

Fast and Accurate Detection of Multiple Quantitative Trait Loci

CARL NETTELBLAD, BEHRANG MAHJANI,* and SVERKER HOLMGREN

ABSTRACT

We present a new computational scheme that enables efficient and reliable quantitative trait loci (QTL) scans for experimental populations. Using a standard brute-force exhaustive search effectively prohibits accurate QTL scans involving more than two loci to be performed in practice, at least if permutation testing is used to determine significance. Some more elaborate global optimization approaches, for example, DIRECT have been adopted earlier to QTL search problems. Dramatic speedups have been reported for high-dimensional scans. However, since a heuristic termination criterion must be used in these types of algorithms, the accuracy of the optimization process cannot be guaranteed. Indeed, earlier results show that a small bias in the significance thresholds is sometimes introduced.

Our new optimization scheme, PruneDIRECT, is based on an analysis leading to a computable (Lipschitz) bound on the slope of a transformed objective function. The bound is derived for both infinite- and finite-size populations. Introducing a Lipschitz bound in DIRECT leads to an algorithm related to classical Lipschitz optimization. Regions in the search space can be permanently excluded (pruned) during the optimization process. Heuristic termination criteria can thus be avoided. Hence, PruneDIRECT has a well-defined error bound and can in practice be guaranteed to be equivalent to a corresponding exhaustive search. We present simulation results that show that for simultaneous mapping of three QTLS using permutation testing, PruneDIRECT is typically more than 50 times faster than exhaustive search. The speedup is higher for stronger QTL. This could be used to quickly detect strong candidate eQTL networks.

Key words: algorithms, branch-and-bound, genetic mapping, genomics, statistical models, statistics.

1. INTRODUCTION

THE RAPID DEVELOPMENT OF OFF-THE-SHELF TECHNOLOGY for molecular genetics has resulted in that dense genetic maps and the corresponding genotype information can be provided in a much easier and cheaper way than before. This development opens new possibilities for analysis of quantitative traits, that is, traits that exhibit a continuous phenotype distribution. Since most important traits in humans, animals, and

Division of Scientific Computing, Department of Information Technology, Uppsala University, Uppsala, Sweden.

plants can indeed be seen as quantitative and affected both by the genetic composition and the environment, genetic mapping of them represents both a major opportunity and a challenge for modern genetics.

The underlying genetic architecture of a quantitative trait can be described by identifying a set of *quantitative trait loci* (QTL) in the genome for a population and attributing effect values to these loci using a suitable statistical model framework. The standard approach for locating a QTL is based on *interval mapping* (IM) (Lander and Botstein, 1989). Evaluating the standard IM model at a given position in the genome involves solving a maximum-likelihood problem based on genotype and phenotype frequencies for the population studied. In a *QTL search*, the evaluation of the statistical model is repeated for a large set of candidate positions in the genome to determine the QTL locations that results in the best model fit. Mathematically, this corresponds to solving a global optimization problem using some optimization scheme.

The result of a QTL mapping procedure is normally useful only if a proper significance threshold can be derived. Already when searching for a single QTL, traditional χ^2 approximations have been shown to have a significant bias (Carbonell et al., 1992). Therefore, randomization testing is frequently used (Churchill and Doerge, 1994), where many permuted datasets and the corresponding QTL searches are employed to empirically derive the distribution of the optimal model fit under the null hypothesis of no QTL being present.

In general, it can be assumed that multiple interacting QTL (epistatic interactions) should be included in a model to fully describe the genetic effect on a trait (Doerge, 2002). However, using a d -QTL model with general interactions, a d -dimensional global optimization problem has to be solved for each QTL search. For QTL models, the optimization landscape is often varied with many local minima, and still today the standard approach for QTL mapping problems is to use a brute-force exhaustive search over a dense lattice covering the search space. For multidimensional searches, this approach rapidly becomes computationally intractable. Beyond $d = 2$, QTL mapping employing true multidimensional optimization has, as far as we know, not been used in practice. This means that geneticists have so far not drawn any firm conclusions on how important epistatic interactions between more than two QTL are for describing quantitative traits. However, there are indications that such interactions may indeed be important (see, e.g., Carlborg and Haley, 2004).

A main reason for the inconclusive situation regarding the importance of more complex epistatic interactions is the lack of efficient and reliable computational tools for performing multidimensional QTL mapping, as well as determining the joint significance of the set of QTL found. In this context, it should also be noted that even under the assumption that the QTL do not interact (i.e., the true population effects are perfectly additive), the estimated QTL effects will be more accurately estimated if all putative loci are included in a single, multilocus model. Hence an efficient computational tool also can be very useful for such settings. If a single-locus model is used repeatedly, correlations between loci could inflate or distort the detected effects. This is especially true for linked loci residing on the same chromosome, but also holds due to random correlations for loci located on different chromosomes.

Some early examples of simultaneous mapping of two QTL are found in Carlborg et al., (2000), where a genetic optimization algorithm is used for solving the optimization problem. Today, performing two-dimensional QTL mapping is regarded as a standard procedure by many researchers in genetics. In Ljungberg et al. (2004), the deterministic optimization algorithm DIRECT (Jones et al., 1993) was introduced for solving QTL search problems and also for $d \geq 2$, but these results have so far not been used to perform high-dimensional QTL mapping experiments of relevance to genetics.

In this article, we present a computational scheme that enables an efficient and reliable solution to QTL mapping problems in experimental populations for high-dimensional general models of interacting QTL. Our scheme is based on an analysis of the behavior of the objective function (the results from a linear regression QTL model fit) and implemented in deterministic global optimization framework. We also show how the result of the analysis and our optimization framework can be used to set up permutation testing in a very efficient way to determine the relevant significance thresholds. Our algorithms are structured in such a way that the problem of determining the set of d QTL resulting in the optimal model fit is separated from the evaluation of the model of the QTL effects. This implies that models based on for example, both the linear regression approximation and the standard interval mapping maximum-likelihood model can easily be included in a production tool for genetic analysis. We focus on linear regression models in this article since they are much less computationally demanding and lend themselves to the type of transformations that are exploited in our analysis of the behavior of the objective function. In this context, it is important to

note that if complete genetic information is available, and general assumptions of normal distributions hold, the linear regression and maximum likelihood are equivalent (Haley and Knott, 1992).

2. LINEAR REGRESSION QTL MODELS

We consider QTL analysis for experimental populations with known relations between individuals, with the founder individuals demonstrating some origin-defining feature. This can be a matter of a known genetic relation (a set of common inbred or outbred lines) between the founders, or founders expressing a specific phenotype. Today, very dense marker maps are available, and we formulate the analysis assuming information on allele origin being available in any position tested.

