CSIRO PUBLISHING
www.publish.csiro.au/journals/fpb *Functional Plant Biology*, 2008, **35**, 837–849

Quasi-Monte Carlo simulation of the light environment of plants

Mikolaj Cieslak[A,E,F], *Christiane Lemieux*[B], *Jim Hanan*[C] *and Przemyslaw Prusinkiewicz*[D]

[A]The University of Queensland, School of Physical Sciences, Qld 4072, Australia.
[B]Department of Statistics and Actuarial Science, University of Waterloo, ON N2L 3G1, Canada.
[C]The University of Queensland, Centre for Biological Information Technology, Qld 4072, Australia.
[D]Department of Computer Science, University of Calgary, AB T2N 1N4, Canada.
[E]The Horticulture and Food Research Institute of New Zealand Limited, Palmerston North Research Centre, Palmerston North 4474, New Zealand.
[F]Corresponding author. Email: cieslak@maths.uq.edu.au

This paper originates from a presentation at the 5th International Workshop on Functional–Structural Plant Models, Napier, New Zealand, November 2007.

Abstract. The distribution of light in the canopy is a major factor regulating the growth and development of a plant. The main variables of interest are the amount of photosynthetically active radiation (PAR) reaching different elements of the plant canopy, and the quality (spectral composition) of light reaching these elements. A light environment model based on Monte Carlo (MC) path tracing of photons, capable of computing both PAR and the spectral composition of light, was developed by Měch (1997), and can be conveniently interfaced with virtual plants expressed using the open L-system formalism. To improve the efficiency of the light distribution calculations provided by Měch's *MonteCarlo* program, we have implemented a similar program *QuasiMC*, which supports a more efficient randomised quasi-Monte Carlo sampling method (RQMC). We have validated *QuasiMC* by comparing it with *MonteCarlo* and with the radiosity-based CARIBU software (Chelle *et al*. 2004), and we show that these two programs produce consistent results. We also assessed the performance of the RQMC path tracing algorithm by comparing it with Monte Carlo path tracing and confirmed that RQMC offers a speed and/or accuracy improvement over MC.

Additional keywords: light simulation, open L-system, PAR, path tracing, red/far red ratio, (randomised) quasi-Monte Carlo sampling, variance reduction, virtual plant modelling.

Introduction

The distribution of light in the canopy is a major factor regulating the growth and development of plants. Consequently, simulation of light environment is an important component of functional-structural plant modelling. The main variables of interest are the amount of photosynthetically active radiation (PAR) reaching different elements of the plant canopy, and the quality (spectral composition) of light reaching those elements, which may be a signal for photomorphogenesis.

Light environment models estimate the radiant energy reaching a plant from direct light sources (e.g. sun and sky), and may include indirect sources (e.g. light reflected from or transmitted through plant organs) for increased accuracy. The estimation techniques are related to the description of the plant canopy, which may be approximated as a turbid medium or specified explicitly as a 3-D geometric structure (virtual plant) (Room *et al*. 1996). Models of light environment operating at the plant structure level are usually based either on the Monte Carlo path tracing method or the radiosity method (Chelle and Andrieu 2007). In the past, attention was given to improvements of the efficiency of the radiosity method (Chelle and Andrieu 1998; Soler *et al*. 2003), with the Monte Carlo path tracing used as a benchmark for comparisons. Soler *et al*. (2003) mentioned two drawbacks of the Monte Carlo approach: slow convergence and poor control over the accuracy of the result. Here, we address these drawbacks by improving the accuracy and efficiency of path tracing in the context of biological applications.

Monte Carlo path tracing is derived from standard ray tracing. Both methods approximate the solution to the rendering equation, which describes the transfer of light energy between a point on a surface and a point on another surface (Kajiya 1986). Path tracing differs from ray tracing by sampling many possible light paths from a surface point, instead of recursively following a single reflected and a single refracted ray. Consequently, path tracing captures some optical phenomena, such as diffuse light reflection and transmission, that ray tracing does not (Watt 2000).

Ross and Marshak (1988) developed a radiative transfer Monte Carlo model to simulate canopy bidirectional reflectance. Their model, however, relies on estimating parameters of plant canopy architecture that allow for the simulation of only one canopy type. The Monte Carlo model by Govaerts (1996) overcomes this limitation by using an explicit 3-D description of canopy architecture [see the review by Disney *et al*. (2000)]. Měch (1997) developed a similar model based on

 10.1071/FP08082 1445-4408/08/100837

Monte Carlo path tracing, which can be interfaced conveniently with virtual plants expressed using the open L-system formalism (Měch and Prusinkiewicz 1996). He applied two variance-reduction techniques to calculate light distribution efficiently. First, importance sampling (Rubinstein 1981) is used to generate rays preferentially in the direction from which (or towards which) most energy propagates, as opposed to generating them with a uniform distribution. Second, individual rays may carry information regarding several wavelengths simultaneously. This technique reduces the variance of the ratios of energy associated with different light wavelengths, which is important in estimating the spectrum of light reaching plant organs.

