bottom Q-Q plots are based on sample sizes of 40,000 and 800,000 loci respectively.

for the estimators of the remaining time and size parameters.

3.2 The data from Wang and Hey (2010)

3.2.1 Maximum-likelihood estimation

To illustrate our method, we apply it to a real, multilocus data set from two closely related species of *Drosophila*. The data set from Wang and Hey (2010) includes an alignment of sequences covering 30247 loci, which consists of two *D. simulans* assemblies, one *D. melanogaster* assembly and one *D. yakuba* assembly. It also includes a smaller alignment of sequences, spanning only 378 loci, which consists of two *D. melanogaster* assemblies and one *D. yakuba* assembly. We will call this last alignment the ‘Hutter subset’, because it was first studied by Hutter et al. (2007). The larger alignment will be termed the ‘Wang subset’. Our models are fitted to the *D. melanogaster* and *D. simulans* sequences from both subsets. The *D. yakuba* sequences are only used as outgroup sequences, to estimate the relative mutation rates at the different loci and to

Table 1: Results for the data of Wang and Hey (2010): maximum-likelihood estimates and values of the maximised log-likelihood, for the models shown in Figure 4.

Model	θ_a	θ	θ_b	θ_{c_1}	θ_{c_2}	T_1	V	M_1	M_2	$\log L(\phi)$
ISO	4.757	5.628	2.665	-	-	-	13.705	-	-	-90879.14
IM ₁	3.974	5.641	2.493	-	-	-	14.965	0.000	0.053	-90276.00
IIM ₁	3.191	5.581	2.589	-	-	6.931	9.928	0.000	0.528	-90069.44
IIM ₂	3.273	3.357	1.929	6.623	2.647	6.930	9.778	0.000	0.223	-89899.22
IIM ₃	3.273	3.357	1.929	6.623	2.647	6.930	9.778	-	0.223	-89899.22

calibrate time.

To estimate the relative mutation rates r_{j_i} , we use the *ad hoc* approach proposed by Yang (2002), which was also used in Wang and Hey (2010) and Lohse et al. (2011). Estimates are computed by means of the following method-of-moments estimator:

$$\hat{r}_{j_i} = \frac{J \bar{k}_{j_i}}{\sum_{m=1}^3 \sum_{n=1}^{J_m} \bar{k}_{n_m}}, \quad (14)$$

where J is the total number of loci, and \bar{k}_{j_i} is the average of the numbers of nucleotide differences observed in pairs of one ingroup sequence and one outgroup sequence, at locus j_i .

Since our method uses only one pair of sequences at each of a large number of independent loci, and requires observations for all initial states, the following procedure was adopted to select a suitable set of data. According to the genome assembly they stem from, sequences in the Wang subset were given one of three possible tags: ‘Dsim1’, ‘Dsim2’ or ‘Dmel’. To each of the 30247 loci in the Wang subset we assigned a letter: loci with positions 1, 4, 7,... in the genome alignment were assigned the letter A; loci with positions 2, 5, 8,... were assigned the letter B; and loci with positions 3, 6, 9,..., the letter C. A data set was then built by selecting observations corresponding to initial states 1 and 3 from the Wang subset (we used the Dsim1-Dsim2 sequences from loci A, the Dmel-Dsim1 sequences from loci B, and the Dmel-Dsim2 sequences from loci C), whilst observations corresponding to initial state 2 were obtained from the Hutter subset by comparing the two *D. melanogaster* sequences available at each locus.

Table 1 contains the maximum-likelihood estimates for the models shown in Figure 4. Note that the parameters of time and population size have been reparameterised as in section 3.1, and recall that M_1 and M_2 are the scaled migration rates backward in time. In the diagrams, the left and right subpopulations represent *D. simulans* and *D. melanogaster* respectively.

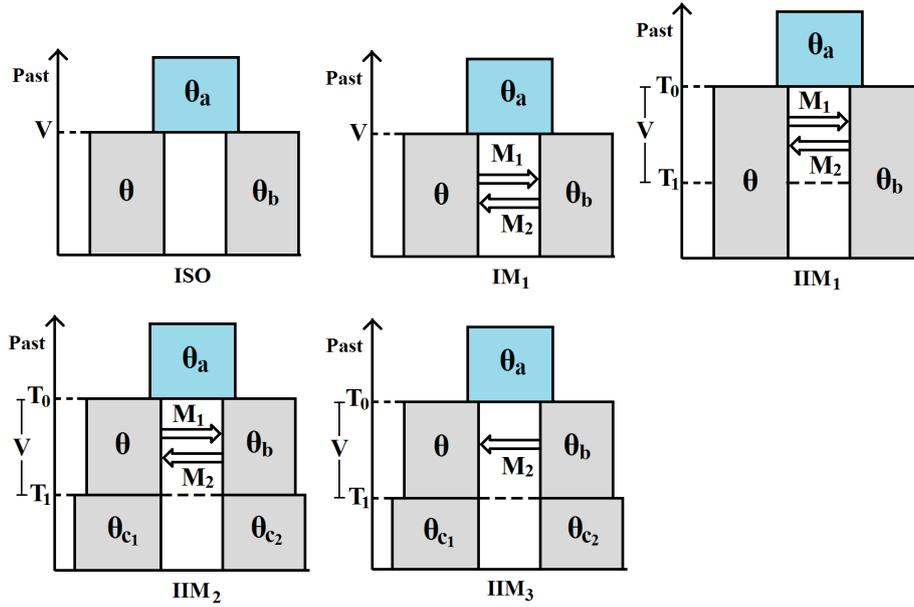


Figure 4: Models fitted to the data of Wang and Hey (2010): $\theta_a = \theta_a$, $\theta_b = \theta_b$, $\theta_{c_1} = \theta_{c_1}$, $\theta_{c_2} = \theta_{c_2}$, $V = T_0 - T_1 = \theta(\tau_0 - \tau_1)$ and $T_1 = \theta\tau_1$.