If a model with a total of d QTL is used, then the corresponding d sites in the genome are assumed to represent the only genetic factors that contribute to the phenotype. One can split individuals based on genotype into 2^d classes in a backcross (3^d for an intercross), since each of the d sites can take two different values (three for an intercross, two homozygote genotypes, and one heterozygote, if allele parental origin is ignored). Within each of these classes the variance is entirely nongenetic. Another important assumption is that phenotypes are normally distributed with different means but identical variance in each class, that is, $y|g_1, g_2, \dots, g_d \sim N(\mu_{g_1, g_2, \dots, g_d}, \sigma^2)$. For details on different experimental cross structures, see, for example, Wu et al. (2007). For ease of presentation, we now consider the typical case in which we have an F_0 generation of individuals that can be considered to belong to either out of two lines, Q and q . Assuming loci act additively, one can model a relation between genotype and phenotype in a backcross based on these founders as:

$$y_i = \mu(x) + \sum_{j=1}^d a_j(x_j) Z_i(x_j) + e_i(x), \text{ where: } Z_i(x_j) = \begin{cases} 1, & \text{Qq at } x_j; \\ 0, & \text{qq at } x_j. \end{cases} \quad (1)$$

In this model, $x = (x_1, \dots, x_d)$ is a vector of d elements that defines the search space, spanning a hypercube where x_i ranges over the length of the genome. In practice, the search space volume can be slightly reduced by employing the symmetries, resulting from the fact that the ordering of the QTL within the model is irrelevant. The phenotypes of all individuals are denoted by $y = (y_1, \dots, y_n)$. e_i is normally distributed with mean 0 and variance σ^2 , μ is the reference effect, and a is the additive effect. The model (1) can be written in a matrix form:

$$y = \mathbf{A}(x)b(x) + e(x), \text{ where } b(x) = (\mu(x), a(x))^T. \quad (2)$$

The least-squares estimate of the QTL effects for this linear model is:

$$\hat{b} = (\mathbf{A}(x)^T \mathbf{A}(x))^{-1} (\mathbf{A}(x)^T y) \quad (3)$$

$$\hat{\sigma}^2 = \frac{1}{n} (y - \mathbf{A}(x)b(x))^T (y - \mathbf{A}(x)b(x)). \quad (4)$$

These QTL positions can now be found by minimizing the residual sum of squares over all x and b :

$$RSS_{opt} = \min_{x, b} (y - \mathbf{A}(x)b(x))^T (y - \mathbf{A}(x)b(x)). \quad (5)$$

The solution to this minimization problem can be separated into two parts: the inner, linear problem:

$$RSS_{opt} = \min_b (y - \mathbf{A}(x)b(x))^T (y - \mathbf{A}(x)b(x)), \quad (6)$$

and the outer, nonlinear problem:

$$RSS_{opt} = \min_x RSS(x). \quad (7)$$

Solving the minimization problems (5) for a multiple QTL mapping problem is computationally heavy since x is a d -dimensional vector, and the optimization landscape for the outer (global) optimization problem is in general quite complex. It is clear that an optimal (albeit not necessarily unique) solution to the QTL search problem always exists, but to determine if a result is statistically valid a significance threshold

has to be determined. If permutation testing is used for determining the genome-wide threshold, several hundreds or thousands of QTL mapping problems of the type (5) must then be solved for the permuted datasets.

3. A GLOBAL OPTIMIZATION ALGORITHM FOR QTL MAPPING

Our new computational scheme for solving the high-dimensional QTL mapping problem is based on a deterministic Lipschitz optimization approach implemented in the DIRECT (Jones et al., 1993) algorithmic framework. In this section, we first review the original DIRECT algorithm and then summarize how this scheme was adapted to solve multidimensional QTL search problems in Ljungberg et al. (2004). We then present the basic idea behind our new optimization scheme, named PruneDIRECT, as a background for the more detailed description and analysis in later sections.

3.1. *The original DIRECT algorithm*

In the DIRECT scheme, the search space is successively divided into progressively smaller boxes. The search effort is focused in the most promising regions, and the subdivision of less promising boxes is postponed. For a traditional exhaustive search over a d -dimensional hypercube, the objective function is evaluated in a brute-force fashion at every point of a fine lattice covering the search space, using no further information on the objective function. If DIRECT is run to completion using a minimum resolution criterion matching the step length in the exhaustive search lattice, it will eventually explore exactly the same points as the corresponding exhaustive search. The efficiency of the DIRECT algorithm comes from exploiting the property that the most promising regions are explored first. A heuristic termination criterion is then needed to stop the search well before it devolves into an exhaustive search. As the heuristic criterion has no firm mathematical foundation, this implies that there is no well-established guarantee that the result from DIRECT is equal to the result from an exhaustive search over the same space. This is a typical result for global optimization schemes aimed at solving general problems. To be able to provide guaranteed accuracy without exploring the full search space, more information on the properties of the objective functions must be provided and used in the optimization procedure.

The original DIRECT algorithm initially creates a single search box covering the full search space. The objective function, that is, in our case the RSS, is evaluated in the center of this box. The box is then split into three equally sized boxes along the majoring dimension. This trinary split results in the centroid of the original box coinciding with the centroid of one of the new boxes. Therefore, only two additional function evaluations are required for the three resulting boxes. DIRECT then continues iteratively splitting the boxes. At the end of each DIRECT iteration, the convex hull is determined among the remaining boxes, in a space of box radii vs. objective function values. This hull determines which boxes to split in the following iteration. By computing the hull, the RSS value for the sequence of boxes will be monotonously decreasing if tracing along the box radii. The idea is that a promising box is characterized by either being large, so there is a high possibility that exploring it further might recover a new optimum, or that the value at the centroid is close to the current optimum and hence a new optimum might be found, even if the radius is small. This qualitative argument can also be represented as different assumptions of the value of a Lipschitz constant K , that is, on the maximal slope of the objective function (Jones et al., 1993). Figure 1 illustrates a few iterations of DIRECT in a simple one-dimensional space. The hull is ‘peeled off’ when those boxes are split, making new boxes available in the next iteration. This process is repeated until a suitable termination condition has been reached.

3.2. *Adaption of DIRECT to solve QTL mapping problems*

In Ljungberg et al. (2004), the original DIRECT algorithm was adapted for solving QTL search problems. A heuristic termination criterion, as described above, was used to stop the search, and the resulting scheme was proven to be orders of magnitudes faster than the corresponding exhaustive search. This made QTL searches for at least $d = 3, 4,$ and 5 possible using a reasonable computational effort. However, the results in Ljungberg et al. (2004) also show that using DIRECT for the permutation tests can result in a bias in the significance thresholds compared to those from an exhaustive search. The reason for this is that the