To further improve the efficiency of light distribution calculations provided by Měch's model (Měch 1997), we have implemented an alternative method for generating reflected or transmitted rays, called the quasi-Monte Carlo method. In the context of computer graphics, quasi-Monte Carlo sampling was surveyed by Owen (2003). Although Monte Carlo (MC) path tracing relies on random sampling of the space of reflected or transmitted rays, and results in a set of independent paths, quasi-Monte Carlo (QMC) is based on a highly regular sampling that produces a set of correlated paths. As shown by Keller (1996) and Veach (1997), this reduces the number of rays required for path-tracing virtual 3-D scenes within given error bounds. To estimate the variance of our samples and provide some measure of the accuracy of the results, we implemented an extension to QMC sampling called randomised quasi-Monte Carlo (RQMC) sampling (L'Ecuyer and Lemieux 2002). QMC methods are deterministic, but in RQMC the sampling points are randomised in a way that preserves their highly regular distribution. This makes it possible to estimate the error of the computation.

In this paper, we present our implementation of the RQMC path tracing algorithm in the context of calculating light distribution within a virtual plant canopy. In particular, we discuss how we coupled path tracing with RQMC sampling. The resulting light simulation program, *QuasiMC*, is intended to operate in concert with a simulator of plant development. Specifically, we interfaced *QuasiMC* with two L-system-based plant simulators, *cpfg* (Prusinkiewicz *et al*. 1997) and *lpfg* (Karwowski and Prusinkiewicz 2004), which are parts of the L-studio/VLAB modelling software (Prusinkiewicz *et al*. 2000). To assess the RQMC path tracing algorithm, we applied *QuasiMC* to compute light distribution in a mix of triangles within a cube, and we compared the results with those obtained using MC path tracing as well as using radiosity-based CARIBU software (Chelle *et al*. 2004). Finally, we evaluated our program in the context of a more realistic application: computation of light distribution in a kiwifruit vine (a broadleaf liana ~2 m tall in cultivation) under various light conditions.

The RQMC path tracing algorithm

To compute light distribution within a scene, the *QuasiMC* program uses the path tracing algorithm. This algorithm was introduced by Kajiya (1986) as a Monte Carlo solution to an integral equation that describes radiance at any point on a surface as a function of direct and indirect light from all other points of all surfaces. We present an overview of this algorithm as it applies to our program.

The path tracing algorithm

The computation proceeds by generating a user-specified number of rays, where each ray is traced through the plant canopy until its radiant energy is absorbed by the plant's organs. A light ray may originate from a light source and be traced towards the plant canopy, or from a plant organ and be traced towards a light source. In computer graphics terminology, both methods are called particle tracing and are classified as view-independent forms of path tracing (Shirley *et al*. 2005). When the canopy is dense, the former method is advantageous, because rays originating at plant organs would rarely reach a light source. In contrast, when organs are small relative to the whole plant, and are highly dispersed, the latter method is advantageous, because rays traced from a light source would often miss organs. In the *QuasiMC* program, the user can choose either of these methods, but bidirectional path tracing (combined tracing of rays towards the plant canopy and towards the light sources) is not available.

The starting point and direction of a ray depend on the type of light source and the type of path tracing (from the light source to the canopy or from the canopy to the light source), as specified by the user. Two models of light sources are available in the *QuasiMC* program: multiple directional light sources and an approximation of the sky. The latter is calibrated by latitude, day of year, time of day, and clear or overcast sky conditions. Each directional light source is sampled according to its intensity and direction for outgoing light. In the case of sampling from the light source to the canopy, all rays originate on the surface of the bounding sphere encompassing the scene (Jensen 2001). In the case of sampling from the canopy to the light source, the ray's starting point is randomly chosen on some surface, and the initial ray direction is stochastically generated on the basis of the user-supplied source light directions and the local light model of the surface. For the sky model, the hemisphere encompassing the plant canopy is sampled using the probability density given by the CIE standard clear sky model or overcast sky model (CIE-110 1994). This ensures more rays are sampled from parts of the sky where the radiant energy is greatest.

Given the starting point and direction of a ray, the next step of the algorithm is to follow the ray through the plant canopy and to find the plant organ surface that the ray intersects at the smallest distance from the starting point. Once such a surface is found, a local light model is applied to calculate how much light is reflected, transmitted and absorbed at the intersection point. The fraction r of the radiant energy reflected from the surface and the fraction t of the energy transmitted through the surface satisfy the inequality $0 \leq r+t \leq 1$, and are specified by the user. The absorbed fraction is equal to $1 - (r+t)$. A reflected or transmitted ray is then generated stochastically with the probabilities $r/(r+t)$ and $t/(r+t)$, respectively. The direction and energy of this ray are determined by the bidirectional reflectance distribution function (BRDF) or bidirectional transmittance distribution function (BTDF) of the organ's surface, as described below. There are many reported forms of these distribution functions, which depend mainly on the surface's material type (Shirley *et al*. 2005). *QuasiMC* supports two distribution functions: the Lambertian function (Shirley *et al*. 2005) or the modified Phong function (Shirley and Wang 1992), which are used to characterise both the BRDF

and the BTDF of the plant organs. The Lambertian BRDF and BTDF are expressed as the probability density function

$$p_{\mathrm{L}}(\theta,\ \phi) = \frac{1}{\pi}\cos\theta, \tag{1}$$

where $\theta \in [0, \pi/2)$ is the angle between the surface normal and the ray direction, and $\phi \in [0, 2\pi)$ is the rotation around the normal (which does not effect the value of $p_{\mathrm{L}}(\theta, \phi)$). The modified Phong BRDF and BTDF are expressed as the probability density function

$$p_{\mathrm{Ph}}(\alpha,\ \beta) = \frac{n+2}{8\pi}\cos^{n}\frac{\alpha}{2}, \tag{2}$$

where α is the angle relative to the ideally reflected/transmitted direction, β is the rotation around this direction, and $n \geq 0$ is a parameter characterising the surface smoothness. The angle α is constrained to values for which the reflected/transmitted ray is on the correct side of the surface. As n becomes large, the surface approaches a mirror (for BRDF) or allows the ray to pass through the surface without scattering (for BTDF).

