



HAL
open science

Likelihood-Free Parallel Tempering

Meili Baragatti, Agnès Grimaud, Denys Pommeret

► **To cite this version:**

Meili Baragatti, Agnès Grimaud, Denys Pommeret. Likelihood-Free Parallel Tempering. 2011. hal-00614873v2

HAL Id: hal-00614873

<https://hal.science/hal-00614873v2>

Preprint submitted on 29 Dec 2011 (v2), last revised 27 Apr 2012 (v3)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

LIKELIHOOD-FREE PARALLEL TEMPERING.

Meili Baragatti^{1,2,*}, Agnès Grimaud², Denys Pommeret²

¹ *Ipsogen SA, Luminy Biotech Entreprises, Case 923, Campus de Luminy, 13288 Marseille Cedex 9, France.*

² *Institut de Mathématiques de Luminy (IML), CNRS Marseille, case 907, Campus de Luminy, 13288 Marseille Cedex 9, France.*

* *baragatt@iml.univ-mrs.fr, baragattmeili@hotmail.com.*

WORKING PAPER DECEMBER 29, 2011

Abstract

Approximate Bayesian Computational (ABC) methods (or likelihood-free methods) have appeared in the past fifteen years as useful methods to perform Bayesian analysis when the likelihood is analytically or computationally intractable. Several ABC methods have been proposed: Monte Carlo Markov Chains (MCMC) methods have been developed by Marjoram et al. [2003] and by Bortot et al. [2007] for instance, and sequential methods have been proposed among others by Sisson et al. [2007], Beaumont et al. [2009] and Del Moral et al. [2012]. Until recently, while ABC-MCMC methods were the reference, sequential ABC methods have appeared to outperform them (see for example McKinley et al. [2009] or Sisson et al. [2007]). In this paper a new algorithm combining population-based MCMC methods with ABC requirements is proposed, using an analogy with the Parallel Tempering algorithm [Geyer, 1991]. Performances are compared with existing ABC algorithms on simulations and on a real example.

Keywords: Approximated Bayesian Computational, likelihood-free, intractable likelihood, parallel tempering, population-based, Monte Carlo Markov chain.

1 Introduction

The principle of Approximate Bayesian Computation (ABC) algorithms is to evaluate the posterior density when the likelihood function is intractable. More precisely, let $x \in \mathbb{D} \subseteq \mathbb{R}^n$ be an observed vector, $\theta \in \mathbb{R}^d$ a parameter of interest, and $\pi(\cdot)$ the prior distribution of θ . The aim of an ABC algorithm is to evaluate the following posterior density:

$$\pi(\theta | x) \propto f(x | \theta)\pi(\theta).$$

The likelihood $f(x | \theta)$ is supposed analytically or computationally intractable. Hence the first requirement is that its calculation is not necessary. The second requirement is the relative ease of simulation from the model, given a parameter θ . These requirements are fulfilled in several areas in which ABC methods have been used, including genetics (Pritchard et al. [1999], Leuenberger et al. [2010]), extreme values theory (Bortot et al. [2007]), Gibbs random fields (Grelaud et al. [2009]) or epidemiology (Tanaka et al. [2006], Sisson et al. [2007]).

The first ABC algorithm has been proposed by Tavaré et al. [1997], and is simply a rejection algorithm. It consists in generating jointly $\theta' \sim \pi$ and $z' \sim f(\cdot | \theta')$. The generated θ' is accepted if the equality $z' = x$ holds, and is then denoted by θ_1 . This process is repeated until n generated θ' have been accepted. The accepted $\{\theta_1, \theta_2, \dots, \theta_n\}$ form an i.i.d. sample from the target posterior $\pi(\cdot | x)$, that is why this algorithm is often called the exact ABC algorithm. It has been extended by Pritchard et al. [1999] in order to be applied in continuous state spaces. The idea is to accept a simulated z' if it is sufficiently close to the observed x . The condition $z' = x$ is replaced by $\rho(z', x) < \varepsilon$, with a distance $\rho(\cdot, \cdot)$ and a tolerance level ε . In case of large datasets z' and x , it can be computationally demanding to calculate $\rho(z', x)$. Hence a second refinement is to compare summary statistics instead of the complete datasets, and the condition becomes $\rho(S(z'), S(x)) < \varepsilon$, with S a statistic which can be multi-dimensional. The algorithm proposed by Pritchard et al. [1999] includes these two refinements and is considered as the standard of the ABC algorithms. This algorithm gives an i.i.d. sample from the following joint posterior distribution:

$$\pi_\varepsilon(z, \theta | x) \propto \pi(\theta) f(z | \theta) \mathbb{1}_{\{\rho(S(z), S(x)) < \varepsilon\}}(z). \quad (1)$$

It is usually the marginal in θ which is of interest; that is $\pi_\varepsilon(\theta | x) \propto \int \pi_\varepsilon(z, \theta | x) dz$. If S provides a good summary of the data and if ε is sufficiently small, then $\pi_\varepsilon(\theta | x)$ can be considered as a reasonable approximation of the target posterior $\pi(\theta | x)$. Of course, the quality of the obtained approximation depends on the calibration of the algorithm by choosing ε , $\rho(\cdot, \cdot)$ and S .

The acceptance rates of the standard ABC algorithm can be very low, hence several improvements have been proposed, using MCMC or sequential schemes.

1.1 Improvements using MCMC schemes and associated issues

Algorithms have been proposed by Marjoram et al. [2003], Bortot et al. [2007] and Ratmann et al. [2007], the reference method being the one of Marjoram et al. [2003] who suggested the use of a Markov kernel $q(\cdot | \cdot)$ to propose a new parameter from a previous accepted one. The validity of this algorithm called ABC-MCMC is ensured as the stationary distribution for the Markov chain $(\theta^{(t)}, z^{(t)})$ is $\pi_\varepsilon(z, \theta | x)$. The acceptance rates of this algorithm are higher than those obtained with a standard ABC algorithm, but two drawbacks are observed:

- In case of small ε the sampler can encounter difficulties to move in the parameter space. Indeed, the acceptance rate is proportional to the probability of simulating a z' such that

$\rho(S(z'), S(x)) < \varepsilon$, which is itself proportional to the likelihood (see Sisson et al. [2007]). This probability is relatively low if ε is small, and in practice when the ABC-MCMC enters an area of relatively low probability, the acceptance probabilities to move elsewhere are quite small and the sampler encounters difficulties to move from this part of the space. As a consequence, not only the sampler sticks in the distribution tails, but also these tails are not well visited.

- The samples obtained are dependent, and often highly correlated. This is inherent in all methods based on MCMC, and not necessarily a problem, but if one needs non-correlated particles, thinning of the chain should be done.

Modified toy example and ABC-MCMC

To illustrate the behavior of the ABC-MCMC in distribution tails, we use a modified toy example inspired from the example studied in Sisson et al. [2007], Beaumont et al. [2009] and Del Moral et al. [2012]. The prior for θ is uniform $\mathcal{U}(-10, 10)$, the model is a mixture of Gaussian distributions given by $x | \theta \sim 0.45\mathcal{N}(\theta, 1) + 0.45\mathcal{N}(\theta, 1/100) + 0.1\mathcal{N}(\theta - 5, 1)$ (1 and 1/100 denoting the variances), and the posterior distribution associated with an observation $x = 0$ is the mixture:

$$\theta | x = 0 \sim \left(0.45\mathcal{N}(0, 1) + 0.45\mathcal{N}(0, 1/100) + 0.1\mathcal{N}(5, 1)\right) \mathbf{1}_{[-10, 10]}(\theta).$$

Using $S(x) = x$ and $\rho(x, z) = |z - x|$, the approximation of this posterior $\pi(\theta | x = 0)$ can be written as:

$$\pi_\varepsilon(\theta | x = 0) \propto 0.45 \left[\Phi(\varepsilon - \theta) + \Phi(10(\varepsilon - \theta)) - \Phi(-\varepsilon - \theta) - \Phi(10(-\varepsilon - \theta)) \right] + 0.1 \left[\Phi(\varepsilon - \theta + 5) - \Phi(-\varepsilon - \theta + 5) \right],$$

where Φ denotes the distribution function of a $\mathcal{N}(0, 1)$. Using a tolerance level $\varepsilon = 0.025$, the approximated posterior is indistinguishable from the exact posterior density. As in Sisson et al. [2007] and Beaumont et al. [2009], we used a $\mathcal{N}(\theta^{(t-1)}, 0.15^2)$ as the proposal $q(\cdot | \theta^{(t-1)})$. The algorithm was run on $20 \cdot 10^6$ iterations with a burn-in of $5 \cdot 10^6$, and it appears that for some runs the chain did not succeed to visit the small mode around 5, while for others it stayed stuck in this local mode and does not visit well the global mode around 0. The Figure 1 shows the results of two runs. It is clear that these results are not satisfactory. As a comparison, the population obtained by a standard ABC is given in Figure 12 in Appendix A. It is quite good, hence the issue observed in the distribution tail is specific to MCMC, it is not due to the ABC approximation of the likelihood.