3.2.2 Model selection

In this section, we use likelihood ratio tests to determine which of the models listed in Table 1 fits the data of Wang and Hey (2010) best. For reasons which we now explain, the use of such tests in the present situation is not straightforward.

We wish to apply a standard large-sample theoretical result which states that, as the number of observations increases, the distribution of the likelihood ratio test statistic given by

$$D = -2 \log \lambda(\mathbf{s}) \quad ,$$

where

$$\lambda(\mathbf{s}) = \frac{\sup_{\phi \in \Phi_0} L(\phi; \mathbf{s})}{\sup_{\phi \in \Phi} L(\phi; \mathbf{s})} \quad (15)$$

approaches a χ^2 distribution. In equation (15), Φ_0 denotes the parameter space according to the null hypothesis (H_0). This space is a proper subspace of Φ , the parameter space according to the alternative hypothesis (H_1). The number of degrees of freedom of the limiting distribution is given by the difference between the dimensions of the two spaces.

A list of sufficient regularity conditions for this result can be found, for example,

in Casella and Berger (2001, p. 516). Two of them are clearly not met in the present problem. All models under study assume independent, but *not* identically distributed observations. In fact, the *pmf* of an observation changes according to which species the sequences are taken from (i.e. according to the initial state) and to which locus they belong to (because each locus has its own rate of mutation). In some cases, the pairwise comparison of models is also affected by another problem: every point of Φ_0 is a boundary point of Φ . In other words, if H_0 is true, the vector of true parameters $\phi^* \in \Phi_0$, whichever it might be, is on the boundary of Φ . This irregularity is present, for example, when $M_1 = M_2 = 0$ according to H_0 and $M_1, M_2 \in [0, \infty)$ according to H_1 .

Both these irregularities have been dealt with in the literature. As to the first one, even if observations are not identically distributed, both the maximum-likelihood estimator and the likelihood ratio test statistic still approach their respective limiting distributions if Lyapunov's condition is satisfied. This requires the very mild assumption that $\text{Var}(S_{j_i})$ is both bounded and bounded away from zero. It also requires $E|S_{j_i} - E(S_{j_i})|^3$ to be bounded (Pawitan, 2001, pp. 233-234). This last condition is satisfied if θ_{j_i} is bounded, which again is a rather mild assumption. The applicability of the central limit theorem is also supported by the simulations in section 3.1, which show that the distribution of the maximum-likelihood estimator indeed approaches normality as we increase the number of observations, so the lack of identically distributed observations should not compromise the applicability of large sample results in the procedures of model selection described in this paper.

The second irregularity, that is, the problem of having parameters on the boundary, has been the subject of papers such as Self and Liang (1987) and Kopylev and Sinha (2011). The limiting distribution of the likelihood ratio test statistic under this irregularity has been derived in these papers, but only for very specific cases. In most of these cases, the use of the naive χ_r^2 distribution, with r being the number of additional free parameters according to H_1 , turns out to be conservative, because the correct null distribution is a mixture of χ_ν^2 distributions with $\nu \leq r$. Our analysis of the data of Wang and Hey (2010) involves two likelihood ratio tests with parameters on the boundary (ISO vs. IM_1 , and IM_1 vs. IIM_1), so we need to check that the naive χ_r^2 distribution is also conservative in these cases. This was verified in a short simulation study which we now describe.

We generated 100 data sets from the ISO model, each one consisting of 40,000 observations, and fitted both the ISO model (H_0) and the IM_1 model (H_1) to obtain a sample of 100 realisations of the likelihood ratio test statistic. A Q-Q plot (Figure 5, left

boxplot) shows that the estimated quantiles of the null distribution are smaller than the corresponding theoretical quantiles of the χ^2 distribution with 2 degrees of freedom (the difference between the dimensions of Φ_0 and Φ in this particular case). In other words, the use of the naive χ^2 distribution is conservative in this case. Using χ^2_2 instead of the correct null distribution, at a significance level of 5%, the null hypothesis (i.e. the ISO model) was falsely rejected in only 1 out of the 100 simulations performed.