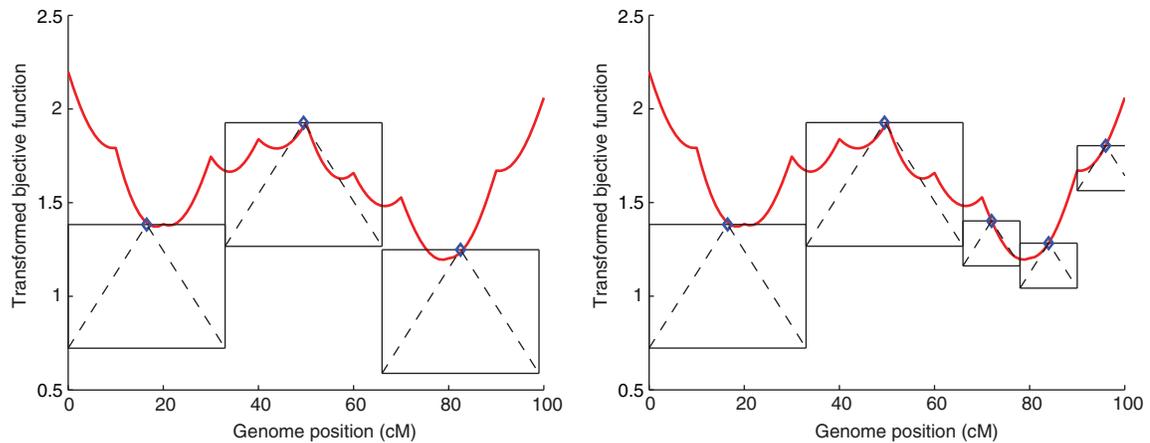


FIG. 1. To the left, three boxes in the optimization space are illustrated. As all boxes are of equal size, only one will be selected for splitting in the convex hull, resulting in the boxes on the right. If the splitting continues, two boxes would be split, one from each size, as the smallest function value in the smaller box size is slightly lower than the smallest value in the larger one. Dashed lines indicate possible minimum function values found at each distance from the box centroid, assuming a strict Lipschitz bound of $K = 0.04$.

termination criterion used in Ljungberg et al. (2004) results in premature termination for the very flat optimization landscape present in most permuted cases. The end result is that a putative 95% threshold instead would give, for example, 94.8% significance. It should also be noted that no bias was detected in the DIRECT searches for nonpermuted phenotype data, probably due to the more structured nature of these optimization landscapes. A main result of the analysis presented in later sections is that it is possible to terminate our new PruneDIRECT process at a much earlier point than the corresponding exhaustive search while still being able to *guarantee*, up to some defined residual probability, that the same global optimum is found as for the exhaustive search.

We argue that the successful results for DIRECT presented in Ljungberg et al. (2004) originate from the fact that the probability of co-inheritance of two genomic loci on a single chromosome is related to their physical distance from each other. The standard unit for genetic distances, mapping distance, is even directly defined from such probabilities. In Ljungberg et al. (2004), this fact is not explicitly used in the algorithm, but since the end result is that the slope of the objective function is limited, the performance of DIRECT ends up being quite good. It is also clear that this notion of co-inheritance cannot be extended across chromosome boundaries. Separate chromosomes segregate independently and can be considered to have an infinite mapping distance. In the search space for the QTL mapping problems, all genome locations are lined up as a single line, one chromosome after another. There is a clear disjunction at the chromosome borders, corresponding to a discontinuity of the RSS at these boundaries. In Ljungberg et al. (2004), this situation is handled by introducing what is essentially several DIRECT searches governed by a common priority queue. In this version of the algorithm, each chromosome combination is considered to correspond to a separate search space, called a chromosome combination box (cc-box). In the first step of the DIRECT scheme, the RSS is evaluated at the centroids of all these cc-boxes.

3.3. The PruneDIRECT algorithm

The original DIRECT algorithm is based on an assumption of Lipschitz continuity, where the value of the Lipschitz constant is unknown. This means that we know there exists a constant K such that no partial first-order derivative of the RSS will exceed K in any position but that the value on this bound of the slope of the objective function is not known.

Our improvement of the DIRECT procedure for QTL mapping is based on the fact that a Lipschitz bound for a QTL problem can be computed based on the relation between co-inheritance and physical distance mentioned above. Parts of the search space can then be permanently excluded from further subdivisions at each iteration in the DIRECT procedure, resulting in an efficient algorithm where a heuristic termination criterion does not have to be used. The resulting scheme has a well-defined error bound and is equivalent to the corresponding exhaustive search.

If a bound on the Lipschitz constant K is known, it is possible to compute upper and lower bounds for the objective function within any box in the search space, given knowledge of the value at the centroid. When using the original DIRECT algorithm, K is unknown, but it is still possible to impose a partial ordering of boxes. If a box A is both smaller and has a larger value of RSS than another box B , then no value of K can result in a smaller minimum bound within A than within B . This partial ordering gives rise to the choice of a convex hull of boxes to split in the next iteration.

However, DIRECT can also be considered in the more general context of Lipschitz optimization schemes. Here, the more traditional schemes (Shubert, 1972; Pinter, 1986; Galerpin, 1985; Mladineo, 1987) assume that a bound of the Lipschitz constant is known. By introducing an upper bound on K in DIRECT, the resulting algorithm is brought close to classic Lipschitz optimization. A known value of the objective function in small boxes will also enable exclusion of larger boxes, hereby introducing a pruning of the search tree, indicating the choice of name for our new algorithm. In contrast to the peeling effect of successive convex hulls in DIRECT, which will eventually result in every existing box getting split, box volumes pruned due to a Lipschitz criterion are permanently irrelevant for further evaluation of the RSS and can be removed from all data structures.

The natural termination condition for the PruneDIRECT algorithm is given by the finite resolution criterion corresponding to the lattice in the underlying exhaustive search. This results in an error bound of Kh on the value of the objective function, where h is the step length in the lattice. Thus, a bound on the Lipschitz constant K leads to the new PruneDIRECT algorithm having improved performance compared to exhaustive search (by excluding parts of the search space) and a well-defined error bound. Since no heuristic termination criterion is used anymore, a main new result is also that PruneDIRECT can be guaranteed to be equivalent to the corresponding exhaustive search, effectively removing any bias of the results caused by the optimization procedure.

For multidimensional QTL mapping that explores the full search space, it is possible to directly use the permutation test methodology for single-QTL models. However, performing hundreds or thousands of multidimensional QTL searches for permuted data in order to get a significance threshold can be a very computationally demanding task, even when an efficient global optimization scheme is employed. Here, it should be noted that since the purpose of doing the random permutations is to determine a significance threshold empirically, it is not necessary to locate the location of the best model fit for every set of randomized data. Rather, it is enough to answer the yes/no question if it is possible to find a location in the permuted dataset with a residual variance below that of the putative set of QTL. If the optimum value used to determine pruning is replaced by the value of the QTL candidate tested, rather than the optimum found in the current permute search, significant decreases in computational effort are possible, even compared to using the same algorithm to solve the full QTL search problem for each permuted dataset. The significance levels derived are identical. We use this approach to derive a very fast scheme for performing permutation tests, described in more detail later.

4. PRESENTING THE LOGVAR OBJECTIVE FUNCTION AND ITS LIPSCHITZ BOUND

In this section, the behavior of a transformed QTL search objective function at, and in the vicinity of, a QTL is considered. For a QTL search, the explainable genetic variance can be considered a natural objective function since the position with minimum residual variance can be defined as the location of a putative QTL. The benefit of using our transformed objective function, which we call LogVar, is the possibility to calculate a bound for its Lipschitz constant. The derivation of the bound is based on calculating the explainable variance as a function of genetic distance from the true QTL. We start by presenting the transformation and deriving the bound for an infinite size population and then move to the more realistic situation where QTL mapping for a finite size population is performed.

4.1. Infinite size population

Consider an infinite size population. The total variance is the sum of genetic variance and environmental variance. If there are d QTL, then all genetic variance is explainable by them. We start by analyzing $d = 1$ and then show how these results can be generalized to higher dimensions. Assume there is a QTL at

position x_0 . The total variance is the ‘‘mean squared error.’’ We split the total variance into a sum of undiscovered genetic variance and the discovered genetic variance at position $x + x_0$:

$$V_{total} = V_{genetic} + V_{enviromental} = V_{g-discovered}(x+x_0) + V_{g-undiscovered}(x+x_0) + V_e, \quad (8)$$

which can be abbreviated into $V_t = V_{gd}(x+x_0) + V_{gu}(x+x_0) + V_e$.