1.2 Contribution of the paper

Until now ABC methods using MCMC schemes were not completely satisfactory. Hence sequential algorithms have been proposed by different authors: Sisson et al. [2007], Beaumont et al. [2009] (similar algorithms have been proposed by Sisson et al. [2009] and Toni et al. [2009]), and more recently Del Moral et al. [2012] (a similar algorithm has been proposed by Drovandi and Pettitt [2011]). The ABC-PMC of Beaumont et al. [2009] is simple to use and to implement, and it gives non-biased results from the approximate posterior $\pi_\varepsilon(\theta | x)$. A more recent algorithm is the ABC-SMC of Del Moral et al. [2012]: both tolerance levels and transition kernels are automatically calibrated at each iteration, and several datasets are generated given a proposed parameter. These sequential ABC schemes tend to become the references in ABC methods. However, to our knowledge, there is nothing in the literature about ABC methods based on

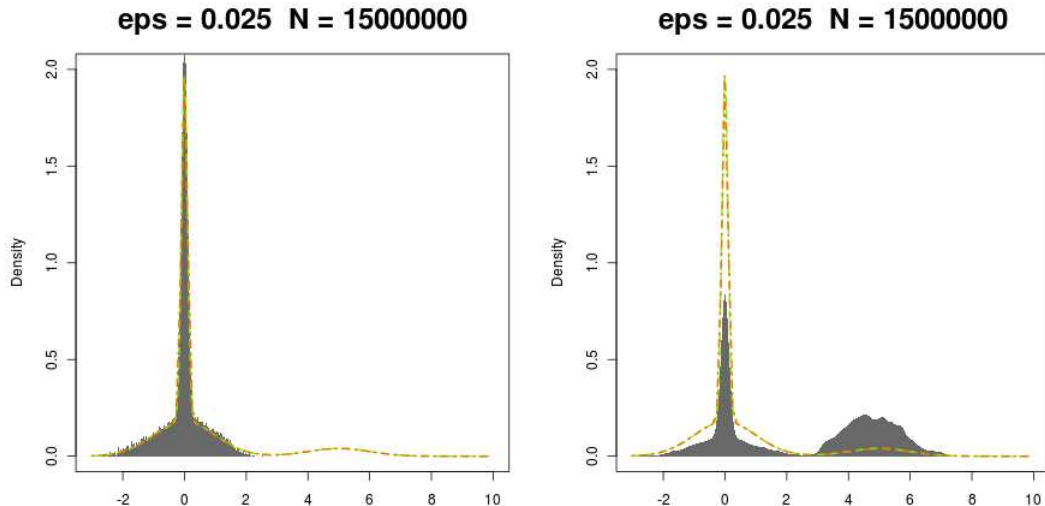


Figure 1: ABC-MCMC for the modified toy example: samples obtained by two runs ($20 \cdot 10^6$ iterations with $5 \cdot 10^6$ iterations of burn-in) for tolerance levels of 0.025. The orange dotted curves stand for the exact posterior distribution, the green dotted curves for the target approximation, and densities of obtained samples are represented in gray.

MCMC schemes and using a population of chains.

In this paper we are interested in developing such a method, by doing an analogy with the Parallel Tempering algorithm used in classical MCMC methods [Geyer, 1991, Geyer and Thompson, 1995]. In particular, we want to know if using a population-based approach the drawback of ABC-MCMC schemes concerning distribution tails can be avoided. Moreover, we would like to compare the behavior of such an approach with the ABC-PMC and ABC-SMC, especially in distribution tails.

The paper is organized as follows. In Section 2 we introduce the new population-based ABC-MCMC scheme, that we call ABC-PT algorithm. In Section 3 this algorithm is illustrated on the modified toy example and on real data, and comparisons with the ABC-MCMC, ABC-PMC and ABC-SMC algorithms are done. Finally, the method and the results are discussed in Section 4.

2 The ABC-PT algorithm

2.1 Standard Parallel Tempering (PT)

When the goal is to sample from a multi-modal target density π , classical MCMC methods like Metropolis-Hastings algorithm [Metropolis et al., 1953, Hastings, 1970] or Gibbs sampler [Gelfand and Smith, 1990] for instance, are often trapped into local modes from where they cannot escape in reasonable time. To avoid this problem, the principle of PT is to introduce N temperatures and to run in parallel N associated MCMC chains with target distributions being tempered distributions of the target π (the i th chain is associated with $\pi_i \propto \pi^{1/T_i}$). The first chain targets π , and since the tempered distributions become flatter as the temperature increases, the chains at

high temperatures can move easily between modes and explore the whole parameter space. Each iteration of the PT algorithm is decomposed into two types of moves: local moves via classical MCMC algorithms to update the different chains, and global moves allowing swaps between two chains. The use of these swaps enables new modes to be propagated through the different chains, thereby improving mixing. The first chain associated with the target distribution will then be able to escape from local modes.

2.2 Analogy with the PT in an ABC framework

The analogy with the PT is the following: a population of N chains is used, chains of higher orders should be able to move easily in the parameter space, and chains of lower orders should give precise approximations of the target posterior. These last chains can have some difficulties to exit local modes, so exchange moves between chains should be defined so that the chains of lower orders can visit all the parameter space, and hence can exit low probability areas.

To achieve this in an ABC framework, chains of higher orders are associated with high tolerance levels, while chains of lower orders are associated with low tolerance levels. A sequence of tolerance levels is then defined as $\boldsymbol{\varepsilon} = (\varepsilon_1, \varepsilon_2, \dots, \varepsilon_N)$, with $\varepsilon_1 < \varepsilon_2 < \dots < \varepsilon_N$. The i th chain is associated with ε_i . In particular the first chain is associated with ε_1 the tolerance level required for the approximation of the posterior of interest, and the chain associated with the higher tolerance level ε_N should be able to exit areas of low posterior probabilities.

The set of parameters associated to the N chains is written $\boldsymbol{\theta} = (\theta_1, \theta_2, \dots, \theta_N)$ with $\theta_i \in \Theta \subseteq \mathbb{R}^d$, and $\boldsymbol{z} = (z_1, z_2, \dots, z_N)$ denotes the associated simulated datasets. The state associated to the i th chain ($i \in \{1, \dots, N\}$) is (z_i, θ_i) .

2.2.1 Local moves

Each chain is locally updated using an iteration of an ABC-MCMC algorithm, with associated tolerance level and transition kernel. In order to improve the exploration of the parameter space by higher order chains, the transition kernels associated with these chains should enable large moves. As the tolerance levels associated with these chains are high, the probabilities to accept the proposed moves can remain reasonably high. On the opposite, transition kernels associated with lower order chains should propose smaller moves, to have more chance to accept the proposed moves. We then propose to use tempered transition kernels, using a sequence of temperatures $T_1 = 1 < T_2 < \dots < T_N$ (such tempered transitions kernels have been used by Ratmann et al. [2007]). In particular the i th chain is associated with ε_i , T_i and the transition kernel $q_i(\cdot | \cdot)$. For instance $q_i(\cdot | \cdot)$ can be a Gaussian transition kernel tempered by T_i . This i th chain is associated with the following joint distribution:

$$\pi_{\varepsilon_i}(z_i, \theta_i | x) \propto \pi(\theta_i) f(z_i | \theta_i) \mathbb{1}_{\{\rho(S(z_i), S(x)) < \varepsilon_i\}}(z_i).$$

We also introduce:

$$\pi_{\boldsymbol{\varepsilon}}^*(\boldsymbol{z}, \boldsymbol{\theta} | x) = \prod_{i=1}^N \pi_{\varepsilon_i}(z_i, \theta_i | x).$$

2.2.2 Exchange moves

To perform an exchange move between two chains, a pair of chains is chosen uniformly. The orders of these chains are denoted by i and j ($i < j$), and the states proposed by the exchange

move are written $\boldsymbol{\theta}'$ and \mathbf{z}' . We denote by $q(\mathbf{z}', \boldsymbol{\theta}' | \mathbf{z}, \boldsymbol{\theta})$ the probability of proposing $(\mathbf{z}', \boldsymbol{\theta}')$ from $(\mathbf{z}, \boldsymbol{\theta})$. We propose to exchange the state (z_i, θ_i) of chain i with the state (z_j, θ_j) of chain j . We then have:

$$\boldsymbol{\theta}' = (\theta_1, \dots, \theta_j, \dots, \theta_i, \dots, \theta_N) \quad \text{and} \quad \mathbf{z}' = (z_1, \dots, z_j, \dots, z_i, \dots, z_N).$$

The acceptance probability to ensure reversibility of the move will be the following (using the classical detailed balance condition):

$$\alpha = 1 \wedge \frac{\pi_{\epsilon}^*(\mathbf{z}', \boldsymbol{\theta}' | x) q(\mathbf{z}, \boldsymbol{\theta} | \mathbf{z}', \boldsymbol{\theta}')}{\pi_{\epsilon}^*(\mathbf{z}, \boldsymbol{\theta} | x) q(\mathbf{z}', \boldsymbol{\theta}' | \mathbf{z}, \boldsymbol{\theta})} = 1 \wedge \frac{\pi_{\epsilon}^*(\mathbf{z}', \boldsymbol{\theta}' | x)}{\pi_{\epsilon}^*(\mathbf{z}, \boldsymbol{\theta} | x)}. \quad (2)$$

Indeed, as the pair of chains is chosen uniformly among all the possible pairs, it is clear that $q(\mathbf{z}, \boldsymbol{\theta} | \mathbf{z}', \boldsymbol{\theta}') = q(\mathbf{z}', \boldsymbol{\theta}' | \mathbf{z}, \boldsymbol{\theta})$. In addition we have:

$$\begin{aligned} \frac{\pi_{\epsilon}^*(\mathbf{z}', \boldsymbol{\theta}' | x)}{\pi_{\epsilon}^*(\mathbf{z}, \boldsymbol{\theta} | x)} &= \frac{\pi_{\epsilon_i}(z'_i, \theta'_i | x) \pi_{\epsilon_j}(z'_j, \theta'_j | x)}{\pi_{\epsilon_i}(z_i, \theta_i | x) \pi_{\epsilon_j}(z_j, \theta_j | x)} \\ &= \frac{f(z'_i | \theta'_i) \mathbb{1}_{\{\rho(S(z'_i), S(x)) < \epsilon_i\}}(z'_i)}{f(z_i | \theta_i) \mathbb{1}_{\{\rho(S(z_i), S(x)) < \epsilon_i\}}(z_i)} \times \frac{f(z'_j | \theta'_j) \mathbb{1}_{\{\rho(S(z'_j), S(x)) < \epsilon_j\}}(z'_j)}{f(z_j | \theta_j) \mathbb{1}_{\{\rho(S(z_j), S(x)) < \epsilon_j\}}(z_j)} \\ &= \frac{\mathbb{1}_{\{\rho(S(z_j), S(x)) < \epsilon_i\}}(z_j) \mathbb{1}_{\{\rho(S(z_i), S(x)) < \epsilon_j\}}(z_i)}{\mathbb{1}_{\{\rho(S(z_i), S(x)) < \epsilon_i\}}(z_i) \mathbb{1}_{\{\rho(S(z_j), S(x)) < \epsilon_j\}}(z_j)}. \end{aligned} \quad (3)$$

By construction $\epsilon_i < \epsilon_j$, and $\mathbb{1}_{\{\rho(S(z_i), S(x)) < \epsilon_i\}}(z_i) \mathbb{1}_{\{\rho(S(z_j), S(x)) < \epsilon_j\}}(z_j) = 1$. Therefore (3) simply reduces to $\mathbb{1}_{\{\rho(S(z_j), S(x)) < \epsilon_i\}}(z_j)$, and the exchange move is accepted if $\{\rho(S(z_j), S(x)) < \epsilon_i\}$. This exchange move satisfies ABC requirements as there is no need to compute any likelihood. We can note that this acceptance rate depends only on the tolerance level ϵ_i .