Although it is possible to sample the direction of rays reflected/transmitted using Eqn (1) or (2) uniformly, it is more efficient to use importance sampling. The ray's energy is then assumed to be constant (equal to the energy that is not absorbed by the surface), while the ray's direction is generated using the inverse of the cumulative distribution function of $p_{\mathrm{L}}(\theta, \phi)$ or $p_{\mathrm{Ph}}(\alpha, \beta)$ (Shirley *et al*. 2005). Specifically, the ray's direction is determined by the formula,

$$(\theta,\ \phi) = (\arccos\sqrt{\xi_1}, 2\pi\xi_2), \tag{3}$$

for the Lambertian BRDF/BTDF, where ξ_1 and ξ_2 are uniformly distributed random numbers in the interval (0,1], or

$$(\alpha,\ \beta) = (2\arccos[(1-\xi_1)^{\frac{1}{n+2}}], 2\pi\xi_2), \tag{4}$$

for the modified Phong BRDF/BTDF, where ξ_2 is a uniformly distributed random number in the interval (0,1], and ξ_1 is constrained to guarantee that the ray will appear on the correct side of the surface. Different values of the scattering exponent n can be used to generate reflected and transmitted rays.

A ray is terminated when its radiant energy is close to zero or when it does not intersect any of the plant's organs. To ensure that the tracing of a path will stop at some point, the user may specify the maximum number of reflections and transmissions. Unfortunately, this method may introduce a statistical bias, because it disregards the ray's radiant energy upon termination. A solution to this problem is offered by the 'Russian roulette' method (Arvo and Kirk 1990). In that case, if the radiant energy of a ray falls below a user-defined threshold, the ray is terminated with some probability p, otherwise its radiant energy is increased by $1/(1-p)$. This increase in a non-terminated ray's energy ensures that the solution converges to the correct result in the limit (Arvo and Kirk 1990).

Thus far, we have ignored the dependence of a ray's radiant energy and a surface's BRDF and BTDF on the spectrum of light being considered. In reality, the BRDF and BTDF parameters for light of two different wavelengths may be different (see the BRDF measurements by Bousquet *et al*. (2005) and the BRDF/BTDF measurements by Breece and Holmes (1971), or the review by Jacquemoud and Ustin (2001) on the optical properties of leaves). One strategy to incorporate spectral effects is to apply path tracing to each wavelength independently. A more efficient strategy is to use a single ray that carries the radiant energy for several wavelengths (Měch 1997; Evans and McCool 1999). The improvement in efficiency comes from the reduction in the number of rays and the strong positive correlation of the radiant energy associated with each wavelength, if the angular distribution [e.g. the scattering exponent n in Eqn (2)] at these wavelengths is similar. We assume that the scattering exponent is the same for all wavelengths in our implementation; however, an extension to the general case is possible (Jensen 2001). This type of Monte Carlo calculation falls into the general category of 'correlated sampling' (Spanier and Gelbard 1969; Rubinstein 1981), and is based on using the same random numbers in two similar processes to reduce the difference in variance between them, with respect to two independent simulations.

The principal change to the path tracing algorithm, needed to capture spectral effects with correlated sampling, affects the component responsible for stochastic generation of reflected or transmitted rays. The direction of a newly reflected/transmitted ray must now take into account the radiant energy associated with each wavelength; furthermore, it is necessary to keep track of the radiant energy of those wavelengths in one ray. In *QuasiMC*, the reflected/transmitted ray is chosen based on a weighted probability over all wavelengths. The Russian roulette method is extended to decide if the ray should be reflected, transmitted, or absorbed, and the radiant energy of all wavelengths is scaled accordingly (Jensen 2001). Assuming a spectrum with M wavelengths, for each spectral component $\lambda = 1, \ldots, M$ we consider products $r_\lambda \Phi_{i,\lambda}$ and $t_\lambda \Phi_{i,\lambda}$ of the fraction of radiant energy reflected (r_λ) or transmitted (t_λ) by the surface, and the radiant energy reaching the surface (incident energy), $\Phi_{i,\lambda}$. We then define

$$p_r = \sum_{\lambda=1}^{M} r_\lambda \Phi_{i,\lambda} \text{ and } p_t = \sum_{\lambda=1}^{M} t_\lambda \Phi_{i,\lambda}, \tag{5}$$

as the sums of these products over all wavelengths, and calculate the probabilities of generating a reflected or transmitted ray as $p_r/(p_r+p_t)$ and $p_t/(p_r+p_t)$, respectively. The direction of the reflected/transmitted ray is calculated according to either the Lambertian or modified Phong model (Eqn 3 or 4), where we assume the scattering exponent n is the same for all wavelengths. At each intersection of a ray with a surface, the reflected or refracted ray is generated in a manner fully consistent with one particular wavelength. Following the Russian roulette technique, the radiant energy associated with all wavelengths is then adjusted to conserve energy (Měch 1997; Jensen 2001). We scale the reflected ($\Phi_{r,\lambda}$) or transmitted ($\Phi_{t,\lambda}$) energy of the ray using the following scheme:

$$\begin{aligned} \Phi_{r,\lambda} &= \frac{r_\lambda \Phi_{i,\lambda}}{p_r/(p_r+p_t)} \\ \Phi_{t,\lambda} &= \frac{t_\lambda \Phi_{i,\lambda}}{p_t/(p_r+p_t)}, \end{aligned} \tag{6}$$

where $\lambda = 1, \ldots, M$. For example, if 1000 rays are traced towards a surface that reflects 50% of the incoming light, the Russian

roulette technique allows us to reflect 500 rays with full power instead of 1000 rays with half power (Jensen 2001).