Then define $V_r(x+x_0) = V_{gu}(x+x_0) + V_e$ as the unexplainable variance at position $x+x_0$. The goal is to express $V_r(x+x_0)$ as a function of the recombination frequency and then find a bound for it. For simplicity, we start by letting $V_e = 0$ and later we add V_e back to the calculation.

Assume that the phenotype values for the two QTL genotypes are 0 and 1. Due to the symmetric structure of individuals with phenotype value 0 and individuals with phenotype value 1, $V_r(x+x_0)$ for each class is equal to the total residual variance. Hence:

$$V_r(x+x_0) = V_r(x+x_0 \text{ for class 1}) = \text{Prob}(\text{no recombination}) \times (1 - \mu(x+x_0 \text{ for class 1}))^2 + \text{Prob}(\text{recombination}) \times (0 - \mu(x+x_0 \text{ for class 1}))^2 \quad (9)$$

where:

$$\begin{aligned} \mu(x+x_0 \text{ for class 1}) &= 1 \times \text{Prob}(\text{No recombination}) + 0 \times \text{Prob}(\text{recombination}) \\ &= p(x+x_0). \end{aligned} \quad (10)$$

Hence, from substituting (10) in (9) and denoting the recombination frequency at position $x+x_0$ by $p(x+x_0)$ we get:

$$\begin{aligned} V_r(x+x_0) &= (1-p(x+x_0))(1-p(x+x_0))^2 + p(x+x_0)(0-p(x+x_0))^2 \\ &= p(x+x_0) - p(x+x_0)^2. \end{aligned} \quad (11)$$

This result for V_r should be related to the total phenotypic variance, which is the genetic variance at position x_0 since we defined $V_e = 0$. We know that all the variance at a QTL point is the discovered variance, hence:

$$V_t = V_g(x_0) = V_{gd}(x_0) + v_{gu}(x_0) = [0.5 \times (1-0.5)^2 + 0.5 \times (0-0.5)^2] + 0 = 0.25. \quad (12)$$

At last, we need to have a recombination map, relating genetic distance to the recombination frequency. We use Haldane’s mapping function (Haldane, 1919):

$$p(x+x_0) = 0.5 + 0.5e^{-\frac{2x}{100}}, \quad (13)$$

where x is the genetic distance from a fixed point x_0 measured in centimorgan. Inserting $p(x+x_0)$ from (13) into equation (11) we get:

$$V_r(x+x_0) = (0.5 + 0.5e^{-\frac{2x}{100}})(1 - (0.5 + 0.5e^{-\frac{2x}{100}})) \quad (14)$$

$$= 0.25 - 0.25e^{-\frac{4x}{100}} \quad (15)$$

$$= V_t - V_g(x_0)e^{-\frac{4x}{100}} \text{ (from 12)}. \quad (16)$$

Here, V_e can be added to the final calculation since variance is additive. Alternatively, it is safe to assume that V_e is already included in calculating V_r . Now that we have the formula for calculating the explainable variance as a function of genetic distance, we can introduce a transformation of the objective function in the DIRECT optimization. Ignoring a scaling factor, the RSS is equivalent to the residual variance V_r . Adding a constant will not affect the location of minima, so we can instead consider $g(x) = V_r(x) - V_t = -V_g e^{-\frac{4|x|}{100}}$ based on (16), assuming $x_0 = 0$. Furthermore, the function $g(x)$ is always negative, so $f(x) = -\ln(-g(x))$ will always be defined and share the locations of minima with $g(x)$. Hence, we can introduce a transformed objective function; call it LogVar objective function:

$$f(x) = -\ln(-V_r(x) + V_t) = -\ln\left(V_g e^{-\frac{4|x|}{100}}\right). \quad (17)$$

Here, it is easy to verify the derivative of the transformed objective function is bounded:

$$\frac{df(x)}{dx} = \frac{d}{dx} - \ln \left(V_g e^{-\frac{4|x|}{100}} \right) = \pm 0.04. \quad (18)$$

The use of $|x|$ in the definition of $g(x)$ is related to the fact that x is defined as the distance from the QTL, while a position in the chromosome can naturally be both upstream and downstream from this position. It is possible to shift the function by introducing the true QTL position y , resulting in $f(x) = \ln V_g - 0.04|x - y|$.

The result above can be further generalized, maintaining the bound on K . First, we have assumed that all genetic variance was attributable to a single locus (at y). We can now assume a single-locus model for analysis, but that the true QTL, with respective components of V_g are represented as a vector $\mathbf{y} = \{y_1, y_2, \dots, y_n\}$, resulting in the following expression for the residual variance (assuming all y_i being unlinked):

$$V_r(x) = V_t - \sum_{i=1}^n V_{g_i} e^{-\frac{4|x-y_i|}{100}}. \quad (19)$$

If all QTL are indeed unlinked, the positions y_i relative to any point considered in a single-QTL model will be $\pm\infty$, except for at most one y_j . Since linkage is transitive, the observation position x can at most be linked to a single QTL. Thus, (19) reduces to:

$$V_r(x) = V_t - V_{g_j} e^{-\frac{4|x-y_j|}{100}}, \quad (20)$$

and the derivative bound on (20) follows from the result in (18).

The next extension is to make the search landscape itself multidimensional, replacing the scalar x with a vector $\mathbf{x} = \{x_1, x_2, \dots, x_m\}$. Assume that each x_i is defined with a point of reference in linkage with the corresponding QTL z_i . Modeled locations not in linkage with any true QTL will result in no explainable variance, and therefore do not need to be considered. As the numbering of both vectors is essentially arbitrary, all other cases are also symmetrical to this one. Using the earlier result, each x_i will only have a term for the corresponding z_i , as all other mapping distances $|x_i - z_j|$ would be $+\infty$. This results in:

$$V_r(\mathbf{x}) = V_t - \sum_{i=1}^n V_{g_i} e^{-\frac{4|x_i-z_i|}{100}}. \quad (21)$$

We then have:

$$f(\mathbf{x}) = -\ln(-V_r(\mathbf{x}) + V_t), \quad (22)$$

where the logarithm does not affect the location of minima.

Equation (21) can be rewritten as:

$$V_r(\mathbf{x}) = V_t - V_{g_j} e^{-\frac{4|x_j-z_j|}{100}} - \underbrace{\sum_{i=1, i \neq j}^n V_{g_i} e^{-\frac{4|x_i-z_i|}{100}}}_C, \quad (23)$$

for any arbitrary $j \in N, j \leq n$. Based on (23), (22) becomes:

$$f(\mathbf{x}) = -\ln \left(V_{g_j} e^{-\frac{4|x_j-z_j|}{100}} + C \right). \quad (24)$$

From basic calculus, we know that $\left| \frac{d}{dx} \ln(e^{kx} + C) \right| = \left| \frac{ke^{kx}}{e^{kx} + C} \right| \leq \frac{d}{dx} \ln(e^{kx}) = \frac{ke^{kx}}{e^{kx}} = k$ for any $C \in R, C \geq 0$. Hence, all partial derivatives $\frac{\partial f}{\partial x_j}$ are confined by the bound given for the single-dimensional first derivative presented earlier. Thus, the LogVar objective function $f(\mathbf{x})$ as defined above will again have a well-defined Lipschitz bound.