2.3 ABC-PT algorithm

Define a sequence of tolerance levels $\epsilon_1 < \epsilon_2 < \dots < \epsilon_N$, and a sequence of temperatures $T_1 = 1 < T_2 < \dots < T_N$.

ABC-PT

To obtain a Markov chain of length n :

1. $t = 1$. Use the standard ABC algorithm to obtain N realizations:

For i in $\{1, \dots, N\}$, generate $(\theta_i^{(0)}, z_i^{(0)})$ from $\pi_{\epsilon_i}(z, \theta | x)$.

Set $\boldsymbol{\theta}^{(0)} = (\theta_1^{(0)}, \theta_2^{(0)}, \dots, \theta_N^{(0)})$ and $\mathbf{z}^{(0)} = (z_1^{(0)}, z_2^{(0)}, \dots, z_N^{(0)})$.

2. For t in $\{2, \dots, n\}$:

- (a) LOCAL MOVES ABC-MCMC: For i in $\{1, \dots, N\}$:

i. Generate θ'_i from $q_i(\cdot | \theta_i^{(t-1)})$ depending on T_i .

ii. Generate z'_i given θ'_i , from the model $f(\cdot | \theta'_i)$.

Set $(\theta_i^{(t)}, z_i^{(t)}) = (\theta_i', z_i')$ with probability

$$\alpha = \min \left\{ 1, \frac{\pi(\theta_i') q_i(\theta_i^{(t-1)} | \theta_i')}{\pi(\theta_i^{(t-1)}) q_i(\theta_i' | \theta_i^{(t-1)})} \right\} \mathbb{1}_{\{\rho(S(z_i'), S(x)) < \varepsilon_i\}}(z_i').$$

Else set $(\theta_i^{(t)}, z_i^{(t)}) = (\theta_i^{(t-1)}, z_i^{(t-1)})$.

(b) **EXCHANGE MOVES:** N exchange moves are proposed:

N pairs of chains are chosen uniformly in all possible pairs with replacement. For each of them: Denote by i and j ($i < j$) the orders of a chosen pair of chains.

Exchange (z_i, θ_i) with (z_j, θ_j) with probability $\mathbb{1}_{\{\rho(S(z_j), S(x)) < \varepsilon_i\}}(z_j)$.

At each iteration N exchange moves are proposed. Indeed, it appears from our experience that only one proposition is not sufficient to enable a good mixing of the first chain of interest. Moreover increasing the number of exchange moves improves the mixing with the first chain of interest and allows to obtain less correlated samples.

3 Illustrations on modified toy example

3.1 ABC-PT

We use the modified toy example presented in the Introduction to study the behavior of the ABC-PT algorithm, especially in distribution tails and in local posterior modes. We take $N = 15$ chains, and imitating the ABC-MCMC algorithm, the tempered transition kernels used are $q_i(\cdot | \theta_i^{(t-1)}) \sim \mathcal{N}(\theta_i^{(t-1)}, 0.15^2 T_i)$, $i = 1, \dots, N$. Concerning the sequence of temperatures, we choose $T_1 = 1$ and $T_N = 4$. The other temperatures can be chosen by evenly spacing them on a logarithmic scale, by evenly spacing their inverses, or by evenly spacing their inverses geometrically (see for instance Kou et al. [2006], Nagata and Watanabe [2008] or Neal [1996]). In this example we decide to space them on a logarithmic scale. Concerning the tolerance levels, once ε_1 and ε_N are chosen, it gives satisfactory results to choose the others by using a sequence of proportionality constants (approach similar to the one of Beaumont et al. [2009]) or by regularly spacing them on a logarithmic scale (which is a special case of a sequence of proportionality constants). In this example they are chosen evenly spaced on a logarithmic scale between $\varepsilon_1 = 0.025$ and $\varepsilon_N = 2$. We checked that the highest order chain associated with $T_N = 4$ and $\varepsilon_N = 2$ easily explores the parameter space. The algorithm is run for 600 000 iterations with a burn-in period of 150 000 iterations.

Figures 2 and 3 show the obtained approximations $\pi_{\varepsilon_i}(\cdot | x)$ for the N chains in parallel. They are satisfactory, the two modes being well visited. Table 1 presents the acceptance rates and the numbers of accepted exchange moves for six of the 15 chains: on average, 2.05 moves are accepted per iteration (out of 15).

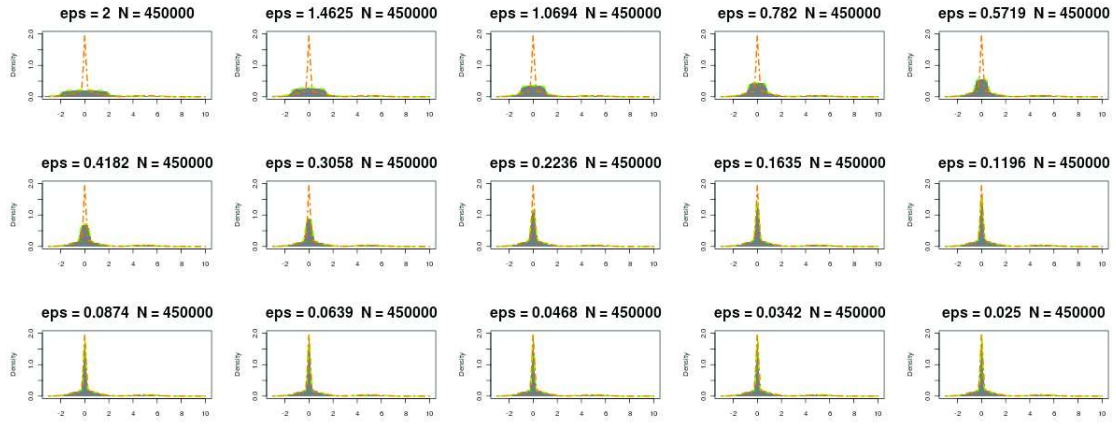


Figure 2: ABC-PT for the modified toy example: approximations obtained by each of the N chains, the associated tolerance levels going from 0.025 to 2 (600 000 iterations with a burn-in of 150 000). The orange dotted curves stand for the exact posterior distribution, the green dotted curves for the target approximation, and densities of obtained samples are represented in gray.

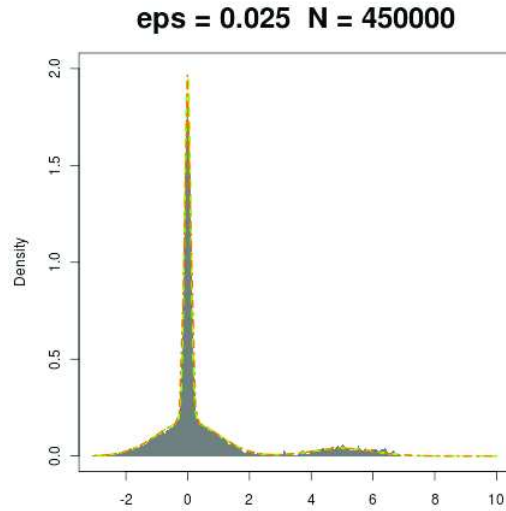


Figure 3: ABC-PT for the modified toy example: approximations obtained by the chain associated with tolerance level of 0.025 (600 000 iterations with a burn-in of 150 000). The orange dotted curves stand for the exact posterior distribution, the green dotted curves for the target approximation, and densities of obtained samples are represented in gray.

ϵ	2	0.78	0.31	0.12	0.047	0.025
Local acceptance rates	65.8%	45.6%	26.0%	12.5%	5.4%	3.0%
Accepted exchange moves	106 711	171 678	191 311	187 557	156 676	107 330

Table 1: Modified toy example: acceptance rates for local moves and number of accepted exchange moves, for chains 15, 12, 9, 6, 3 and 1 of the ABC-PT.