Stochastic sampling methods

The *QuasiMC* program implements two stochastic methods for generating sampling rays in path tracing. The first one is the Monte Carlo method, in which s uniformly distributed random numbers are used for generating random variates from different distributions to trace one ray through the plant canopy. Samples from these distributions are generated by an inverse transform method or acceptance-rejection method (Rubinstein 1981; Chapter 3) using pseudo-random numbers generated with a combined multiple recursive generator by L'Ecuyer (1999). For convenience, we represent these random numbers by a vector $\mathbf{u}=(u_1,\ldots,u_s)$ uniformly distributed over $[0,1)^s$. An approximation of radiant energy absorbed by each leaf within the plant canopy is made by generating N such vectors (i.e. by constructing a point set $P_N=\{\mathbf{u}_1,\ldots,\mathbf{u}_N\}\subset[0,1)^s$) and tracing N rays according to this point set. The second method is the quasi-Monte Carlo method, which may be considered as a deterministic counterpart of the Monte Carlo method. In this case, the point set P_N is constructed with a more regular distribution than the random point set used in Monte Carlo. Both methods were reviewed in the context of light transport simulation by Veach (1997).

Several algorithms for generating sets or sequences of regularly spaced sampling points have been proposed for use in QMC computations; the most commonly used in practice are by Korobov (1959), Sobol' (1967), Halton (1960) and Faure (1982). We chose Korobov's algorithm, because it can generate sampling points dynamically, as the tracing proceeds, without knowing in advance how many ray-surface intersections will occur in each path, and thus, how many numbers $u_1,\ldots,u_s$ will be needed to trace it. In other words, this method can be used when the dimension s of the vector $u_1,\ldots,u_s$ of uniformly distributed numbers is not known *a priori*. Korobov's algorithm requires choosing an integer $a\in\{1,\ldots,N-1\}$ relatively prime to the assumed sample size N. The point set is then calculated as

$$P_N=\left\{\mathbf{u}_i=\left(\frac{i-1}{N}\cdot(1,a,a^2 \bmod N,\ldots,a^{s-1} \bmod N)\right) \bmod 1,\right.$$
$$\left. i=1,\ldots,N\right\}, \qquad (7)$$

where the modulo 1 operation is applied component-wise after multiplication of the sequence by $(i-1)/N$ (Korobov 1959).

To estimate variance/error of the computation performed by *QuasiMC*, the QMC point set is randomised in a way that preserves the highly regular distribution of the sampling points. As the light paths generated from the QMC point set are not independent, simple error estimation as in MC is otherwise not possible. For the Korobov method, a very simple way to randomise the underlying (deterministic) point set P_N is to generate a random vector $\mathbf{v}$ with a uniform distribution in $[0,1)^s$, and then add that same vector – modulo 1 component-wise – to each point in P_N (Cranley and Patterson 1976). The rays in that case are traced using the randomised points.

$$\tilde{\mathbf{u}}_i=(\mathbf{u}_i+\mathbf{v}) \bmod 1, \qquad (8)$$

where $i=1,\ldots,N$. Figure 1 shows a QMC point set in two dimensions, which has been randomised using this approach. Note that each vector $\tilde{\mathbf{u}}_i$ is uniformly distributed over $[0,1)^s$, since $\mathbf{v}$ is. A ray traced with $\tilde{\mathbf{u}}_i$, thus, has the same statistical properties as one traced using Monte Carlo sampling. The difference is that with RQMC sampling, the N rays traced using vectors $\tilde{\mathbf{u}}_1,\ldots,\tilde{\mathbf{u}}_N$ are dependent, and designed to provide a more representative sample of light than N independent random rays. Note that, when using RQMC sampling, we have to decide in advance how to assign coordinates in an RQMC vector to random variables in the algorithm. This is because