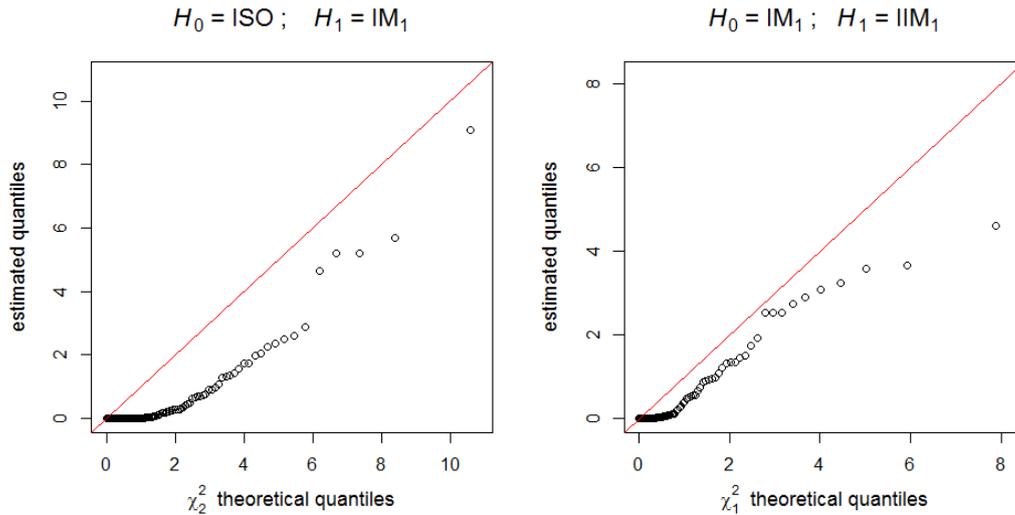


Figure 5: Q-Q plots of the estimated quantiles of the likelihood-ratio test statistic null distribution against the χ^2 distribution theoretical quantiles. Left plot: $H_0 = \text{ISO}$ model, $H_1 = \text{IM}_1$ model. Right plot: $H_0 = \text{IM}_1$ model, $H_1 = \text{IIM}_1$ model.

A similar simulation was carried out with respect to another pair of nested models: the IM_1 model (now as H_0), in which $\tau_1 = 0$, and the IIM_1 model (H_1), in which $\tau_1 \geq 0$. Again the naive χ^2 distribution (this time with only one degree of freedom) was found to be conservative (Figure 5, right boxplot). And once more, only in one out of the 100 simulations performed is the null hypothesis (the IM_1 model) falsely rejected at the 5% significance level, if χ^2_1 is used instead of the correct null distribution.

To select the model that best fitted the data of Wang and Hey (2010), we performed the sequence of pairwise comparisons shown in Table 2. For any sensible significance level, this sequence of comparisons leads to the choice of IIM_2 as the best fitting model. In fact, assuming the naive χ^2 as the null distribution, a significance level as low as 1.2×10^{-74} is enough to reject H_0 in each of the three tests. However, since $\hat{M}_1 = 0$ for this model (see Table 1), a final (backward) comparison is in order: that between IIM_2 and IIM_3 (which corresponds to fixing M_1 at zero in IIM_2). The nested model in

this comparison has one parameter less and, as can be seen in Table 1, has the same likelihood. So, in the end, we should prefer IIM₃ to IIM₂.

Table 2: Forward selection of the best model for the data of Wang and Hey (2010).

H₀	H₁	− log 2λ(S)	P-value
ISO	IM ₁	603.14	1.147E-262
IM ₁	IIM ₁	413.120	7.673E-92
IIM ₁	IIM ₂	340.440	1.187E-74

3.2.3 Confidence intervals for the selected model

The Wald confidence intervals are straightforward to calculate whenever the vector of estimates is neither on the boundary of the model’s parameter space, nor too close to it. In that case, it is reasonable to assume that the vector of *true* parameters does not lie on the boundary either. As a consequence, the vector of maximum-likelihood estimators is consistent and its distribution will approach a multivariate Gaussian distribution as the sample size grows (see, for example, Pawitan, 2001, p. 258). The confidence intervals can then be calculated using the inverted Hessian matrix.

In the case of the data of Wang and Hey (2010), the vector of estimates of the selected model (IIM₃) is an interior point of the parameter space. Assuming that the vector of true parameters is also away from the boundary, we computed the Wald 95% confidence intervals shown in Table 3 using the inverted Hessian. In agreement with our assumption, we note that none of the confidence intervals include zero.

For large sample sizes, and for true parameter values not too close to the boundary of the parameter space, the Wald intervals are both accurate and easy to compute. To check how well the Wald intervals for the IIM₃ model fare against the more accurate (see Pawitan, 2001, pp. 47-48), but also computationally more expensive, profile likelihood intervals, we included these in Table 3. The two methods yield very similar confidence intervals for all parameters except θ_b . The cause of this discrepancy should lie in the fact that we only had pairs of *D. melanogaster* sequences available from a few hundred loci (θ_b is the size of the *D. melanogaster* subpopulation during the migration stage).

Table 3: Results for the data of Wang and Hey (2010): point estimates and confidence intervals under the model IIM₃.

Parameter	Estimate	95% Confidence intervals	
		Wald	Profile likelihood
θ_a	3.273	(3.101, 3.445)	(3.100, 3.444)
θ	3.357	(3.139, 3.575)	(3.097, 3.578)
θ_b	1.929	(0.079, 3.779)	(0.672, 5.010)
θ_{c_1}	6.623	(6.407, 6.839)	(6.415, 6.843)
θ_{c_2}	2.647	(2.304, 2.990)	(2.331, 3.021)
T_1	6.930	(6.540, 7.320)	(6.542, 7.319)
V	9.778	(9.457, 10.099)	(9.456, 10.098)
M_2	0.223	(0.190, 0.256)	(0.186, 0.259)

3.2.4 Conversion of estimates

The conversion of the point estimates and confidence intervals to more conventional units is based on the estimates of Powell (1997) of the duration of one generation ($g = 0.1$ years) and the speciation time between *D. yakuba* and the common ancestor of *D. simulans* and *D. melanogaster* (10 million years); see also Wang and Hey (2010) and Lohse et al. (2011). Using these values, we estimated μ , the mutation rate per locus per generation, averaged over all loci, to be $\hat{\mu} = 2.31 \times 10^{-7}$.