Compared with the approximations obtained with the ABC-MCMC algorithm, the results obtained with ABC-PT are greatly improved in terms of exploration of the parameter space. Indeed, the first chain of interest succeeds to visit the small mode around 5, and it is not trapped in this local mode. Moreover, the results of different runs appeared very similar indicating that the results of ABC-PT are reproducible. Table 2 shows that there is also an improvement in terms of autocorrelations. In Figure 4 the traces of the two algorithms are given, and it is observed that the mixing of the chain is improved by ABC-PT compared to ABC-MCMC.

	thinning	Nb of samples	Order 1	Order 10	Order 20
ABC-MCMC	none	$15 \cdot 10^6$	1	1	1
	100	$15 \cdot 10^4$	0.999	0.994	0.990
	1000	$15 \cdot 10^3$	0.994	0.968	0.951
ABC-PT (chain 1)	none	$45 \cdot 10^4$	0.842	0.367	0.284
	10	$45 \cdot 10^3$	0.361	0.198	0.142
	50	$9 \cdot 10^3$	0.215	0.093	0.082

Table 2: Modified toy example: autocorrelations of order 1, 10 and 20 of the chain obtained by ABC-MCMC, and of the first chain obtained by ABC-PT (with or without thinning the chains).

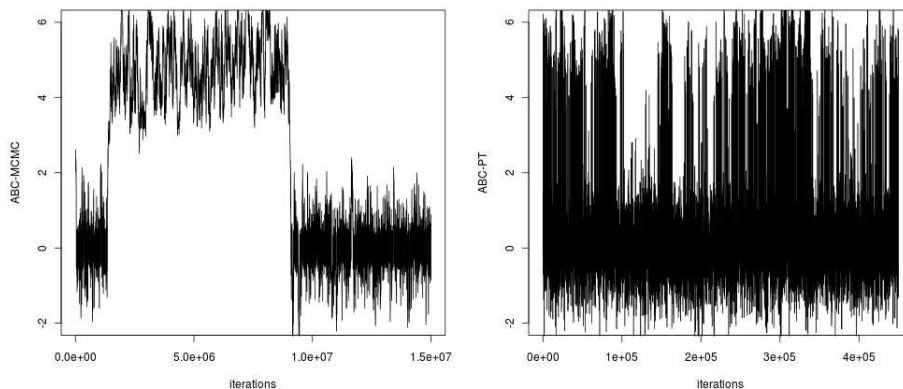


Figure 4: Modified toy example: On the left the trace of a chain obtained by ABC-MCMC (example corresponding to the right part of Figure 1), on the right the trace of a first chain obtained by ABC-PT.

3.2 Comparison with ABC-PMC

The ABC-PMC of Beaumont et al. [2009] is run on the modified toy example. The decreasing sequence of tolerance levels used is the same than in Beaumont et al. [2009], that is $(2, 1.5, 1, 0.5, 0.1, 0.025)$. Six weighted populations of size 5000 are generated, and represented in Figure 5. The results are much better than those obtained by the ABC-MCMC. In particular, the sampler is not stuck in distribution tails. However it seems that as the tolerance level decreases, the tails are under-covered and in particular the small mode around 5, see the left of Figure 6. That is also observed for the standard toy example (without the small mode around 5), see the right of Figure 6.

Remark 3.1 *The ABC-PMC algorithm has been run with the same sequence of tolerance levels*

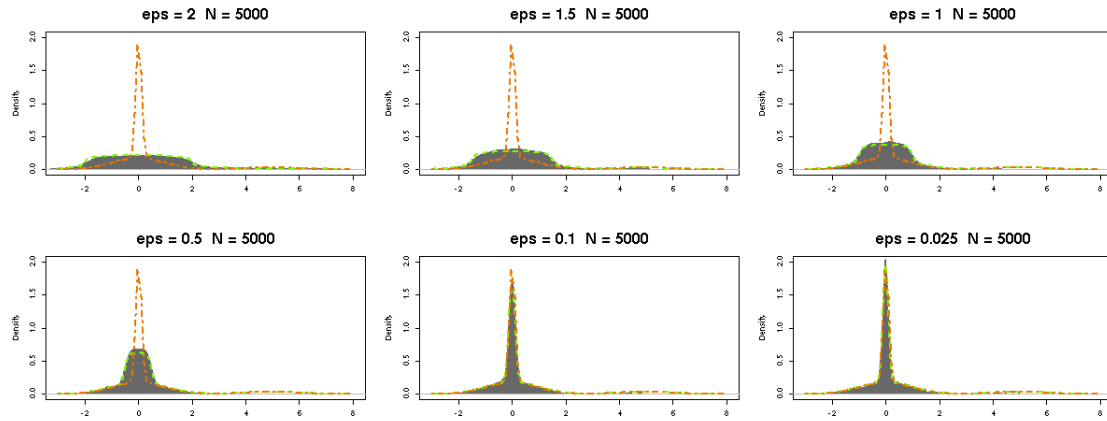


Figure 5: ABC-PMC for the modified toy example: the six populations associated with tolerance levels of 2, 1.5, 1, 0.5, 0.1 and 0.025 (5000 samples). The orange dotted curves stand for the exact posterior distribution, the green dotted curves for the target approximation, and the weighted populations are represented in gray.

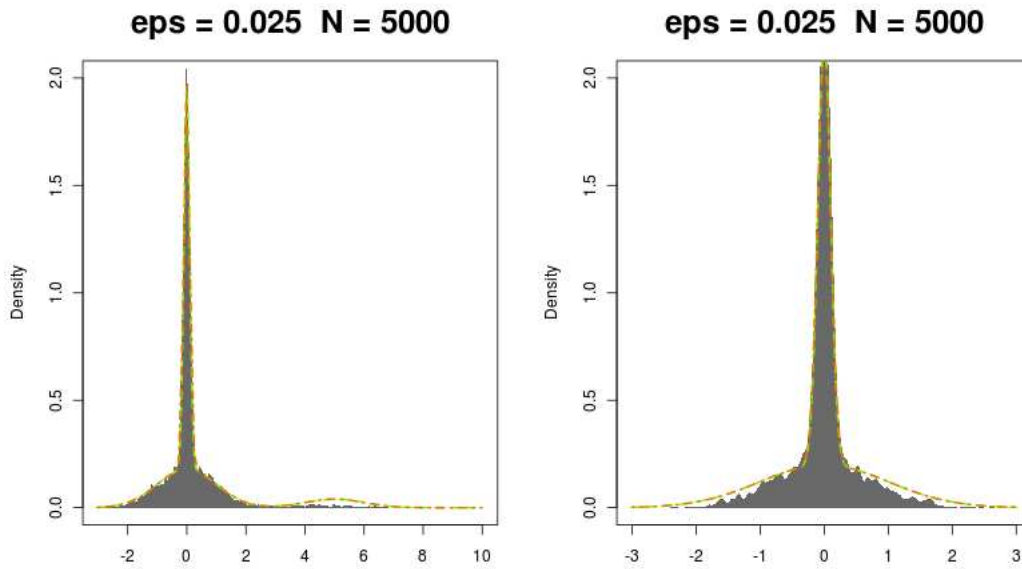


Figure 6: ABC-PMC zoom: populations associated with a tolerance level of 0.025 (5000 samples), for the modified toy example on the left, and for the standard toy example on the right. The orange dotted curves stand for the exact posterior distributions, the green dotted curves for the target approximation, and the weighted populations are represented in gray.

than those used for the ABC-PT: 15 populations were generated, hence it took more time than a run in which only a sequence of six tolerance levels is used. It appeared that the 15th weighted population associated with $\varepsilon_1 = 0.025$ did not cover better the local mode around 5. Moreover the diversity of this population was decreased, as its Effective Sample Size was of 3858.295,

compared to 4164.55 for the 6th population when six populations were generated. Note that like the ABC-PT, the results of the ABC-PMC were reproducible.

3.3 Comparison with ABC-SMC

We also compare the results of the ABC-PT with those obtained with the ABC-SMC (Del Moral et al. [2012]). We take $\epsilon = 0.025$ and $\alpha = 0.95$. To run ABC-SMC and ABC-PMC with similar CPU times, we first take $M = 1$ and 15 000 particles, see the left of Figure 7. The Effective Sample Size is of 11242. We also run the ABC-SMC for $M = 10$ and 20 000 particles, see the right of 7. The Effective Sample Size is of 14666.98. Generally, it appears that the results of the ABC-SMC are reproducible and similar to those obtained by the ABC-PT. However, we can note two differences: using the ABC-SMC the global mode around 0 is slightly under-covered, and the ABC-SMC has the advantage to give non-correlated samples.

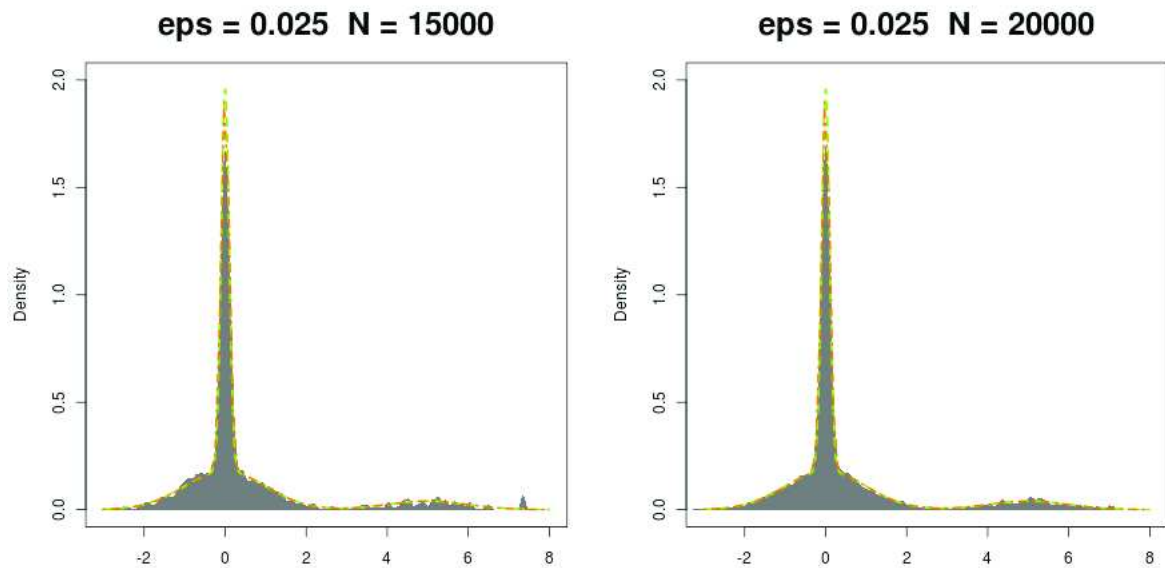


Figure 7: ABC-SMC for the modified toy example: On the left the results for $M = 1$ and 15 000 particles, on the right the results for $M = 10$ and 20 000 particles. The orange dotted curves stand for the exact posterior distribution, the green dotted curves for the target approximation, and densities of obtained samples are represented in gray.