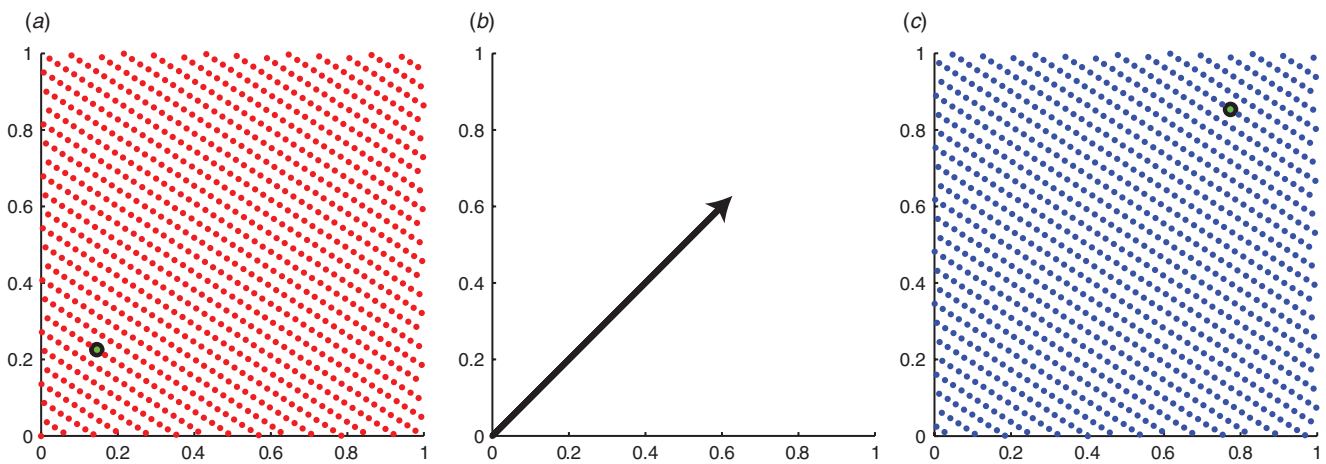

Fig. 1. Randomisation of a QMC point set. (*a*) The QMC set P_N of pairs of numbers ($s=2$), generated using the Korobov method. Each pair is represented as a point in the square $[0,1)\times[0,1)$. (*b*) A sample random vector $\mathbf{v}$. (*c*) The RQMC point set obtained by translating set P_N by vector $\mathbf{v}$. Black dots indicate one of the points before and after translation.

the full benefit of using RQMC sampling can only be obtained if the numbers $\tilde{u}_{i,1}, \ldots, \tilde{u}_{i,s}$ are used for the same purpose across all $i = 1, \ldots, N$.

By repeating the process of ray-tracing the canopy M times – using M independent random vectors $\mathbf{v}_1, \ldots, \mathbf{v}_M$, we can estimate statistical properties of the quantities of interest, such as their variance or standard error (the variance is simply estimated over M samples) (L'Ecuyer and Lemieux 2002). In conclusion, randomised quasi-Monte Carlo sampling is a general variance-reduction technique, which we couple with a more problem-specific method, namely importance sampling, to improve upon Monte Carlo ray-tracing.

The QuasiMC program for simulating light distribution in a canopy

The L-system-based plant model and the light environment model are executed as two separate processes that communicate using the open L-system formalism (Měch and Prusinkiewicz 1996). The plant simulator, in our experiments the L-studio/VLAB program *cpfg* or *lpfg* (Prusinkiewicz 2004), sends information about the location, size and orientation of the virtual plant's organs to the light environment. Each plant organ can be represented as a triangle, a parallelogram, a user-defined polygon, or a Bezier surface. The light environment simulator, *QuasiMC*, returns light distribution among those organs on request from the plant simulator. Thus, *QuasiMC* can dynamically estimate light distribution in a canopy during the simulated development of a plant.

The interplay between *lpfg* and *QuasiMC* is illustrated in Fig. 2. Both simulators can visualise the current state of the model in separate windows on the screen. *QuasiMC* shows the virtual plant's organs that are sent from the plant simulator, and shades each one according to the amount of radiant energy absorbed by it. The user can independently manipulate the view of the model in both simulators, with the same type of interface. The parameters relevant to the light distribution simulation are specified in a separate file that is read by *QuasiMC* at the beginning of the simulation. These parameters are described in detail in the *QuasiMC* user manual (Cieslak 2004) and are summarised below.

Configurable parameters of the QuasiMC program

Operation of the *QuasiMC* program is controlled by a set of parameters read from a file. One class of parameters characterises physical attributes of the simulation, in particular the lighting conditions and the optical properties of plant organs. Another class controls computational aspects of the light simulation. If *QuasiMC* is used in the context of a simulation of development, the values of all parameters are fixed over the entire simulation.

As mentioned before, *QuasiMC* supports two types of light sources: directional sources, with all rays from the same source having the same initial direction, and a hemispherical approximation of the sky based on the CIE standard clear sky model and overcast sky model (CIE-110 1994). The user can specify parameters of each light source, for example the radiant power of each directional light source, and the time of day, location, and weather (clear or overcast) for the sky model. The CIE models are based on a comparison with measured sky conditions, and do not explicitly model clouds. The overcast sky model accounts for brightening of the sky from the horizon towards the zenith (so that horizon luminance is ~1/3 of that at the zenith). The clear sky model accounts for a bright region around the sun, and a slight brightening around the horizon.

QuasiMC supports user-specified materials, which characterise the amount of radiant energy reflected and transmitted from a surface and the scattering direction of a reflected/transmitted ray according to the Lambertian model or modified Phong model (see the previous section on the path tracing algorithm). The user specifies the fraction of reflected and transmitted light for each wavelength, and the scattering exponent for the Phong model. However, if one ray is used to carry information about several wavelengths, the program assumes that this exponent is the same for all wavelengths. The L-system model can then associate the adaxial and abaxial sides of each plant organ with the specific materials.