In Tables 4, 5 and 6, we show the converted estimates for the best fitting model IIM₃. The effective population size estimates, in units of diploid individuals, are all based on estimators of the form $\hat{N} = \frac{1}{4\hat{\mu}} \times \hat{\theta}$, for example, the estimate of the ancestral population effective size N_a is given by $\frac{1}{4\hat{\mu}} \times \hat{\theta}_a$. The estimates in years of the time since the onset of speciation and of the time since the end of gene flow are given by $\hat{t}_0 = \frac{g}{2\hat{\mu}} \times (\hat{T}_1 + \hat{V})$ and $\hat{t}_1 = \frac{g}{2\hat{\mu}} \times \hat{T}_1$ respectively. With respect to gene flow, we use $\hat{q}_1 = \hat{\mu} \times \frac{\hat{M}_2\hat{b}}{\hat{\theta}}$ as the estimator of the *fraction* of subpopulation 1 that migrates to subpopulation 2 in each generation, forward in time, and $\hat{s}_1 = \frac{\hat{M}_2\hat{b}}{2}$ as the estimator of the *number* of migrant sequences from subpopulation 1 to subpopulation 2 in each generation, also forward in time.

If g and $\hat{\mu}$ are treated as constants, then each of the estimators just given can be expressed as a constant times a product – or a ratio – of the estimators of non-converted parameters. For example, we have that

$$\hat{q}_1 = \hat{\mu} \times \frac{\hat{M}_2\hat{b}}{\hat{\theta}} = \text{constant} \times \frac{\hat{M}_2\hat{b}}{\hat{\theta}} \quad ,$$

Table 4: Effective population size estimates for the data of Wang and Hey (2010) under the model IIM₃ (values in millions of diploid individuals).

Population	Population size	95% Confidence intervals	
		Wald	Profile likelihood
Ancestral population (N_a)	3.549	(3.362, 3.736)	(3.362, 3.735)
<i>D. simulans</i> , migration stage (N)	3.640	(3.404, 3.877)	(3.359, 3.880)
<i>D. melanogaster</i> , migration stage (N_b)	2.092	(0.085, 4.099)	(0.729, 5.433)
<i>D. simulans</i> , isolation stage (N_{c_1})	7.182	(6.949, 7.415)	(6.957, 7.421)
<i>D. melanogaster</i> , isolation stage (N_{c_2})	2.871	(2.498, 3.243)	(2.528, 3.276)

and

$$\hat{N}_a = \frac{\hat{\theta}_a}{4\hat{\mu}} = \text{constant} \times \hat{\theta}_a \quad .$$

Hence if we denote the vector of estimators of the converted parameters by $\hat{\phi}_c$, then $\hat{\phi}_c = \mathbf{W}\hat{\phi}$, where \mathbf{W} is a diagonal matrix and $\hat{\phi} = (\hat{M}_2\hat{b}/\hat{\theta}, \hat{\theta}_a, \dots)^T$. Because $\hat{\phi}$ is a maximum-likelihood estimator (of a reparameterised model), it approaches, as the sample size increases, a multivariate Gaussian distribution with some covariance matrix Σ ; hence $\hat{\phi}_c$ also approaches a multivariate Gaussian distribution, but with covariance matrix $\mathbf{W}\Sigma\mathbf{W}^T$. To calculate the Wald confidence intervals of Tables 4, 5 and 6, we reparameterised the IIM model in terms of $\phi = (M_2b/\theta, \theta_a, \dots)^T$ and found $\hat{\Sigma}$ as the inverse of the observed Fisher information. An estimate of $\mathbf{W}\Sigma\mathbf{W}^T$ followed trivially.

Table 5: Divergence time estimates for the data of Wang and Hey (2010) under the model IIM₃ (values in millions of years ago).

Event	Time since occurrence	95% Confidence intervals	
		Wald	Profile likelihood
Onset of speciation (t_0)	3.624	(3.559, 3.689)	(3.561, 3.691)
Complete isolation (t_1)	1.503	(1.419, 1.588)	(1.419, 1.587)

Note: These are the converted estimates of τ_0 and τ_1 (see Figure 1).

Profile likelihood confidence intervals were also computed for the parameterisation $\phi = (M_2b/\theta, \theta_a, \dots)^T$. Then, if $\hat{\mathbf{u}}$ (or $\hat{\mathbf{l}}$) is the vector of estimated upper (or lower) bounds for the parameters in ϕ , $\mathbf{W}\hat{\mathbf{u}}$ (or $\mathbf{W}\hat{\mathbf{l}}$) will be the vector of estimated upper (or lower) bounds for the converted parameters. This follows from the likelihood ratio invariance – see, for example, Pawitan (2001, p. 47-48).

Table 6: Converted migration rates for the data of Wang and Hey (2010) under the model IIM₃.