Remark 3.2 *In an ABC framework, all sequential methods require simulating a number of pseudo-observations proportional to the number of particles. But their computational costs differ concerning the calculation of some importance weights: this cost is quadratic in the number of particles for the algorithms of Beaumont et al. [2009] (ABC-PMC), Sisson et al. [2007] and Toni et al. [2009], while it is linear for the adaptive ABC-SMC of Del Moral et al. [2012] (see also Beskos et al. [2011]). Comparison with the ABC-MCMC or the ABC-PT is not easy. Indeed, MCMC methods give unweighted correlated samples, hence the use of the effective sample size is not relevant. To compare the sequential and MCMC approaches, comprehensive convergence studies of the MCMC methods in an ABC framework should be achieved, taking into account*

the number N of chains, the temperature spacings [see Atchadé et al., 2011], as well as the autocorrelations.

3.4 Comparison with N independent parallel chains

Finally we compare the results of the ABC-PT with those obtained with N independent parallel chains. The 15 independent chains are locally updated by an ABC-MCMC algorithm during 600 000 iterations with a burn-in period of 150 000 iterations, with $\varepsilon = 0.025$ and $\mathcal{N}(\theta^{(t-1)}, 0.15^2)$ as the proposal $q(\cdot | \theta^{(t-1)})$. Figure 8 shows on three runs that the results obtained for the N independent chains are not reproducible: the run represented on the left enables to detect the small mode around 5 (which is a little bit over-covered), the run on the middle detect the small mode which is quite under-covered, while the run on the right did not detect it at all (none of the 15 chains succeeded to visit this mode). The traces of each of the 15 chains are similar to those of the ABC-MCMC algorithm, see the left of Figure 4. On the opposite, it appears that thanks to the exchange moves, the ABC-PT always succeeds to visit the small mode, and to escape it.

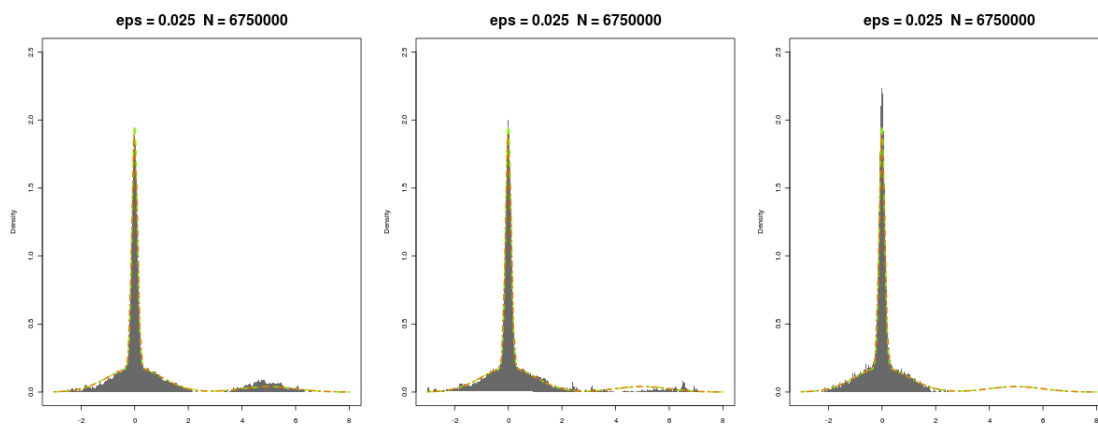


Figure 8: Modified toy example: three runs of 15 independent parallel chains.

3.5 Improving acceptance rates in ABC-PT

In order to increase the acceptance rate of the exchange moves in the ABC-PT algorithm, we consider a K -partition of the tolerance levels space, for a fixed integer K . In this way we obtain K disjoint subsets E_1, \dots, E_K , that we shall call *rings*. A chain of order i is associated to a ring E_j if $\rho(S(z_i), S(x)) \in E_j$. In the exchange move part of the ABC-PT algorithm, two chains are then chosen such that they are associated to the same ring. Thus, the second part of the algorithm is modified as follows:

(b) EXCHANGE MOVES WITH RINGS: N exchange moves are proposed:
 For each exchange move a ring with at least two associated chains is chosen randomly and two associated chains are chosen uniformly in this ring.
 For each of them: Denote by i and j ($i < j$) the orders of a chosen pair of chains. Exchange (z_i, θ_i) with (z_j, θ_j) with probability $\mathbb{1}_{\{\rho(S(z_j), S(x)) < \varepsilon_i\}}(z_j)$.

In practice, taking into account these rings, if an exchange between two selected chains i and j is proposed, then the acceptance probability $\mathbb{1}_{\{\rho(S(z_j), S(x)) < \epsilon_i\}}(z_j)$ is generally greater than in the case where the two chains are chosen uniformly (without rings). We illustrate this approach on the modified toy example. As previously we use $N = 15$, $T_N = 4$, $\epsilon_1 = 0.025$ and $\epsilon_N = 2$. We choose $K = 3$ rings containing the same number of tolerance levels. We get $E_1 = [0, 0.103]$, $E_2 =]0.103, 0.495]$, and $E_3 =]0.495, 2]$. The algorithm is run for 600 000 iterations with a burn-in period of 150 000 iterations. The results are given in Figure 9. Compared to the ABC-PT without rings, we observe that for an average of about +13% of CPU times, an increase of about +83% of accepted exchange moves is obtained. All the $N(N - 1)/2$ exchange rates between chains increase with the use of rings, and especially the rates between chains with high temperatures. Matrices of accepted exchange rates obtained on two runs of ABC-PT (with or without rings) are given in Appendix (Tables 7 and 8). An improvement in terms of autocorrelation is also observed as shown in Table 3. Note that this idea of partitioning can be related to the construction of rings of energy of the states space proposed in Kou et al. [2006].

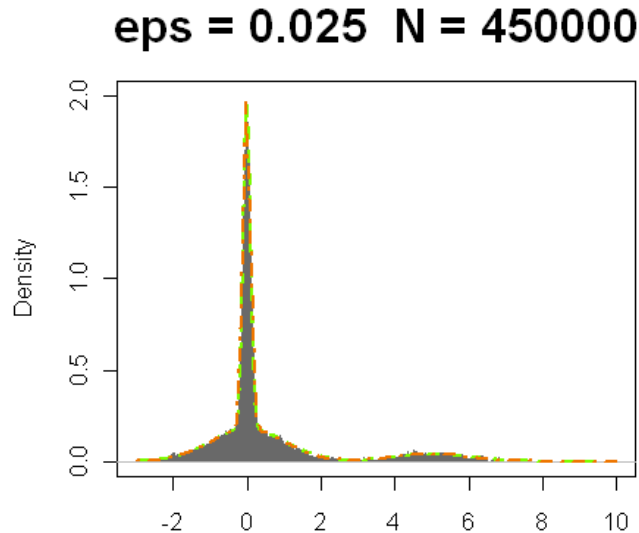


Figure 9: Approximations obtained by the first chain after an ABC-PT algorithm with $N = 15$ chains, with 3 rings of tolerance levels (600 000 iterations with a burn-in of 150 000). The orange dotted curves stand for the exact posterior distribution, the green dotted curves for the target approximation, and densities of obtained samples are represented in gray.

4 Illustration on tuberculosis data

We consider the real example of tuberculosis transmission studied by Tanaka et al. [2006], Sisson et al. [2007] and Blum [2010]. The data come from a study of a tuberculosis epidemic in San Francisco during 1991 and 1992 (Small et al. [1994]). They consist of DNA fingerprint at the

	thinning	Nb of samples	Order 1	Order 10	Order 20
ABC-PT (chain 1)	none	$45 \cdot 10^4$	0.622	0.264	0.238
	10	$45 \cdot 10^3$	0.268	0.148	0.111
	50	$9 \cdot 10^3$	0.193	0.070	0.048

Table 3: Modified toy example: autocorrelations of order 1, 10 and 20 of the first chain obtained by ABC-PT with 3 rings of tolerance levels (with or without thinning the chains).

IS6110 marker for 473 isolates, which are grouped in 326 distinct genotypes represented as follows:

$$30^1 23^1 15^1 10^1 8^1 5^2 4^4 3^{13} 2^{20} 1^{282},$$

where n^k indicates that k clusters of size n were observed.

A model of disease transmission and marker mutation is defined, which is an extension of a birth and death process: a birth corresponds to a new infection and a death corresponds to a host dead or healed. A system of stochastic differential equations is used, and is governed by three rates: α the birth rate per case per year, δ the death rate per case per year and θ the mutation rate per case per year (see Appendix C). However interest parameters for biologists are the transmission rate $\alpha - \delta$, the doubling time $\log(2)/(\alpha - \delta)$ and the reproductive value α/δ , results will then be given for these parameters in the following.

The likelihood of this model can not be written explicitly. However it is easy (although long) to simulate datasets from this model (see Tanaka et al. [2006] for details). This example is then particularly suited for the use of ABC methods. We use the same simulation process than in Tanaka et al. [2006].