Parameters controlling the light distribution computation in *QuasiMC* are as follows:

(1) sampling type: Monte Carlo or randomised quasi-Monte Carlo,
(2) number of rays (must be a power of 2 when using QMC with the Korobov generator),

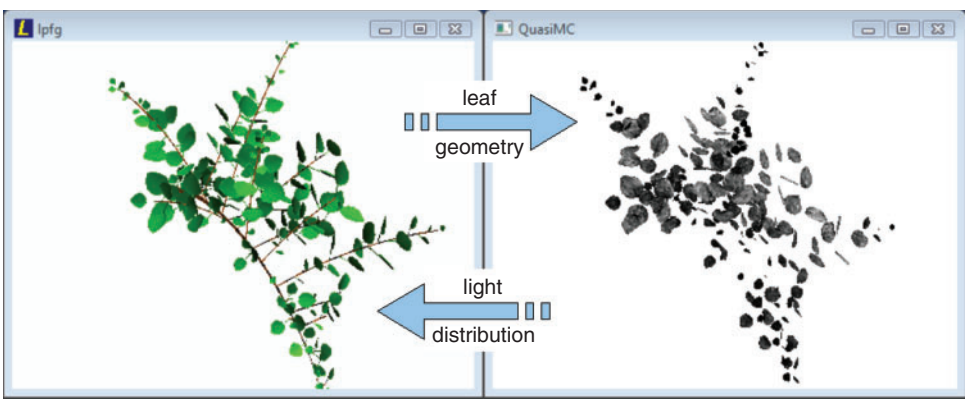

Fig. 2. The interplay between the plant simulator (*lpfg*, left) and the light environment simulator (*QuasiMC*, right). The plant simulator sends information about the size, position and orientation of each leaf, and the light environment simulator returns the absorbed radiant energy available to each leaf.

(3) tracing method (from light sources to the canopy or from the canopy to the sources),
(4) number of runs (with different randomisations of the QMC set, as needed to estimate standard error/variance of results), and
(5) granularity of uniform spatial subdivision of the scene (which is used to speed up computation).

The benefit of using a spatial subdivision comes from the reduced time needed to find intersections between a ray and a plant organ: where a basic method would test each organ for a possible intersection with a ray, spatial subdivision methods only test those organs that are close to the ray (Shirley *et al*. 2005).

Results

The purpose of this section is to illustrate the operation of *QuasiMC* in the context of plant simulations realised using *cpfg* and *lpfg*. We compare the effectiveness of the RQMC sampling *v.* MC sampling, and we test *QuasiMC* by comparing it with CARIBU software, which computes light environment using the radiosity method (Chelle *et al*. 2004). We first considered a simplified plant canopy that is modelled as a mix of triangles within a cube, and then moved on to a more complex model of a kiwifruit vine. In all the following computations, we used the method of tracing rays from the light source towards the plant canopy. The alternative method of tracing rays from the canopy to the light sources was used by Cici *et al*. (2008) in a virtual plant model of crop–weed interaction.

Comparison of light evaluations using RQMC and MC sampling

To compare the RQMC and MC sampling methods in the context of light simulation, we applied both methods to compute the absorbed radiant energy and sample variance for a set of triangles, randomly distributed within a cube. This test set is similar to the virtual canopy proposed by Chelle *et al*. (1998). Our model is expressed by the following open L-system in the *cpfg* language (Měch *et al*. 2005):

```
#define N 5000                  /* number of triangles */
#define x ran (1) * 100−50      /* random position along x-axis */
#define y ran (1) * 100−50      /* random position along y-axis */
#define z ran (1) * 100−50      /* random position along z-axis */
#define α ran (360)             /*random orientation around
                                turtle's H-axis */
#define β ran (360)             /* random orientation around
                                turtle's L-axis */
#define γ ran (360)             /* random orientation around
                                turtle's U-axis */
#define l 5                     /* length of a triangle's edge */
#define h sqrt (3)/2 * 5        /* height of a triangle */

ω : L(N)
p1 : L(n) > 0 → @M(x, y, z)/(α) ^(β) + (γ) ?E(0) T(l, h) L(n−1)
p2 : ?E(light) < T(L, H) : * {printf("light: %g\n", light);}→T(l, h)
```

For a comparison, the specification of the same L-system in the L + C language supported by the *lpfg* plant simulator (Karwowski and Lane 2007) has the form:

```
/* #define section as above */
module L(int);
module E1(float);
module T(float, float);
ω : L(N)
p1 : L(n):{
  if (n > 0)
     produce MoveTo(x, y, z) RollR(α) Up(β)Left(γ)
             E1(0) T(l, h) L(n−1);}
p2 : E1(light) < T(L, H) : {Printf("light: %g\n", light); produce
     T(l, h);}
```

We omitted here the #define statements, which are the same as in *cpfg*. In either case, the axiom ω consists of a module *L* that will serve as the generator of *N* triangles in the virtual canopy. In each application of production p_1, the module L gives rise to an equilateral triangle T with a 5 cm edge length, positioned and oriented at random within a $100 \times 100 \times 100$ cm cube. The triangle module is preceded by a communication module, ?E or E1, forming a pair of modules that is sent to the *QuasiMC* program. Once all the *N* triangles have been created, *QuasiMC* calculates light distribution and returns the absorbed energy value for each triangle through the *light* variable of the corresponding module ?E or E1. These values are output by production p_2 and provide the basis for a further analysis.