4.2. Explainable variance as a function of genetic distance (finite size population)

The bound derived for an infinite size population does not apply directly to the derivative of the actual residual variance in experimental data, but to the expected value of the derivative, corresponding to the relation between the mapping distance and the expected number of crossover events. Depending on what

recombinations are actually present (i.e., in which individuals, with accompanying phenotype values), the actual residual variance can, and will, be different from the relationship predicted. This is an effect of sampling, which should decrease with an increasing size in population and vanish at a theoretical infinite population size. However, experimental populations tend to be rather small, and thus the infinite-size approximation cannot be used directly.

In this section, we derive an approximation to the distribution of the residual sum of errors in the vicinity of any point in the search space for a backcross. We revisit the linear model for a single QTL. Assuming there is a QTL at point $x_0 = 0$, the model is presented as:

$$y_i = \mu(x) + a(x)Z_i(x) + e_i(x), \text{ where: } Z_i(x) = \begin{cases} 1, & \text{Qq at } x; \\ 0, & \text{qq at } x. \end{cases} \quad (25)$$

Our goal is to calculate the distribution of the residual sum of squares (RSS) at point x . At first, we introduce some definitions:

$$n = \text{total number of individuals} \quad (26)$$

$$n_0(x) = \text{number of 0's at } x \quad (27)$$

$$n_1(x) = \text{number of 1's at } x \quad (28)$$

$$m_{10}(x) = \text{number of 0's at } x \text{ that are 1 at zero} \quad (29)$$

$$m_{01}(x) = \text{number of 1's at } x \text{ that were 0 at zero} \quad (30)$$

$$M_{10}(x) = \text{set of indices of individuals that recombined from 1 at zero to 0 at } x \quad (31)$$

$$M_{01}(x) = \text{set of indices of individuals that recombined from 0 at zero to 1 at } x \quad (32)$$

$$p = \text{recombination frequency} \quad (33)$$

$$S_0 = S_1 = \text{within-class environmental variance} = 1 \quad (34)$$

The last definitions are simplifying assumptions relating to the normal residual assumption underlying linear regression. Based on the above definitions we get:

$$n_0(x) = n_0(0) + m_{10}(x) - m_{01}(x) \quad (35)$$

$$n_1(x) = n_1(0) + m_{01}(x) - m_{10}(x) \quad (36)$$

$$\bar{Z}(x) = \frac{n_1(x)}{n} \quad (37)$$

For now, consider $m_{01}(x)$ and $m_{10}(x)$ to be fixed and known. We also assume that $n_0(0) = n_1(0) = n/2$ for simplicity. We introduce the estimates for $\hat{\mu}(x)$, $\hat{a}(x)$, and RSS_x . Set $a_1 = \sum (Z_i(x) - \bar{Z}(x))y_i$ (proportional to the covariance between Z and y) and $a_2 = \sum (Z_i(x) - \bar{Z}(x))^2$ (proportional to the variance of Z), then:

$$\hat{a}(x) = \frac{a_1}{a_2} \quad (38)$$

$$\hat{\mu}(x) = \bar{y} - \hat{a}(x)\bar{Z}(x) \quad (39)$$

After some calculation we get the values for a_1 and a_2 :

$$a_1 = \sum_{i \in N \setminus N_r(x)} (Z_i(x) - \bar{Z}(x))y_i + \sum_{i \in N_r(x)} (Z_i(x) - \bar{Z}(x))y_i \quad (40)$$

$$= \sum_{i \in N_r(x)} (1 - 2Z_i(0))y_i + \sum_{i \in N} (Z_i(0) - \bar{Z}(x))y_i \quad (41)$$

where $N = \{1, 2, \dots, n\}$ and $N_r(x) = \{\text{indices for recombined individuals at point } x\}$

As Equation (41) shows, we split a_1 into two sums, one of recombinants and the other nonrecombinant, which can be summarized as:

$$a_1 = a_{11} + a_{12} \tag{42}$$

where:

$$a_{11} = \sum_{i \in N_r(x)} (1 - 2Z_i(0))y_i = \sum_{i \in M_{01}} y_i - \sum_{i \in M_{10}} y_i \tag{43}$$

$$a_{12} = \sum_{i \in N} (Z_i(0) - \bar{Z}(x))y_i \tag{44}$$

a_{11} is a random weighted sum of phenotypes of recombinants and is the only stochastic variable involved in $\hat{a}(x)$. The stochastic behavior of a_{11} comes from the fact that $N_r(x)$ is a random variable related to the recombination process. Generally speaking, one can say that a_{11} captures all the stochasticity of $\hat{a}(x)$. We approximate the value of a_{11} by assuming y_i 's in a_{11} to be normally distributed:

$$\text{Assume (only for recombinant individuals) : } y_i \sim N(\mu(0) + a(0)Z_i(0), 1) \tag{45}$$

then:

$$a_{11} \sim N(\mu_{a_{11}}, \sigma_{a_{11}}^2) \tag{46}$$

where from (43):

$$\mu_{a_{11}} = \mu(0)(m_{01}(x) - m_{10}(x)) - a(0)m_{10}(x), \tag{47}$$

$$\sigma_{a_{11}}^2 = S_0 m_{01}(x) \left(1 - \frac{m_{01}(x) - 1}{n_0(x) - 1}\right) + S_1 m_{10}(x) \left(1 - \frac{m_{10}(x) - 1}{n_1(x) - 1}\right) \tag{48}$$

$$S_0 = 1, S_1 = 1 \tag{49}$$

We are using population correction coefficients $(1 - \frac{m_{01}(x) - 1}{n_0(x) - 1})$ and $(1 - \frac{m_{10}(x) - 1}{n_1(x) - 1})$, since we are sampling from a finite size population. If, for example, all individuals would recombine, all randomness in a_{11} would also disappear, which this correction reflects.