Migration parameter	Point Estimate	95% Confidence intervals	
		Wald	Profile likelihood
Migration rate (q_1)	8.8E-09	(1.1E-10, 1.8E-08)	(3.2E-09, 2.4E-08)
Number of migrant sequences (s_1)	0.064	(0.001, 0.127)	(0.023, 0.172)

Note: These are forward-in-time parameters; q_1 is the fraction of subpopulation 1 (*D. simulans*) that migrates to subpopulation 2 (*D. melanogaster*) in each generation, during the period of gene flow; s_1 is the number of sequences migrating from subpopulation 1 to subpopulation 2 in each generation, during the period of gene flow.

4 Discussion

We have described a fast method to fit the isolation-with-initial-migration model to large data sets of pairwise differences at a large number of independent loci. This method relies essentially on the eigendecomposition of the generator matrix of the process during the migration stage of the model: for each set of parameter values, the computation of the likelihood involves this decomposition. Nevertheless, the whole process of estimation takes no more than a couple of minutes for a data set of tens of thousands of loci such as that of Wang and Hey (2010). The implementation of the simpler IIM model of Wilkinson-Herbots (2012), with R code provided in Wilkinson-Herbots (2015), is even faster than the more general method presented here, since it makes use of a fully analytical expression for the likelihood (avoiding the need for eigendecomposition of the generator matrix), but it relies on two assumptions which we have dropped here, and which are typically unrealistic for real species: the symmetry of migration rates and the equality of subpopulation sizes during the gene flow period. When compared to implementations of the IM model based on numerical integration (such as *MDIV*, *IM*, *IMa* and *IMa2*), our method is considerably faster, as it does not require the use of high-performance computing resources. It is also more appropriate for studying different species (rather than subpopulations within the same species) than any IM model method, because it drops the assumption of gene flow until the present.

Due to the number of parameters, it is not feasible to assess the performance of our method systematically over every region of the parameter space. However, our experience with simulated data sets suggests that there are two cases in which the variances of some estimators become inflated, in particular the variances of the estimators associated with the gene flow period (\hat{M}_1 , \hat{M}_2 , $\hat{\theta}$, $\hat{\theta}_b$ and \hat{V}). One of such cases arises whenever V is very small or T_1 is very large, making it very unlikely that the genealogy of a pair of sequences under the IIM model is affected by events that occurred during the gene flow

Table 7: Comparison of converted estimates obtained with IM and IIM models

	IM _{wh}	IM ₁	IIM ₃
Time since onset of speciation	3.040	3.240	3.624
Time since isolation	-	-	1.503
Size of ancestral population	3.060	4.310	3.549
Current size of <i>D. sim.</i> population	5.990	6.120	7.182
Current size of <i>D. mel.</i> population	2.440	2.700	2.871
Size of <i>D. sim.</i> population during IIM gene flow period	-	-	3.640
Size of <i>D. mel.</i> population during IIM gene flow period	-	-	2.092
Migration rate (<i>D. sim.</i> → <i>D. mel.</i>)	0.013	0.012	0.064
Migration rate (<i>D. mel.</i> → <i>D. sim.</i>)	0.000	0.000	-

Note: Times are given in millions of years; population sizes are given in millions of individuals; the migration rates stated represent the number of sequences that migrate per generation, forward in time. The model IM_{wh} is the IM model fitted by Wang and Hey (2010).

period. The second case arises when the values of the scaled migration rates are greater than one, so that the two subpopulations during the period of gene flow resemble a single panmictic population. In either of these cases, the very process of model fitting can become unstable, that is, the algorithm of maximisation of the likelihood may have difficulty converging.

It is not the goal of this paper to draw conclusions regarding the evolutionary history of *Drosophila* species. We used the data of Wang and Hey (2010) with the sole objective of demonstrating that our method can be applied efficiently and accurately to real data. In Table 7, we list both our estimates and those of Wang and Hey (2010) for a six-parameter isolation-with-migration model (the IM₁ model – see Figure 4). The same table contains the estimates for our best-fitting IIM model. Our parameter estimates for the IM model agree well with those of Wang and Hey (2010). The reason that they do not match exactly lies in the fact that we have omitted the ‘screening procedure’ described in Wang and Hey (2010) and have therefore not excluded some of the most divergent sequences in the data set. It should also be borne in mind that our model of mutation is the infinite-sites model, whereas Wang and Hey (2010) have worked with the Jukes-Cantor model. Furthermore, our choice of sequence pairs was somewhat different: Wang and Hey (2010) randomly selected a pair of sequences at each locus, whereas we followed the procedure described in Section 3.2.1 above.

There are some notable differences between the estimates for both IM models and those for the IIM model: under the IIM model, the process of speciation is estimated to have started earlier (3.6 million years ago instead of 3.0 or 3.2 million years ago), to have reached complete isolation before the present time (1.5 million years ago), and to

have a higher rate of gene flow (0.064 sequences per generation instead of 0.013 or 0.012 sequences).

The method we used assumes that relative mutation rates are known (see section 3.2.1). In reality, we must deal with estimates of these rates, and this introduces additional uncertainty which is not reflected in the standard errors and confidence intervals obtained. However, it should be noted that, in principle, this uncertainty can be reduced to any extent desired: the method-of-moments estimator of the relative mutation rates – given by equation (14) – is a consistent estimator, so increasing the number of observations will get us arbitrarily close to the true value of the parameter. Ideally, these observations should stem from outgroup sequences only, to avoid any dependence between the estimates of relative mutation rates and the observations on ingroup pairwise differences (this was not possible here since the Wang and Hey (2010) data included exactly one outgroup sequence for each locus).