We compared the RQMC method from our program with the MC method from Měch's *MonteCarlo* program by performing numerical experiments on the virtual canopy model with 5000 triangles (Fig. 3). Each triangle was set to reflect 10% and transmit 10% of the incident radiation, and the Lambertian local light

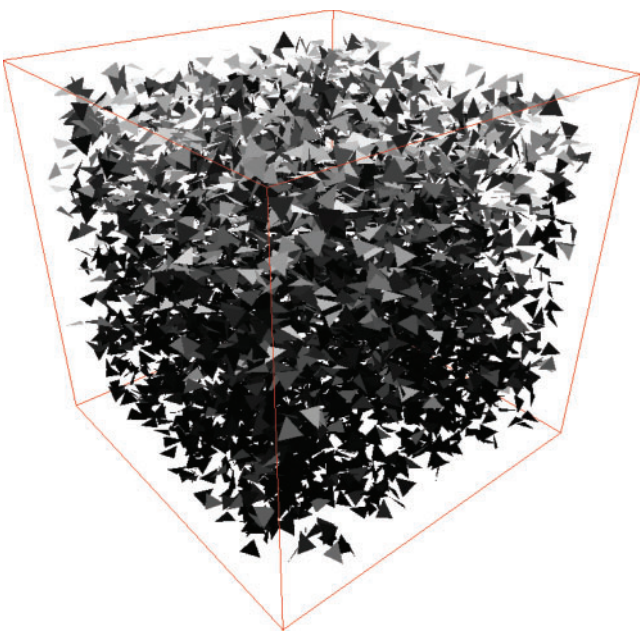

Fig. 3. 3-D canopy mock-up (triangle mix), used to test *QuasiMC*. The canopy consists of 5000 triangles uniformly distributed within a cube. In the example shown, the scene was illuminated by directional light from above. Each triangle is shaded according to the amount of absorbed energy, with lighter shades representing high radiant energy and darker shades representing low radiant energy.

model was applied. The scene was illuminated from above, using a directional light source. All experiments involved the same number of light paths, 262 144. To estimate the error in the results, the experiments were repeated 10 times for both the MC and RQMC methods. In either case, the computation time was ~30 s per experiment on a 3.0 GHz computer.

The results are shown in Fig. 4. We fitted exponential curves to the experimental data, because the attenuation of light in a turbid medium, which can be considered as a continuous approximation of our triangle mix, is described by an exponential equation: the Beer-Lambert law (Thornley and Johnson 1990). Both methods yield closely matching mean values (compare Fig. 4*a* with Fig. 4*b*), but the RQMC method produced significantly smaller variance (compare Fig. 4*c* with Fig. 4*d*). The mean variance over all 5000 triangles in the MC case was 8.06×10^{-2} with standard deviation 15.29×10^{-2}, and in the RQMC case it was 0.75×10^{-2} with standard deviation 1.14×10^{-2}. Thus, RQMC makes it possible to achieve a better accuracy than MC using the same number of ray paths. Alternatively RQMC can be used to achieve the same accuracy as MC with a smaller number of light paths. For our mix of 5000 triangles, the RQMC method required approximately one-quarter of the number of MC rays to achieve approximately the same accuracy (for the RQMC method with 65 536 rays, the mean variance was 5.4×10^{-2} with standard deviation 8.86×10^{-2}).

To evaluate the MC and RQMC methods when computing the distribution of light with multiple wavelengths, we calculated the red/far red ratio (R/FR) using four scenarios: treating each wavelength separately using MC or RQMC, and representing both wavelengths in a single ray (correlated sampling) using MC or RQMC. In this case, the comparison is made within our *QuasiMC* program because *MonteCarlo* does not fully support correlated sampling. Following Gautier *et al*. (2000) we assumed reflectance of 5.3% for red light and 42.6% for far red light, and transmittance of 2% and 40.5%, correspondingly. As in the previous test, we assumed that rays could be scattered with equal probability in any direction. The initial value of the R/FR ratio for the incoming light (directional light source illuminating the scene from above) was set to 1.2 (Chelle *et al*. 2007). The number of rays was 2 097 152 for the experiments with one wavelength per ray and 1 048 576 for the experiments with correlated sampling.

Figure 5 shows the results for the case where separate rays were used. Each graph compares the results obtained with the MC