We also need to simplify a_2 . After some calculations, one gets

$$a_2 = n_0(x)\bar{Z}(x)^2 + n_1(x)(1 - \bar{Z}(x))^2 \tag{50}$$

Now that we have a_1 and a_2 , we should calculate RSS . If a dataset is divided into disjoint categories, it is known that one can write RSS_x as:

$$RSS_x = \sum (y_i - \bar{y})^2 - n_0(x)(\hat{\mu}_0(x) - \bar{y})^2 - n_1(x)(\hat{\mu}_1(x) - \bar{y})^2 \tag{51}$$

where $\hat{\mu}_0(x)$ and $\hat{\mu}_1(x)$ are the means of each category. Using this formula we get:

$$\begin{aligned} RSS_x &= RSS_{total} - n_0(x)(\hat{\mu}_0(x) - \bar{y})^2 - n_1(x)(\hat{\mu}_1(x) + \hat{a}(x) - \bar{y})^2 \\ &= RSS_t - a_2 \hat{a}(x)^2 \end{aligned} \tag{52}$$

Given the above description, we calculate the cumulative distribution function (CDF) of RSS_x :

$$\begin{aligned} F_{RSS_x}(rss_x) &= P_{a_{11}} \left(\hat{a}(x)^2 \leq \frac{rss_x - RSS_t}{-a_2} \right) \\ &= 1 - P_{a_{11}} \left(a_{11} \leq \sqrt{(RSS_t - rss_x)a_2} - a_{12} \right) \\ &\quad + P_{a_{11}} \left(a_{11} \leq -\sqrt{(RSS_t - rss_x)a_2} - a_{12} \right) \end{aligned} \tag{53}$$

Set $R = \sqrt{(RSS_t - rss_x)a_2}$ and then normalize the CDF to get the standard normal:

$$\text{Set: } z = \frac{a_{11} - \mu_{a_{11}}}{\sigma_{a_{11}}}, \text{ then :} \tag{54}$$

$$F_{RSS_x}(rss_x | m_{01}(x), m_{10}(x), y, Z(0), p) = 1 - \Phi \left(\frac{R - a_{12} - \mu_{a_{11}}}{\sigma_{a_{11}}} \right) + \Phi \left(\frac{-R - a_{12} - \mu_{a_{11}}}{\sigma_{a_{11}}} \right) \tag{55}$$

where Φ is the CDF of standard normal distribution. Given this result, we remove the constraints on $m_{01}(x)$ and $m_{10}(x)$ by summing over all values they can take. Since $n_0(0) = n_1(0) = n/2$, we have $m_{01}(x), m_{10}(x) \in \{0, 1, \dots, n/2\}$. Hence, one gets:

$$F_{RSS_x}(rss_x|y, Z(0), p) = \sum_{m_{01}(x)} [B(m_{01}(x), n_0(0), p) \sum_{m_{10}(x)} B(m_{10}(x), n_1(0), p) F_{RSS_x}(rss_x|m_{01}(x), m_{10}(x), y, Z(0), p)] \quad (56)$$

Where B is the binomial probability mass function, and the sums are over all values of $m_{01}(x), m_{10}(x) \in \{0, 1, \dots, n_{0/1}(0)\}$.

In this section, we derived the CDF of $RSS_x(x + x_0)$, the residual sum of squares at point x in the vicinity of x_0 , conditional on known m_{01} and m_{10} , by approximating the phenotypes of recombinants with a normal distribution. Then, we removed the condition on m_{01} and m_{10} by computing a probability-weighted sum over all values they can take. This sum forms an approximation of the CDF for RSS_x for fixed y_i 's. This CDF is then a two-leveled binomially weighted sum of mixture normals. One can find the value of rss_x for which this CDF is arbitrarily close to 1, for example, the $1 - \epsilon$ quantile. We thus have an upper bound for the value of rss_x at location $x + x_0$, which holds with probability $1 - \epsilon$.

4.3. Applications

Based on the derivation in Section 3.3, the PruneDIRECT algorithm, and the Lipschitz bound for the LogVar objective function in Equation (17), a method can be devised for accelerating multidimensional QTL searches and especially permutation testing for significance of an existing candidate QTL.

As indicated in the previous section, Equation (56) can be used to compute a quantile for the distribution of the residual sum of squares for a finite-size population, that is, F_{RSS_x} , at a specific distance from a hypothetical QTL explaining all genetic variance. We define a bound for the LogVar objective function based on the quantile of F_{RSS_x} instead of a Lipschitz bound. All boxes are compared against the LogVar transform for an F_{RSS_x} distribution based on the currently found optimum. Pruning of a box is possible if the $1 - \epsilon$ quantile for the LogVar distribution at distance x lies below the value evaluated in the centroid of a box with Manhattan radius x . If that condition is fulfilled, then the probability that there would be a new optimum, surpassing the current one, within the box is less than ϵ . The value for ϵ will then need to be chosen so that the aggregate probability of missing a minimum over all splits in a QTL search (and any permutation testing) is limited to an acceptable level.

In practice, the set of possible box radii appearing in the DIRECT search for a specific dataset is limited by the structure of the marker map. For a specific minimum value, the set of quantile limits can therefore be reused after first calculated, reducing the load of computing the sum of Gaussian-binomial products. Furthermore, the set of all binomial distribution coefficients for a specific distance can be computed in almost linear time with respect to the total number of individuals (see Appendix B). If, instead, each coefficient would be determined individually when computing the sum, the complexity per term would be on the order of $O(\log n)$, with a rather larger constant term, even when using an efficient implementation such as the ones in the Boost Project (2012) and R Development Core Team (2011).

Finally, the binomial distribution is rather thin-tailed. Therefore, it is not necessary to sum over all m_{01}, m_{10} , but rather only the subset from, for example, $np/2 - 8\sqrt{np(1-p)}/2$ to $np/2 + 8\sqrt{np(1-p)}/2$, still capturing almost all of the probability mass. This reduces the growth rate of the number of terms in the sum from $O(n^2)$ to $O(n)$. The resulting bound is also only slightly more conservative, as the quantiles will get shifted upward. Details for calculating the quantiles can be found in Appendix A.

In all, these adaptations and implementation aspects make it possible to use the defined bound for finite-population online within the Prune-DIRECT algorithm, while avoiding it representing a major part of the total computational load. Using LogVar as the objective function rather than the RSS directly is also beneficial even when the bound is not as simple as the Lipschitz bound, since the expectation will still be a matter of (close to) linear behavior with a bounded slope. This fact makes the splitting strategy used within DIRECT more efficient.

For intercrosses and other configurations with more than two states, we are still using the finite-size distribution determined for the backcross as an approximation, with an added condition. When more parameters are added, the computed RSS can fluctuate more. If the RSS and then LogVar value computed in a box is higher than what would be determined by average in any single linear regression with the same

number of degrees of freedom, that is, the mean of the χ^2 distribution, then the LogVar transform of the χ^2 mean, rather than the actually computed LogVar value, is used in determining the pruning condition. At this level, small disturbances due to the specific phenotype values overtake the general Gaussian assumptions.

5. RESULTS

The use of DIRECT in Ljungberg et al. (2004) only resulted in a very small bias when used for permutation testing. Evaluating our PruneDIRECT method with too few permutations and with a specific experimental dataset might render a false positive, such as the case in which an exhaustive search finds a different optimum that can be assumed to be very rare. Also, using a purely simulated dataset might hide nonideal properties of experimental datasets, giving a validation of our bound that would not work out in practice. For example, the patterns of missing genotype data can give additional confounding and sampling effects.

For these reasons, we decided to use an unpublished experimental dataset with a combination of microsatellite and SNP markers and varying patterns of missing data for our experiments. The dataset is described in Nettelblad et al. (2009). Based on inferred haplotypes in the F_0 generation, derived using the tool presented in that article, multiple replicates based on this population structure were constructed and QTL with varying dimensionality simulated accordingly. These replicates could then be analyzed for main effects and permuted runs. By creating hundreds or thousands of replicates, with hundreds of permutations within each, we can expect to discover any deviations between exhaustive search and running PruneDIRECT with a combined termination condition of minimum resolution equivalent to the exhaustive search lattice, and a pruning of impossible split candidates based on the strategy outlined in Section 4.3.