Some assumptions of our IIM model, such as the infinite-sites assumption and the assumption of free recombination between loci and no recombination within loci, may not be sensible for some real data sets. The appropriateness of other assumptions, for example those regarding the constant size of populations or the constant rate of gene flow, will depend on the actual evolutionary history of the species or populations involved. How robust the IIM model is to severe violations of its assumptions is a question which lies beyond the scope of this paper. It is nevertheless obvious that, in any case, the IIM model is more robust than any model nested in it, including the IM model and models of complete isolation.

As long as the data consist of pairwise differences, the method of eigendecomposition is easily applicable to several other models of speciation. This is true for any model nested in the full IIM model; and it should also be true for any model that consists of a sequence of island models (with or without gene flow) and Wright-Fisher populations.

Acknowledgements

We thank Ziheng Yang for some valuable discussions and helpful suggestions. We thank Yong Wang, Jody Hey, Nick Barton and Konrad Lohse for kindly providing the *Drosophila* DNA sequence data. This research was supported by the Engineering and Physical Sciences Research Council (grant number EP/K502959/1).

Supplementary material

In the ancillary files of this paper, we provide the R code to fit the IIM model and other simpler models, including the IM model, to data sets consisting of observations on the number of segregating sites between pairs of DNA sequences from a large number of independent loci.

Appendix

Proof of equation (8):

This proof has three parts. Part (i) proves the result under two assumptions: a) \mathbf{Q}_{mig} has three strictly negative eigenvalues and one zero eigenvalue, all of them real; and b) \mathbf{Q}_{mig} is diagonalizable. Part (ii) proves assumption a). Part (iii) proves assumption b). To simplify the notation, we denote \mathbf{Q}_{mig} by \mathbf{Q} throughout the proof.

(i)

Consider the continuous-time Markov chain defined by the matrix \mathbf{Q} . Let $P_{ij}(t)$, the (i, j) entry of the matrix $\mathbf{P}(t)$, be the probability that the process is in state j at time t into the past, given that the process starts in state i . $\mathbf{P}(t)$ can be calculated by solving the following initial value problem:

$$\begin{aligned}\mathbf{P}'(t) &= \mathbf{P}(t)\mathbf{Q} \ ; \\ \mathbf{P}(0) &= \mathbf{I}_4 \ ,\end{aligned}$$

where \mathbf{I}_4 is the four by four identity matrix. Under the assumptions that \mathbf{Q} is diagonalizable and that its eigenvalues are real, the solution to this initial value problem is given by:

$$\begin{aligned}\mathbf{P}(t) &= \mathbf{P}(0)e^{\mathbf{Q}t} \\ &= \mathbf{V}^{-1}e^{\mathbf{B}t}\mathbf{V} \ ,\end{aligned}$$

where \mathbf{B} denotes the diagonal matrix containing the real eigenvalues β_j , $j \in \{1, 2, 3, 4\}$, of \mathbf{Q} , and \mathbf{V} is the matrix of left eigenvectors of \mathbf{Q} . Note that $P_{i4}(t)$ is the probability that the process has reached coalescence by time t , if it started in state i . In other words, it is the *cdf* of $T_{\text{mig}}^{(i)}$:

$$P_{i4}(t) = F_{\text{mig}}^{(i)}(t) = \mathbf{v}_i^{-1}e^{\mathbf{B}t}\mathbf{v}_{\cdot 4} \ ,$$

where \mathbf{v}_i^{-1} is the i^{th} row vector of \mathbf{V}^{-1} , and $\mathbf{v}_{\cdot 4}$ the 4th column vector of \mathbf{V} . Differentiating, we get the *pdf*:

$$\begin{aligned}f_{\text{mig}}^{(i)}(t) &= \mathbf{v}_i^{-1}\mathbf{B}e^{\mathbf{B}t}\mathbf{v}_{\cdot 4} \\ &= \sum_{j=1}^4 V_{ij}^{-1}V_{j4}\beta_j e^{\beta_j t} \ .\end{aligned}$$

If we denote the eigenvalue equal to zero by β_4 , and the remaining eigenvalues are strictly negative, this *pdf* can be written as a linear combination of exponential densities:

$$f_{mig}^{(i)}(t) = - \sum_{j=1}^3 V_{ij}^{-1} V_{j4} \lambda_j e^{-\lambda_j t} \quad , \quad (16)$$

where $\lambda_j = |\beta_j|$ for $j \in \{1, 2, 3\}$.

(ii)

As \mathbf{Q} is given by equation (3), its characteristic polynomial, $\mathcal{P}_{\mathbf{Q}}(\beta)$, is of the form $\beta \times \mathcal{P}_{\mathbf{Q}^{(r)}}(\beta)$, where $\mathbf{Q}^{(r)}$ is the three by three upper-left submatrix of \mathbf{Q} , that is:

$$\mathbf{Q}^{(r)} = \begin{bmatrix} -(1 + M_1) & M_1 & 0 \\ M_2/2 & -(M_1 + M_2)/2 & M_1/2 \\ 0 & M_2 & -(1/b + M_2) \end{bmatrix} .$$

Thus the eigenvalues of \mathbf{Q} are the solutions to $\beta \times \mathcal{P}_{\mathbf{Q}^{(r)}}(\beta) = 0$. Consequently, one of them is zero (β_4 , say) and the remaining three eigenvalues are also eigenvalues of $\mathbf{Q}^{(r)}$.