To compare simulated and observed data, two summary statistics considered as relevant by the biologists are used: the number of distinct genotypes in the studied sample g , and the gene diversity H defined by $H = 1 - \sum_{i=1}^g (\frac{n_i}{n})^2$, where n_i is the size of cluster with genotype i and n is the size of the sample. The values g_{obs} and H_{obs} denote the two statistics obtained on the observed data. A simulated sample of size n associated to g and H is considered close to the observed data if $\frac{1}{n}|g_{obs} - g| + |H_{obs} - H| < \varepsilon$, where ε is a suitable tolerance level.

Prior distributions for the parameters α and δ are uniforms on the interval $[0, 5]$ with $\alpha > \delta$. Concerning the mutation rate θ , many studies have been conducted to estimate it (see Tanaka et al. [2006] for references), which enable us to take the following informative prior: $\theta \sim \mathcal{N}(0.198, 0.06735^2)$ truncated on the left at 0.

Concerning the ABC-PT algorithm, the vector of parameters is denoted by $\phi = (\alpha, \delta, \theta)$, and $\phi_i^{(t)}$ is the vector of the chain i at time t , for $i = 1, \dots, N$. As in Tanaka et al. [2006], the transition kernel for the first chain $q_1(\cdot | \phi_1^{(t-1)})$ corresponds to a $\mathcal{N}(\phi_1^{(t-1)}, \Sigma)$, with:

$$\Sigma = \begin{pmatrix} 0.5^2 & 0.225 & 0 \\ 0.225 & 0.5^2 & 0 \\ 0 & 0 & 0.015^2 \end{pmatrix}.$$

For the other chains, $q_i(\cdot | \phi_i^{(t-1)})$ corresponds to a $\mathcal{N}(\phi_i^{(t-1)}, \Sigma_i)$, with $\Sigma_i = \Sigma^{1/T_i}$. We took $N = 7$ chains, associated with temperatures between $T_1 = 1$ and $T_N = 2$ evenly spaced on a logarithmic scale. We used a Quad-Core Xeon E5320 1.86GHz processor with 8GB of RAM and we chose a minimum tolerance level of 0.01 to have reasonable acceptance rates for generated chains (to be compared with the tolerance level of 0.0025 taken by Tanaka et al. [2006]). The

sequence of tolerance levels was evenly spaced on a logarithmic scale as follows: $\varepsilon_1 = 0.01, \varepsilon_2 = 0.03, \dots, \varepsilon_7 = 0.1$.

Results

The algorithm was run for 20 000 iterations with a burn-in period of 2 000 iterations, it took about 4.5 days. The first chain was locally updated in 3.8 % of the iterations, and accepted on average 0.25 exchange move per iteration (5095 total moves). Concerning the set of $N = 7$ chains, on average 3.07 exchange moves were accepted per iteration (among 7). Table 4 gives the local acceptance rates for each chain, and proportions of accepted exchange moves for each pair of chain (among the proposed exchange moves).

	chain 1	chain 2	chain 3	chain 4	chain 5	chain 6	chain 7
chain 1	3.8	22.6	17.0	13.2	10.7	7.6	5.6
chain 2	.	13.9	76.1	58.8	44.5	34.0	25.2
chain 3	.	.	16.2	77.0	58.9	43.9	33.6
chain 4	.	.	.	17.9	77.3	56.3	45.0
chain 5	18.9	74.3	59.9
chain 6	19.9	78.3
chain 7	20.4

Table 4: ABC-PT for the tuberculosis example: the diagonal gives local acceptance rates, and outside the diagonal are given the proportions of accepted exchange moves.

Table 5 gives the posterior estimates of the interest parameters, obtained from the first chain. These estimates are of the same orders than those obtained by Tanaka et al. [2006], Sisson et al. [2007] and Blum [2010]. The mean transmission rate is around 0.60, describing the rate of increase of the number of cases in the population. The doubling time gives the same information, as it is just a transformation of the transmission rate. This parameter is used by biologists, as it can be easily interpreted (period of time required to double the percentages of cases in the population). Here it is estimated around 1.35 years. The credibility interval of the reproductive value (number of new infections generated by one case) is quite large, as we have positive posterior probabilities for few very small values of δ , corresponding to high values of this rate. The mean is then unstable. However, most of the posterior values are around the median (see Figure 10), which should then be preferred compared to the mean.

	Mean	Median	95% CI
Transmission rate	0.59	0.58	(0.29,0.92)
Doubling time	1.35	1.20	(0.76,2.41)
Reproductive rate	5.98	2.29	(1.20,17.35)
Mutation rate	0.25	0.25	(0.15,0.35)

Table 5: ABC-PT for the tuberculosis example: posterior estimates of the parameters of interest, on the 18 000 post-burn-in iterations of the first chain: mean, median and 95% credibility intervall.

We then compared the results of the ABC-PT algorithm with those of the standard ABC, ABC-MCMC and ABC-PMC algorithms. The same minimum tolerance level of 0.01 is used

for all the algorithms. Concerning standard ABC, 1000 samples were generated; concerning ABC-MCMC 350 000 iterations were performed with a burn-in period of 50 000 iterations; and concerning ABC-PMC 10 populations of 1000 samples were generated using the following sequence of tolerance levels (from Sisson et al. [2007]): $\varepsilon_1 = 1, \varepsilon_{10} = 0.01$ and $\varepsilon_k = 0.5(\varepsilon_{k-1} + \varepsilon_{10}), k = 2, \dots, 9$. Note that for the ABC-PMC algorithm we used weighted covariance matrices as suggested in Filippi et al. [2011] to get optimal kernels.

Posterior estimates of parameters obtained by these three algorithms were close to those of the ABC-PT, see Figure 10.

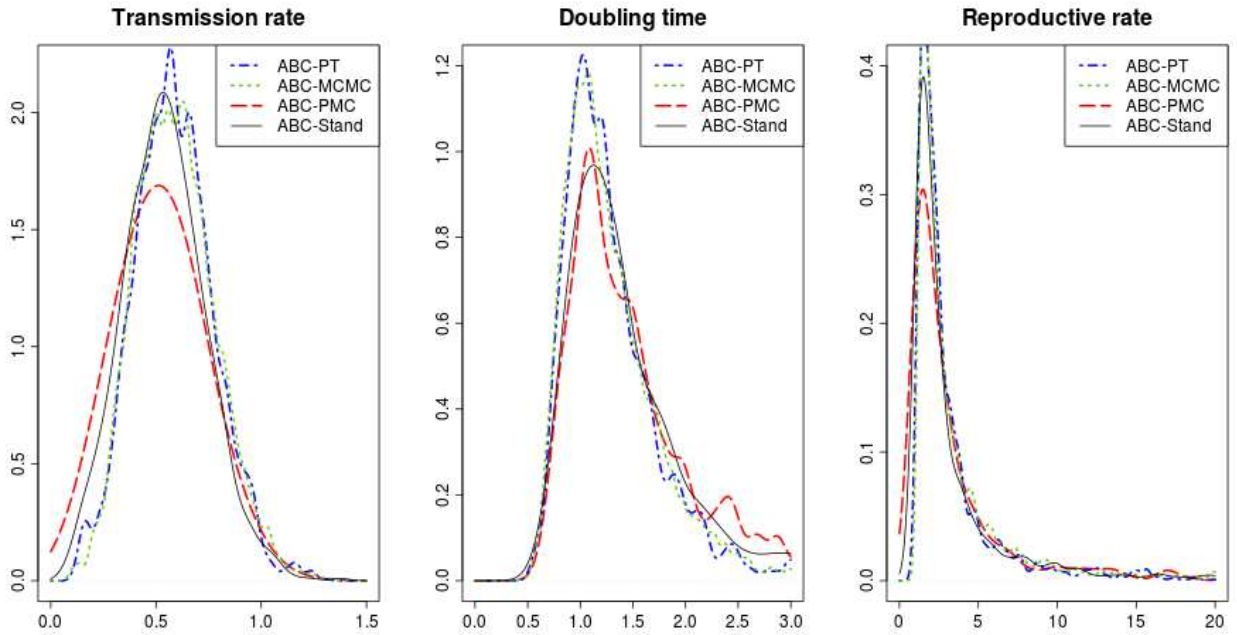


Figure 10: Posterior densities of the transmission rate, the doubling time, and the reproductive rate, obtained by the algorithms ABC-PT, ABC-MCMC, ABC-MCMC and standard ABC.

Concerning the standard ABC algorithm, an average number of 408 data generation steps were needed to obtain one sample. This algorithm took approximately 6.5 days.

Concerning the ABC-PMC, 203 data generation steps were necessary on average to obtain one particle and computing time was about 4.5 days.

Concerning the ABC-MCMC algorithm, 9 days were required. The chain was updated in 3.8 % of iterations. The obtained simulations are more correlated than those of the first chain of the ABC-PT algorithm. See Table 6, which gives for the α parameter the autocorrelations of the chain obtained by ABC-MCMC, and of the first chain obtained by ABC-PT, with or without thinning. Similar results were obtained for the parameters δ et θ . Hence an improvement of the ABC-PT over the ABC-MCMC in terms of autocorrelations was observed. If one is interested in having samples not too correlated, it appears from Table 6 that in order to have autocorrelations decreasing rapidly and non significant from the fifth order, every 500th sample must be kept for the ABC-MCMC algorithm, and every 15th sample must be kept for ABC-PT. As a consequence, 583 and 277 data generation steps are required on average to obtain one realization of the target posterior from the ABC-MCMC and the ABC-PT respectively. Figure

11 gives the autocorrelation plot for the ABC-PT when all the samples are kept, and when a thinning of 15 is used.

	thinning	Nb of samples	Order 1	Order 10	Order 20
ABC-MCMC	none	300 000	0.991	0.919	0.850
	100	3 000	0.513	0.107	0.032
	500	600	0.162	0.056	0.017
ABC-PT (chain 1)	none	18 000	0.812	0.456	0.296
	10	1 800	0.473	0.007	-0.040
	15	1 200	0.402	0.005	0.029

Table 6: For the α parameter: autocorrelations of order 1, 10 and 20 of the chain obtained by ABC-MCMC, and of the first chain obtained by ABC-PT (with or without thinning the chains).