5.1. Simulations

The simulations were performed on the Tintin cluster at the UPPMAX computational resource center, running as single threads on nodes with AMD Opteron 6220 CPUs. The code is parallel, but since a very large number of replicates were used, many jobs corresponding to individual replicates were executed simultaneously, where each job executed the serial version of the code.

For each run, first the non permuted QTL model of specified size was fit, allowing all levels of interaction (each genotype–phenotype mean was a free parameter). Then, permutations were created. For exhaustive search, all possible candidate loci sets were explored in a 1 cM lattice. For PruneDIRECT, the minimum resolution was that same lattice, but in addition the bound was used to avoid splitting of some boxes, if it was definite that the values within that box could not improve upon the minimum found in the original main run. The ϵ value used in PruneDIRECT was 10^{-9} .

Table 1 presents the specific number of replicates, the simulated broad-sense heritability h^2 , and the number of permutations done for each replicate. Table 2 presents average, minimum and median number of function evaluations for complete sets of main run plus permutations. Timings show that objective function evaluations exceed 90% of the time used in the PruneDIRECT version. The maximum number of function evaluations can exceed the number of an exhaustive search, as the DIRECT implementation we started out from did not natively implement a discretized search grid and thus points coinciding in the discrete lattice could be evaluated multiple times.

Out of 500 simulated 2D QTL, the loci recovered in 25 of them resulted in more than 10 permuted datasets surpassing the simulated QTL, that is, less than 99% significant. These runs were also more time-

TABLE 1. SIZE OF VALIDATION TESTS FOR TWO AND THREE DIMENSIONS

d	Heritability (h^2)	Number of replicates	Number of permutations
2	0.09	500	1000
3	0.14	500	100

Time use for validating exhaustive search runs limits the possible size for 3D. Permutations were done per replicate.

TABLE 2. NUMBER OF FUNCTION EVALUATIONS FOR TWO AND THREE DIMENSIONS WITH EXHAUSTIVE SEARCH AND PRUNEDIRECT, RESPECTIVELY

<i>d</i>	<i>Method</i>	<i>Min</i>	<i>Median</i>	<i>Avg</i>	<i>Maximum</i>
2	Exhaustive	76,126	76,126	76,126	76,126
	PruneDIRECT	453	2,480	11,012	201,514
3	Exhaustive	998,536	998,536	998,536	998,536
	PruneDIRECT	1,565	17,826	54,457	323,628

Minimum, mean, median and maximum number of function evaluations per full run of main QTL search followed by the number of permutations given in Table 1 are reported. A very low number of PruneDIRECT runs resulted in full exhaustive searches. Note that the 3-dimensional scan used a lower number of permutations. All numbers are in thousands of function evaluations.

consuming, since the efficiency from the PruneDIRECT algorithm arises from the optimum being rare. If the QTL is not significant, then more boxes will be similar by random and thus fewer boxes can be pruned. If these are removed from the computed number of function evaluations, the average number of evaluations decreases to 7529 thousand. Acceleration compared to exhaustive search is only possible when the result from the main run rises above the noise floor defined by the quantiles and χ^2 distributions. The possible accelerations increase rapidly with stronger QTL signals. Finally, we tested simulating a single four-locus network with a total heritability of 0.30. Finding this took 9.0 million function evaluations, which compares favorably to the 963 millions just finding that network, that is, omitting any permutation testing, would require using an exhaustive search.

We also created 10,000 random QTL searches in a backcross templated on this intercross dataset (by fixing allele origin for either parent). Full concordance was achieved in classical one-dimensional QTL searches between the minimum found through exhaustive search and using PruneDIRECT.

6. DISCUSSION

This article presents an efficient scheme, named PruneDIRECT, for simultaneous mapping of several QTL, including permutation testing to determine significance.

The global optimization scheme DIRECT has earlier been adopted to QTL analysis in Ljungberg et al. (2004), and the reported speedups are several orders of magnitude Ljungberg et al. (2004). The work presented in this article does not improve on those speedup results. Instead, PruneDIRECT provides an option for performing QTL scans in settings where accuracy and guaranteed results are of imperative importance. We do so by letting the DIRECT process continue, executing until a full exhaustive search has in some sense been performed, but with a pruning taking place removing boxes that are not possible optima. This allows PruneDIRECT to be used not only for finding QTL candidates, but also to be reliably used in permutation testing to find the extreme end of the null hypothesis distribution. These computations can be used to compute significance thresholds as well as assessing the extreme value distribution, something which could form a basis for comparisons between models with different dimensionality and parametrization.

We finally propose to use a bound based on quantiles of an analytically derived distribution of the objective function at any distance from a minimum, in order to handle the fact that any real finite-sized population will be affected by the random and discrete nature of recombination and thus not be perfectly linear. We then used this bound in pruning the search trees. Our new objective function accelerates the performance of the DIRECT process, even when the pruning step is not added. The reason is that DIRECT performs best if the function is linear (finding the top of a single triangle in only a few iterations), and our transformed objective function will in general be almost piecewise linear. This understanding of the expected local form of the objective function, when full genotype information is available, could also be used to better assess probable QTL locations in cases of partial and incomplete genetic information.

It should be noted that the performance for PruneDIRECT is dependent on the heritability. For a trait with no heritability at all, finding the true optimum can in principle only be done by an exhaustive search, as very limited correlations are expected in the objective function between loci. Previous incarnations of DIRECT would have failed in those cases, while PruneDIRECT is adaptive and will perform more function evaluations. If one knows beforehand that QTL with a heritability below some limit h_{min}^2 are not relevant, then such information can be added to the pruning process and give better performance even in those cases.

We propose that in the future such a scheme could be used to effectively and effortlessly scan for highly significant multidimensional QTL in expression QTL data, where tens of thousands of putative phenotypes should be tested.

If the goal is to establish the exact significance level of a highly significant set of QTL, then our PruneDIRECT approach will excel. If most null hypothesis permutations for a dataset have a minimum that is inferior to that determined for the main model, the permuted runs can exit after only a hundred or so function evaluations, even in multiple dimensions. This allows doing runs equivalent to 100,000 or more permutations for loci, where significance levels above 99.9% would be relevant.

It should be noted that our approach for providing accuracy guarantees is based on some underlying assumptions. If phenotype distributions are far from normal, our approximation for the finite-size objective function distribution will not be accurate. It is also necessary that the number of recombinants in any interval is reflected by the mapping distances given. Our approximation takes the randomness of recombination into account, but if a marker map, for example, states a zero or very small distance in a region where the actual number of recombinants is much higher, then the approximation breaks down. Hence, it is important to use a genetic map that reflects the actual population studied. Naturally, additional tolerances could be entered into the method by, for example using a higher p value than that provided by the map. The results presented here could easily be extended to more generations than an F_2 intercross or a backcross, as the main difference will be that the rate of crossovers with regard to founder origin is multiplied.