Now consider the similarity transformation

$$\mathbf{S} = \mathbf{D}\mathbf{Q}^{(r)}\mathbf{D}^{-1} = \begin{bmatrix} -(1 + M_1) & \sqrt{\frac{M_1 M_2}{2}} & 0 \\ \sqrt{\frac{M_1 M_2}{2}} & -(M_1 + M_2)/2 & \sqrt{\frac{M_1 M_2}{2}} \\ 0 & \sqrt{\frac{M_1 M_2}{2}} & -(1/b + M_2) \end{bmatrix} ,$$

$$\text{where } \mathbf{D} = \begin{bmatrix} \sqrt{\frac{M_2}{2M_1}} & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & \sqrt{\frac{M_1}{2M_2}} \end{bmatrix} .$$

Because \mathbf{S} is a symmetric matrix, its eigenvalues are real. Therefore, all the eigenvalues of $\mathbf{Q}^{(r)}$ are real (a similarity transformation does not change the eigenvalues). \mathbf{S} is also a negative definite matrix, since its first, second and third upper-left determinants are respectively negative, positive, and negative. Hence its eigenvalues are all strictly negative, and so are the eigenvalues of $\mathbf{Q}^{(r)}$. Hence \mathbf{Q} has one zero eigenvalue (β_4) and

three real, strictly negative eigenvalues (β_1, β_2 and β_3).

(iii)

Being a symmetric matrix, \mathbf{S} has three independent eigenvectors. A similarity transformation preserves the number of independent eigenvectors, so $\mathbf{Q}^{(r)}$ has three independent eigenvectors as well. We denote by $\mathbf{V}^{(r)}$ the matrix whose rows are the left eigenvectors of $\mathbf{Q}^{(r)}$.

By definition, any left eigenvector \mathbf{v}_j of \mathbf{Q} satisfies the system of equations $\mathbf{x}(\mathbf{Q} - \mathbf{I}\beta_j) = \mathbf{0}$, where $\mathbf{x} = [x_1 \ x_2 \ x_3 \ x_4]$. The first three linear equations of this system are identical to $\mathbf{x}^{(r)}(\mathbf{Q}^{(r)} - \mathbf{I}\beta_j) = \mathbf{0}$, for $j \in \{1, 2, 3\}$ and $\mathbf{x}^{(r)} = [x_1 \ x_2 \ x_3]$, which is solved by $\mathbf{x}^{(r)} = \mathbf{v}_j^{(r)}$. So this implies that, for $\beta_j \in \{\beta_1, \beta_2, \beta_3\}$, any row vector \mathbf{x} in \mathbb{R}^4 that has $\mathbf{v}_j^{(r)}$ as its first three elements will solve the first three equations of the system, whatever the value of x_4 . If $x_4 = (V_{j1}^{(r)} + \frac{1}{b}V_{j3}^{(r)})/\beta_j$, that vector will be an eigenvector of \mathbf{Q} , because it also solves the fourth equation of the system:

$$\left[\text{---} \mathbf{v}_j^{(r)} \text{---} \quad \frac{V_{j1}^{(r)} + \frac{1}{b}V_{j3}^{(r)}}{\beta_j} \right] \begin{bmatrix} -(1 + M_1) - \beta_j & M_1 & 0 & 1 \\ \frac{M_2}{2} & -\frac{(M_1 + M_2)}{2} - \beta_j & \frac{M_1}{2} & 0 \\ 0 & M_2 & -(\frac{1}{b} + M_2) - \beta_j & \frac{1}{b} \\ 0 & 0 & 0 & -\beta_j \end{bmatrix} \\ = \begin{bmatrix} 0 & 0 & 0 & 0 \end{bmatrix},$$

for $\beta_j \in \{\beta_1, \beta_2, \beta_3\}$. For the case of $\beta_j = \beta_4 = 0$, a row eigenvector is $[0 \ 0 \ 0 \ 1]$. Collecting these row eigenvectors in a single matrix, we get \mathbf{V} . So,

$$\mathbf{V} = \begin{bmatrix} \text{---} \mathbf{v}_1^{(r)} \text{---} & \frac{(V_{11}^{(r)} + \frac{1}{b}V_{13}^{(r)})}{\beta_1} \\ \text{---} \mathbf{v}_2^{(r)} \text{---} & \frac{(V_{21}^{(r)} + \frac{1}{b}V_{23}^{(r)})}{\beta_2} \\ \text{---} \mathbf{v}_3^{(r)} \text{---} & \frac{(V_{31}^{(r)} + \frac{1}{b}V_{33}^{(r)})}{\beta_3} \\ 0 & 0 & 0 & 1 \end{bmatrix}.$$