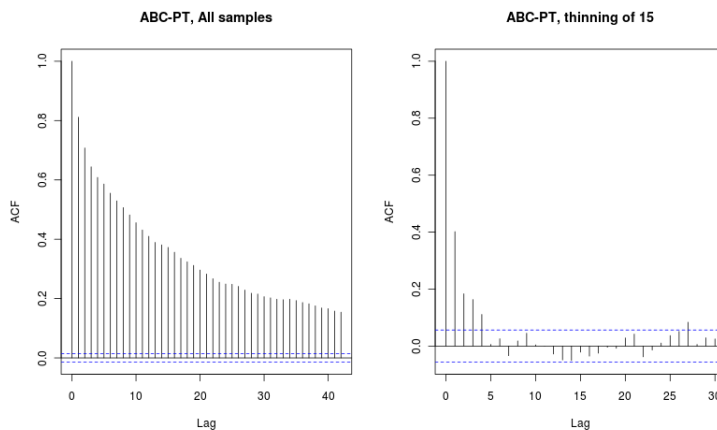


Figure 11: For the α parameter, autocorrelation plots of the first chain obtained by ABC-PT: on the left all samples are kept, on the right every 15th sample is kept.

Note that the obtention of correlated samples is part of any MCMC method, and satisfactory estimates can nevertheless be achieved. Hence, for the ABC-PT algorithm, if every 10th sample is kept for example, the algorithm is equivalent to the ABC-PMC algorithm in terms of number of required data generation steps to obtain one realization or one particle (203 vs 185).

5 Discussion

The ABC-PT algorithm developed in this paper is a new likelihood-free method based on the Markov Chain Monte Carlo theory. On the studied examples, results obtained by the algorithm are better than those obtained by the classical ABC-MCMC algorithm, in terms of exploration of the parameter space, and of autocorrelations of the generated chain. The performances of this new algorithm appear equivalent to those of the ABC-PMC and of the ABC-SMC, which are sequential methods. These sequential methods have the advantage to generate non-correlated samples. However, on the modified toy example the ABC-PMC appeared to slightly under-cover posterior distribution tails and the local mode, while the ABC-SMC appeared to slightly under-cover the global mode. The choice between ABC-PT, ABC-PMC and ABC-SMC should

then depend on the simulation’s objective. We would recommend to use the ABC-PT or the ABC-SMC if we suspect that the posterior of interest is multi-modal, with small modes in low probability areas. Concerning the computational cost of the ABC-PT, a convergence study should be achieved, taking into account the number N of chains, the temperature spacings [see Atchadé et al., 2011], as well as the autocorrelations.

From a theoretical point of view the N chains of the ABC-PT all together can be considered as a Markov chain, and the ABC-PT algorithm generates a Markov chain for (z, θ) with stationary distribution $\pi_{\varepsilon}^*(z, \theta | x)$. The chain obtained for ε_1 is the chain of interest, because it provides samples corresponding to $\pi_{\varepsilon_1}(z_1, \theta_1 | x)$ which is the target distribution. The marginal in θ_1 $\pi_{\varepsilon_1}(\theta_1 | x)$ can be considered as a good approximation of the posterior of interest $\pi(. | x)$ if ε_1 is small enough and S well chosen according to our data.

From a practical point of view, the ABC-PT requires calibration of the sequence of tolerance levels, and of the sequence of temperatures. Practical guidelines have been given in the modified toy example, see section 3.1. As seen in section 3.5, it is possible to improve the acceptance rate of the exchange moves by considering a partition of the tolerance levels space to choose the pairs of chains to exchange. The total number of accepted exchange moves was increased by about 83%, with only 13% of additional CPU time.

In this paper the ABC-PT algorithm only uses one type of global move, which is the exchange move. However, it is always possible to define other types of global moves, like those of the Evolutionary Monte Carlo of Liang and Wong [2001] for instance. In this case the moves should be performed in two steps to fulfill ABC requirements: proposal of new parameters, then updating of the associated simulated datasets. It is also possible to define delayed rejection moves, see for instance Green and Mira [2001] and Jasra et al. [2007]. It would be of particular interest to use the simulations generated by the other chains than the first chain, by developing a method in the same spirit than those of Beaumont et al. [2002] or Blum and François [2010]. Moreover, as proposed by a reviewer, we can take advantage of using a schedule of varying ε_1 for the first chain of interest, to improve its mixing. Finally, we think that using a pre-determined sequence of tolerance levels is not the best way to obtain a pre-determined number of samples with a tolerance level as small as possible, for a given computational time. In the lines of Müller et al. [2004] it would then be interesting to use a non pre-determined sequence of tolerance levels, and hence to have an adaptive algorithm.

References

- Y.F. Atchadé, G.O. Roberts, and .S. Rosenthal. Towards optimal scaling of Metropolis-coupled Markov chain Monte Carlo. *Statistics and Computing*, 21:555–568, 2011.
- M. Beaumont, W. Zhang, and D. Balding. Approximate Bayesian computation in population genetics. *Genetics*, 162:2025–2035, 2002.
- M. Beaumont, J.M. Cornuet, J. M. Marin, and C.P. Robert. Adaptive approximate Bayesian computation. *Biometrika*, 96(4):983–990, 2009.
- A. Beskos, D. Crisan, and A. Jasra. On the stability of sequential Monte Carlo methods in high dimensions. Technical report, Imperial College, London, 2011.
- M. Blum. Approximate Bayesian computational: a non-parametric perspective. *Journal of the American Statistical Association*, 491:1178–1187, 2010.

- M. Blum and O. François. Non-linear regression models for Approximate Bayesian computation. *Statistics and Computing*, 20(1):63–73, 2010.
- P. Bortot, S. Coles, and S. Sisson. Inference for stereological extremes. *Journal of the American Statistical Association*, 102:84–92, 2007.
- P. Del Moral, A. Doucet, and A. Jasra. An adaptive sequential Monte Carlo Method for approximate Bayesian computation. *Statistics and Computing*, to appear, 2012.
- C. Drovandi and A. Pettitt. Estimation of parameters for macroparasite population evolution using approximated Bayesian computation. *Biometrics*, 67:225–233, 2011.
- S. Filippi, C. Barnes, and M. Stumpf. On optimal kernels for ABC SMC. *arXiv:1106.6280,v2*, 2011.
- A.E. Gelfand and A.F.M. Smith. Sampling-based approaches to calculating marginal densities. *Journal of the American Statistical Association*, 85(410):398–409, 1990.
- C.J. Geyer. Markov chain Monte Carlo maximum likelihood. *Computing Science and Statistics: Proceedings of the 23rd Symposium on the interface*, pages 156–163, 1991.
- C.J. Geyer and E.A. Thompson. Annealing Markov chain Monte Carlo with applications to ancestral inference. *Journal of the American Statistical Association*, 90:909–920, 1995.
- P.J. Green and A. Mira. Delayed rejection in reversible jump Metropolis-Hastings. *Biometrika*, 88:1035–1053, 2001.
- A. Grelaud, J. M. Marin, C.P. Robert, F. Rodolphe, and F. Tally. Likelihood-free methods for model choice in Gibbs random fields. *Bayesian Analysis*, 3(2):427–442, 2009.
- W.K. Hastings. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 88:1035–1053, 1970.
- A. Jasra, D.A. Stephens, and C.C. Holmes. Population-based reversible jump Markov chain Monte Carlo. *Biometrika*, 94:787–807, 2007.
- S.C. Kou, Q. Zhou, and W.H. Wong. Equi-Energy sampler with application in statistical inference and statistical mechanics. *The Annals of Statistics*, 34(4):1581–1619, 2006.
- C. Leuenberger, D. Wegmann, and L. Excoffier. Bayesian computation and model selection in population genetics. *Genetics*, 184:243–252, 2010.
- F. Liang and W.H. Wong. Real-parameter evolutionary Monte Carlo with applications to Bayesian Mixture models. *Journal of the American Statistical Association*, 96:653–666, 2001.
- P. Marjoram, J. Molitor, V. Plagnol, and S. Tavaré. Markov chain Monte Carlo without likelihoods. *Proceedings of the National Academy of Sciences of the USA*, 100(26):15324–15328, 2003.
- T. McKinley, A. Cook, and R. Deardon. Inference in epidemic models without likelihoods. *The international Journal of Biostatistics*, 5(1), 2009.