For linked QTL, PruneDIRECT does well as long as the effects do not cancel each other out. When a single QTL effect is fitted, the effects from different QTL are, at least partially, confounded, as there is only a single variable (the indicator position within the linkage group), relating back to both components. Among other things, this means that the total explainable variance can be 0 at some point in the region between two linked QTL if the effects at the QTL have opposite signs. However, outside of the interval between the linked QTL, the behavior is completely identical to the presence of a single QTL at the position of the closest QTL in the set, with an effect equivalent to the combined average effect of all the linked QTL observed from that position. This can intuitively be understood from the memoryless nature of the exponential function. To avoid these issues, one could, for example, imagine enforcing a coarse-grain splitting of all boxes down to some resolution level and only then applying the prune mode in PruneDIRECT.

7. APPENDICES

A. Quantile calculations

We want to find the value of r_{SS_x} such that:

$$F_{RSS_x}(r_{SS_x}|y, Z(0), p) = 1 - \epsilon \quad (57)$$

One knows that $B(x, n, p)$ for x outside the interval $[\mu - 8\sigma, \mu + 8\sigma]$ is almost 0, where μ and σ are the mean and standard deviations of the binomial distribution considered. By using this, one gets:

$$F_{RSS_x}(r_{SS_x}|y, Z(0), p) = \sum_{m'_{01}(x)} [B(m'_{01}(x), n_0(0), p) \sum_{m'_{10}(x)} B(m'_{10}(x), n_1(0), p) F_{RSS_x}(r_{SS_x}|m'_{01}(x), m'_{10}(x), y, Z(0), p)] + \delta \quad (58)$$

$$\text{where: } \delta \approx 0, \quad (59)$$

$$m'_{01}(x), m'_{10}(x) \in \left[\frac{np}{2} - 8\sqrt{\frac{np(1-p)}{2}}, \frac{np}{2} + 8\sqrt{\frac{np(1-p)}{2}} \right] \quad (60)$$

An alternate region for $m'_{01}(x), m'_{10}(x)$ is to choose $m'_{10}(x) \in [np/2 - 8\sqrt{np(1-p)/2}, np/2 + 8\sqrt{np(1-p)/2}]$ and then $m'_{01}(x)$ in a circle around $m'_{10}(x)$ with radius $8\sqrt{np(1-p)/2}$. After all, one should find r_{SS_x} such that:

$$\sum_{m'_{01}(x)} [B(m'_{01}(x), n_0(0), p) \sum_{m'_{10}(x)} B(m'_{10}(x), n_1(0), p) F_{RSS_x}(r_{SS_x}|m'_{01}(x), m'_{10}(x), y, Z(0), p)] = 1 - \epsilon \quad (61)$$

This equation can be solved efficiently numerically by, for example, the bisection method. In our implementation, we only solve it down to a resolution level of 0.04, overestimating the location of the

LogVar quantile by at most the equivalent of 1 cM in the infinite-size population case. Furthermore, since we know the quantile location to be monotonous in terms of increasing x , our precalculation scheme for different box radii can select the low and high initial bisection bounds based on previously calculated solutions for lower and higher x , respectively. In all, this can reduce the amortized number of bisection steps for computing the sum to a handful per unique x . For different x , the ϕ values going into the sum are also unchanged and can thus be stored precalculated, since evaluating the normal CDF is relatively time-consuming.

B. Calculating the full set of binomial coefficients

Most library implementations supporting the binomial distribution compute single values individually. Even if one is using, for example, R (R Development Core Team, 2011) and provides a vector as x when computing $B(x, n, p)$, each step will be performed independently. The computational complexity for these steps is nontrivial and on the order of $O(\log(x) + \log(n))$.

If one knows that all binomial coefficients will be used in a sum, it is instead far more efficient to compute them all as part of the same process. The total computational complexity then becomes linear, with only a few multiplications and divisions per iteration, augmented by a few exponentializations to renormalize the carry values in order to avoid falling outside the dynamic range of conventional floating point implementations.

ACKNOWLEDGMENTS

Per Jensen, Leif Andersson, and Olle Kämpe are acknowledged for sharing experimental data for evaluation of the method. The computations were performed on resources provided by SNIC through Uppsala Multidisciplinary Center for Advanced Computational Science (UPPMAX).

AUTHOR DISCLOSURE STATEMENT

The authors declare that no competing financial interests exist.

REFERENCES

- Boost Project. 2012. Boost C++ Libraries. Available at www.boost.org
- Carbonell, E.A., Gerig, T.M., Balansard, E., and Asins, M.J. 1992. Interval mapping in the analysis of nonadditive quantitative trait loci. *Biometrics* 48, 305–315.
- Carlborg, O., Andersson, L., and Kinghorn, B. 2000. The use of a genetic algorithm for simultaneous mapping of multiple interacting quantitative trait loci. *Genetics* 155, 2003–2010.
- Carlborg, O., and Haley, C.S. 2004. Epistasis: too often neglected in complex trait studies? *Nat Rev Genet* 5, 618–25.
- Churchill, G., and Doerge, R. 1994. Empirical threshold values for quantitative trait mapping. *Genetics* 138, 963–971.
- Doerge, R. 2002. Mapping and analysis of quantitative trait loci in experimental populations. *Nature reviews-Genetics* 3, 43–52.
- Galerpin, E. 1985. The cubic algorithm. *J of Mathematical Analysis and Applications* 112, 635–640.
- Haldane, J.B.S. 1919. The combination of linkage values, and the calculation of distance between the loci of linked factors. *J Genet* 8, 299–309.
- Haley, C.S., and Knott, S.A. 1992. A simple regression method for mapping quantitative trait loci in line crosses using flanking markers. *Heredity* 69, 315–24.
- Jones, D., Perttunen, C., and Stuckman, B. 1993. Lipschitzian optimization without the lipschitz constant. *J. Optimization Theory App* 79, 157–181.
- Lander, E.S., and Botstein, D. 1989. Mapping Mendelian Factors Underlying Quantitative Traits Using RFLP Linkage Maps. *Genetics* 121, 185–199.
- Ljungberg, K., Holmgren, S., and Carlborg, O. 2004. Simultaneous search for multiple QTL using the global optimization algorithm DIRECT. *Bioinformatics* 20, 1887–1895.
- Mladineo, R. 1987. An algorithm for finding the global maximum of a multimodal, multivariate function. *Mathematical Programming* 34(1986), 188–200.

- Nettelblad, C., Holmgren, S., Crooks, L., and Carlborg, O. 2009. cnf2freq: Efficient determination of genotype and haplotype probabilities in outbred populations using Markov models, 307–319. In Rajasekaran, S., ed. *BICoB 2009, Lecture Notes in Computer Science*, Vol. 5462. Springer New York.
- Pinter, J. 1986. Globally convergent methods for n-dimensional multiextremal optimization. *Optimization* 17, 187–202.
- R Development Core Team. 2011. R: A language and environment for statistical computing. Available online at www.R-project.org/
- Shubert, B. 1972. A sequential method seeking the global maximum of a function. *SIAM J. on Numerical Analysis* 9, 379–388.
- Wu, R., Ma, C., and Casella, G. 2007. *Statistical Genetics of Quantitative Traits: Linkage, Maps and QTL*. Springer, San Francisco, CA.

Address correspondence to:

Mr. Behrang Mahjani
Department of Information Technology
Uppsala University
Box 337
SE-751 05 Uppsala
Sweden

E-mail: behrang.mahjani@it.uu.se