If the matrix \mathbf{V} can be shown to be invertible, then \mathbf{Q} is diagonalizable. This will be the case if the system $\mathbf{x}\mathbf{V} = \mathbf{0}$ can only be solved by $\mathbf{x} = [0 \ 0 \ 0 \ 0]$. Now since the three

by three upper-left submatrix of \mathbf{V} , $\mathbf{V}^{(r)}$, is full-ranked, $x_1 = x_2 = x_3 = 0$ is a necessary condition for $\mathbf{xV} = \mathbf{0}$. But then $x_4 = 0$, from the last equation of the system. Thus we have shown that \mathbf{Q} is diagonalizable. \square

Bibliography

- Andersen, L., T. Mailund, and A. Hobolth (2014). Efficient computation in the IM model. *Journal of Mathematical Biology* 68(6), 1423–1451.
- Becquet, C. and M. Przeworski (2009). Learning About Modes of Speciation By Computational Approaches. *Evolution* 63(10), 2547–2562.
- Casella, G. and R. Berger (2001). *Statistical Inference* (2nd ed.). Duxbury.
- Fisher, R. A. (1930). *The Genetical Theory of Natural Selection* (1st ed.). Oxford: Clarendon Press.
- Hein, J., M. H. Schierup, and C. Wiuf (2005). *Gene Genealogies, Variation and Evolution*. Oxford.
- Hey, J. (2005). On the Number of New World Founders: A Population Genetic Portrait of the Peopling of the Americas. *PLoS Biol* 3(6), e193.
- Hey, J. (2010). Isolation with Migration Models for More Than Two Populations. *Molecular Biology and Evolution* 27(4), 905–920.
- Hey, J. and R. Nielsen (2004). Multilocus Methods for Estimating Population Sizes, Migration Rates and Divergence Time, With Applications to the Divergence of *Drosophila pseudoobscura* and *D. persimilis*. *Genetics* 167(2), 747–760.
- Hey, J. and R. Nielsen (2007). Integration within the Felsenstein equation for improved Markov chain Monte Carlo methods in population genetics. *Proceedings of the National Academy of Sciences* 104(8), 2785–2790.
- Hutter, S., H. Li, S. Beisswanger, D. De Lorenzo, and W. Stephan (2007). Distinctly different sex ratios in african and european populations of drosophila melanogaster inferred from chromosomewide single nucleotide polymorphism data. *Genetics* 177(1), 469–480.
- Kingman, J. F. (1982a). The Coalescent. *Stochastic Processes and Their Applications* 13(3), 235–248.
- Kingman, J. F. C. (1982b). On the Genealogy of Large Populations. *Journal of Applied Probability* 19, 27–43.

- Kopylev, L. and B. Sinha (2011). On the asymptotic distribution of likelihood ratio test when parameters lie on the boundary. *Sankhya B* 73(1), 20–41.
- Lohse, K., R. J. Harrison, and N. H. Barton (2011). A general method for calculating likelihoods under the coalescent process. *Genetics* 189(3), 977–987.
- Nielsen, R. and J. Wakeley (2001). Distinguishing migration from isolation: A Markov chain Monte Carlo approach. *Genetics* 158(2), 885–896.
- Notohara, M. (1990). The coalescent and the genealogical process in geographically structured population. *Journal of Mathematical Biology* 29(1), 59–75.
- Pawitan, Y. (2001). *In All Likelihood*. Oxford University Press.
- Pinho, C. and J. Hey (2010). Divergence with Gene Flow: Models and Data. *Annual Review of Ecology, Evolution, and Systematics* 41(1), 215–230.
- Powell, J. R. (1997). *Progress and Prospects in Evolutionary Biology: The Drosophila Model*. Oxford University Press.
- Self, S. G. and K.-Y. Liang (1987). Asymptotic properties of maximum likelihood estimators and likelihood ratio tests under nonstandard conditions. *Journal of the American Statistical Association* 82(398), 605–610.
- Sousa, V. C., A. Grelaud, and J. Hey (2011). On the nonidentifiability of migration time estimates in isolation with migration models. *Molecular Ecology* 20(19), 3956–3962.
- Strasburg, J. L. and L. H. Rieseberg (2011). Interpreting the estimated timing of migration events between hybridizing species. *Molecular Ecology* 20(11), 2353–2366.
- Wang, Y. and J. Hey (2010). Estimating divergence parameters with small samples from a large number of loci. *Genetics* 184(2), 363–379.
- Watterson, G. (1975). On the number of segregating sites in genetical models without recombination. *Theoretical Population Biology* 7(2), 256–276.
- Wilkinson-Herbots, H. (2012). The distribution of the coalescence time and the number of pairwise nucleotide differences in a model of population divergence or speciation with an initial period of gene flow. *Theoretical Population Biology* 82(2), 92–108.
- Wilkinson-Herbots, H. (2015). A fast method to estimate speciation parameters in a model of isolation with an initial period of gene flow and to test alternative evolutionary scenarios. *ArXiv e-prints*. URL: <http://arxiv.org/abs/1511.05478> .

Wright, S. (1931). Evolution in mendelian populations. *Genetics* 16(2), 97–159.

Yang, Z. (2002). Likelihood and Bayes estimation of ancestral population sizes in hominoids using data from multiple loci. *Genetics* 162(4), 1811–1823.

Zhu, T. and Z. Yang (2012). Maximum likelihood implementation of an isolation-with-migration model with three species for testing speciation with gene flow. *Molecular Biology and Evolution* 29(10), 3131–3142.