- N. Metropolis, A.W. Rosenbluth, M.N. Rosenbluth, A.H. Teller, and E. Teller. Equations of state calculations by fast computing machines. *Journal of Chemical Physics*, 21(6):1087–1092, 1953.
- P. Müller, B. Sanso, and M. De Iorio. Optimal Bayesian design by inhomogeneous Markov chain simulation. *Journal of the American Statistical Association*, 99(467):788–798, 2004.
- K. Nagata and S. Watanabe. Asymptotic behavior of exchange ratio in exchange Monte Carlo method. *Neural Networks*, 21:980–988, 2008.
- R.M. Neal. Sampling from multimodal distributions using tempered transitions. *Statistics and computing*, 6(4):353–366, 1996.
- J. Pritchard, M. Seielstad, A. Perez-Lezaun, and M. Feldman. Population growth of human Y chromosomes: a study of Y chromosome microsatellites. *Molecular Biology and Evolution*, 16: 1791–1798, 1999.
- O. Ratmann, O. Jorgensen, T. Hinkley, M. Stumpf, S. Richardson, and C. Wiuf. Using likelihood-free inference to compare evolutionary dynamics of the protein networks of *H. pylori* and *P. falciparum*. *PLoS Computational Biology*, 3(11):2266–2278, 2007.
- S.A. Sisson, Y. Fan, and M.M. Tanaka. Sequential Monte Carlo without likelihood. *Proceedings of the National Academy of Sciences of the USA*, 104:1760–1765, 2007.
- S.A. Sisson, Y. Fan, and M.M. Tanaka. Sequential Monte Carlo without likelihood: Errata. *Proceedings of the National Academy of Sciences of the USA*, 106, 2009.
- P.M. Small, P.C. Hopewell, S.P. Singh, A. Paz, J. Parsonnet, D.C. Ruston, G.F. Schecter, C.L. Daley, and G.K. Schoolnik. The epidemiology of tuberculosis in San Francisco: a population-based study using conventional and molecular methods. *The New England Journal of Medicine*, 330:1703–1709, 1994.
- M.M. Tanaka, A.R. Francis, F. Luciani, and S.A. Sisson. Using approximate Bayesian computation to estimate tuberculosis transmission parameters from genotype data. *Genetics*, 173: 1511–1520, 2006.
- S. Tavaré, D. Balding, R. Griffith, and P. Donnelly. Inferring coalescence times from DNA sequence data. *Genetics*, 145:505–518, 1997.
- T. Toni, D. Welch, N. Strelkowa, A. Ipsen, and M. Stumpf. Approximate Bayesian computation scheme for parameter inference and model selection in dynamical systems. *Journal of the Royal Society Interface*, 6:187–202, 2009.

A Standard ABC results

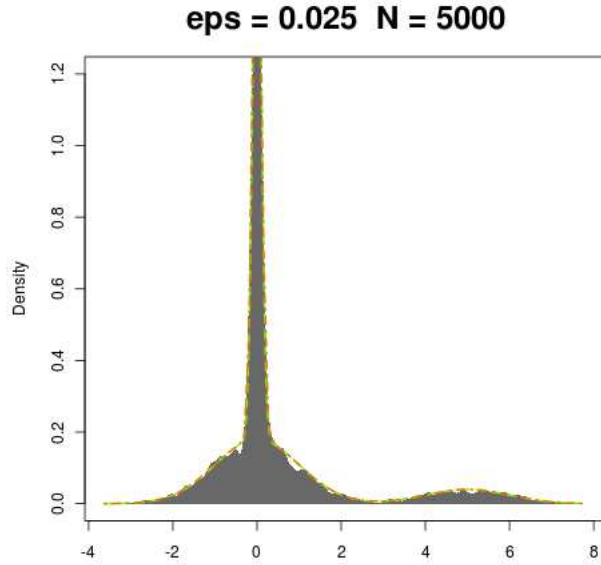


Figure 12: Modified Toy Example: population associated with a tolerance level of 0.025 (5000 samples) obtained with a standard ABC. The orange dotted curve stands for the exact posterior distributions, the green dotted curve for the target approximation, and the weighted population are represented in gray.

B Matrices of exchange rates for ABC-PT with and without rings

Tables 7 and 8 compare accepted exchange rates between chains based on the modified toy example, using ABC-PT algorithms with or without rings. The rate between chains i and j is equal to the number of accepted exchange moves between these chains divided by the total number of iterations. It is observed that the use of rings allows more accepted exchange moves, especially between high tempered chains.

	2	3	4	5	6	7	8	9	10	11	12	13	14	15
chain 1	0.16	0.12	0.09	0.06	0.04	0.03	0.02	0.01	0.01	0.01	0.01	0.00	0.00	0.00
chain 2	.	0.16	0.12	0.09	0.06	0.04	0.03	0.02	0.01	0.01	0.01	0.01	0.01	0.00
chain 3	.	.	0.16	0.12	0.08	0.056	0.04	0.03	0.02	0.01	0.01	0.01	0.01	0.00
chain 4	.	.	.	0.16	0.11	0.077	0.05	0.04	0.03	0.02	0.01	0.01	0.01	0.01
chain 5	0.15	0.10	0.07	0.05	0.04	0.03	0.02	0.02	0.01	0.01
chain 6	0.11	0.08	0.05	0.04	0.03	0.02	0.02	0.01	0.01
chain 7	0.10	0.07	0.06	0.04	0.03	0.02	0.01	0.01
chain 8	0.15	0.11	0.08	0.05	0.04	0.02	0.02
chain 9	0.23	0.16	0.11	0.07	0.05	0.04
chain 10	0.30	0.19	0.13	0.09	0.06
chain 11	0.22	0.15	0.11	0.08
chain 12	0.24	0.18	0.14
chain 13	0.43	0.34
chain 14	0.77

Table 7: ABC-PT with 3 rings for the modified toy example: the diagonal gives local acceptance rates, and outside the diagonal are given the proportions of accepted exchange moves.

	2	3	4	5	6	7	8	9	10	11	12	13	14	15
chain 1	0.10	0.08	0.06	0.04	0.03	0.02	0.02	0.01	0.01	0.01	0.00	0.00	0.00	0.00
chain 2	.	0.10	0.08	0.06	0.04	0.03	0.02	0.02	0.01	0.01	0.01	0.00	0.00	0.00
chain 3	.	.	0.11	0.08	0.06	0.04	0.03	0.02	0.02	0.01	0.01	0.01	0.00	0.00
chain 4	.	.	.	0.10	0.08	0.06	0.04	0.03	0.02	0.02	0.01	0.01	0.01	0.00
chain 5	0.10	0.08	0.05	0.04	0.03	0.02	0.02	0.01	0.01	0.01
chain 6	0.10	0.08	0.06	0.04	0.03	0.02	0.02	0.01	0.01
chain 7	0.10	0.08	0.06	0.04	0.03	0.02	0.01	0.02
chain 8	0.10	0.08	0.06	0.04	0.03	0.02	0.02
chain 9	0.10	0.08	0.06	0.04	0.03	0.02
chain 10	0.10	0.08	0.06	0.04	0.03
chain 11	0.22	0.15	0.06	0.04
chain 12	0.11	0.08	0.05
chain 13	0.10	0.07
chain 14	0.10

Table 8: ABC-PT for the modified toy example: the diagonal gives local acceptance rates, and outside the diagonal are given the proportions of accepted exchange moves.

C Formula for the Tuberculosis example

We used the same notations than Tanaka et al. [2006]. The number of cases of genotype i at time t is denoted by $X_i(t)$, $G(t)$ is the number of distinct genotypes that have existed in the population up to and including time t , and $N(t)$ is the total number of cases at time t .

$$N(t) = \sum_{i=1}^{G(t)} X_i(t).$$

The genotypes are labeled $1, 2, 3, \dots$ for convenience, although the ordering has no meaning, except that $i = 1$ represents the parental type from which others are descended (directly or indirectly). The three rates of the system are the birth rate per case per year α , the death rate per case per year δ , and the mutation rate per case per year θ . Tanaka et al. [2006] define the following probabilities:

$$\begin{aligned} P_{i,x}(t) &= P(X_i(t) = x), \\ \bar{P}_n(t) &= P(N(t) = n), \\ \tilde{P}_g(t) &= P(G(t) = g). \end{aligned}$$

The time evolution of $P_{i,x}(t)$ is described by the following differential equations:

$$\frac{dP_{i,x}(t)}{dt} = \underbrace{-(\alpha + \delta + \theta)xP_{i,x}(t)}_{\text{no event}} + \underbrace{\alpha(x-1)P_{i,x-1}(t)}_{\text{birth}} + \underbrace{(\delta + \theta)(x+1)P_{i,x+1}(t)}_{\text{death or mutation}}, \quad (4)$$

for $x = 1, 2, 3, \dots$

$$\frac{dP_{i,0}(t)}{dt} = (\delta + \theta)P_{i,1}(t). \quad (5)$$

Initially there is only one copy of the ancestral genotype, hence the initial conditions are: $P_{i,x}(0) = 0$ for all (i, x) , except $P_{1,1}(0) = 1$, and for $i = 2, 3, 4, \dots$, $P_{i,0}(0) = 1$.

To take into account the creation of new genotypes, the probability $\tilde{P}_g(t)$ is described by the following differential equations (only a mutation can create a new genotype):

$$\frac{d\tilde{P}_g(t)}{dt} = \underbrace{-\theta N(t)\tilde{P}_g(t)}_{\text{no mutation}} + \underbrace{\theta N(t)\tilde{P}_{g-1}(t)}_{\text{mutation}}, \quad \text{for } g = 2, 3, 4, \dots \quad (6)$$

$$\frac{d\tilde{P}_1(t)}{dt} = -\theta N(t)\tilde{P}_1(t). \quad (7)$$

The initial condition is $G(0) = 1$. Let t_g be the time when a new genotype g is created, we have $P_{g,1}(t_g) = P(X_g(t_g) = 1) = 1$, and $P_{g,x}(t_g) = P(X_g(t_g) = x) = 0$ for $x \neq 1$.

The total number of cases $N(t)$ is described by the following differential equations (only a birth or a death influence changes in this number):

$$\frac{d\bar{P}_n(t)}{dt} = \underbrace{-(\alpha + \delta)n\bar{P}_n(t)}_{\text{no birth, no death}} + \underbrace{\alpha(n-1)\bar{P}_{n-1}(t)}_{\text{birth}} + \underbrace{\delta(n+1)\bar{P}_{n+1}(t)}_{\text{death}}, \quad (8)$$

for $n = 1, 2, 3, \dots$

$$\frac{d\bar{P}_0(t)}{dt} = \delta\bar{P}_1(t). \quad (9)$$

The initial conditions are $\bar{P}_1(0) = 1$ and $\bar{P}_n(0) = 0$ for $n \neq 1$.