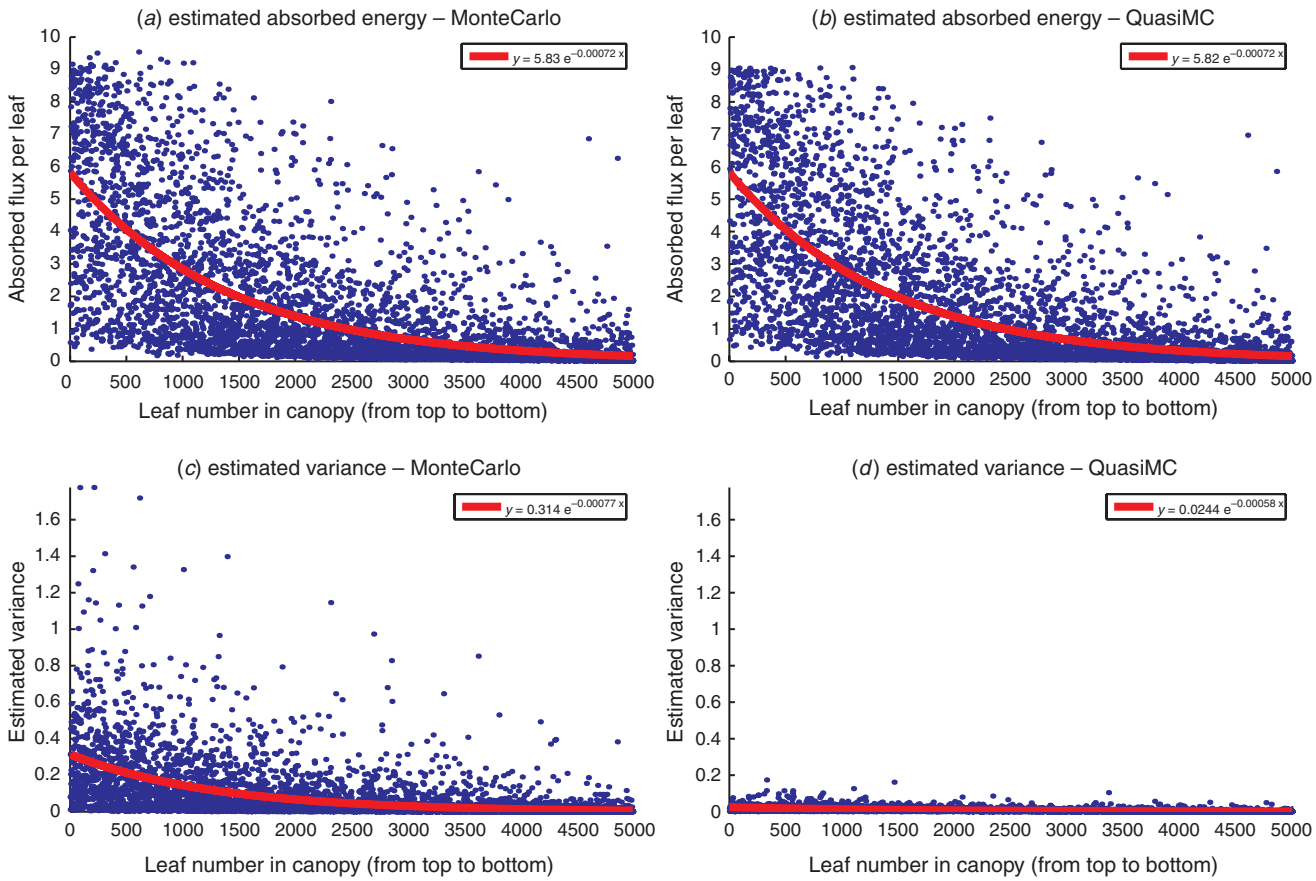

Fig. 4. Comparison of the absorbed energy per leaf in the plant canopy model from Fig. 3. The absorbed energy per leaf is a mean estimated over 10 simulations. The curves described by equations of the form $y = Ae^{-\lambda x}$ were fitted to these mean values. Parameters A and λ were estimated by minimising the sum-of-squares error for $\sum_i Ae^{-\lambda x_i} - y_i$, where i is the leaf number. (*a*) Monte Carlo estimate, (*b*) Randomised quasi-Monte Carlo estimate, (*c*) estimated variance in Monte Carlo calculation, and (*d*) estimated variance in randomised quasi-Monte Carlo calculation.

and RQMC methods for individual triangles (thus, if the results for some triangle are identical, the corresponding point will lie on the diagonal line). Table 1 summarises our statistical analysis of these results. The mean values of the incident irradiance for red and far red light, and the R/FR ratio, obtained using both methods are similar, but the variances of the values returned by the RQMC method are smaller than for the MC method. There is a strong positive linear correlation between values obtained using MC and RQMC for all three variables under consideration (red, far red, and R/FR ratio).

Figure 6 shows the corresponding results for correlated sampling. The mean incident irradiance for red and far red light, and the sample variance of these values match closely those for separate rays. The mean values of the R/FR ratio are also similar (Table 1), which implies that the correlated sampling method did not introduce any error into the computation. In contrast, the variance in the values of the R/FR ratio is reduced, compared with uncorrelated sampling. Thus we conclude that correlated sampling does indeed reduce the variance in calculating ratios in the radiative energy associated with different wavelengths, and, in particular, in calculating the R/FR ratio. At the same time, it halves the number of rays needed when using uncorrelated sampling.

Comparison of QuasiMC and CARIBU

To further test *QuasiMC*, we compared it with CARIBU (Chelle *et al*. 2004), an independently developed program that calculates the distribution of light energy using the radiosity method (Goral *et al*. 1984). We used the virtual canopy model from our previous comparison (with identical leaf R/FR optical properties) and used the Lambertian model in *QuasiMC* to match the one used in CARIBU. The results for absorbed energy per leaf of red and far red light are shown in Fig. 7 (including the energy contribution from indirect light only).

Table 2 presents the statistics we collected from our comparison of *QuasiMC* and CARIBU. Generally, the values of the absorbed radiant energy per leaf returned using RQMC path tracing agree with those obtained using radiosity. *QuasiMC* computed the light distribution in ~3.5 min for 1 048 576 rays with 10 randomisations, and CARIBU took ~10 min on the same computer. The mean variance over all triangles (reported in Table 2) indicates that a sufficient number of rays was traced for this comparison to be valid. However, the interpretation of these results is complicated by different assumptions underlying each program. CARIBU was designed for computing light distribution in infinite canopies using the nested radiosity

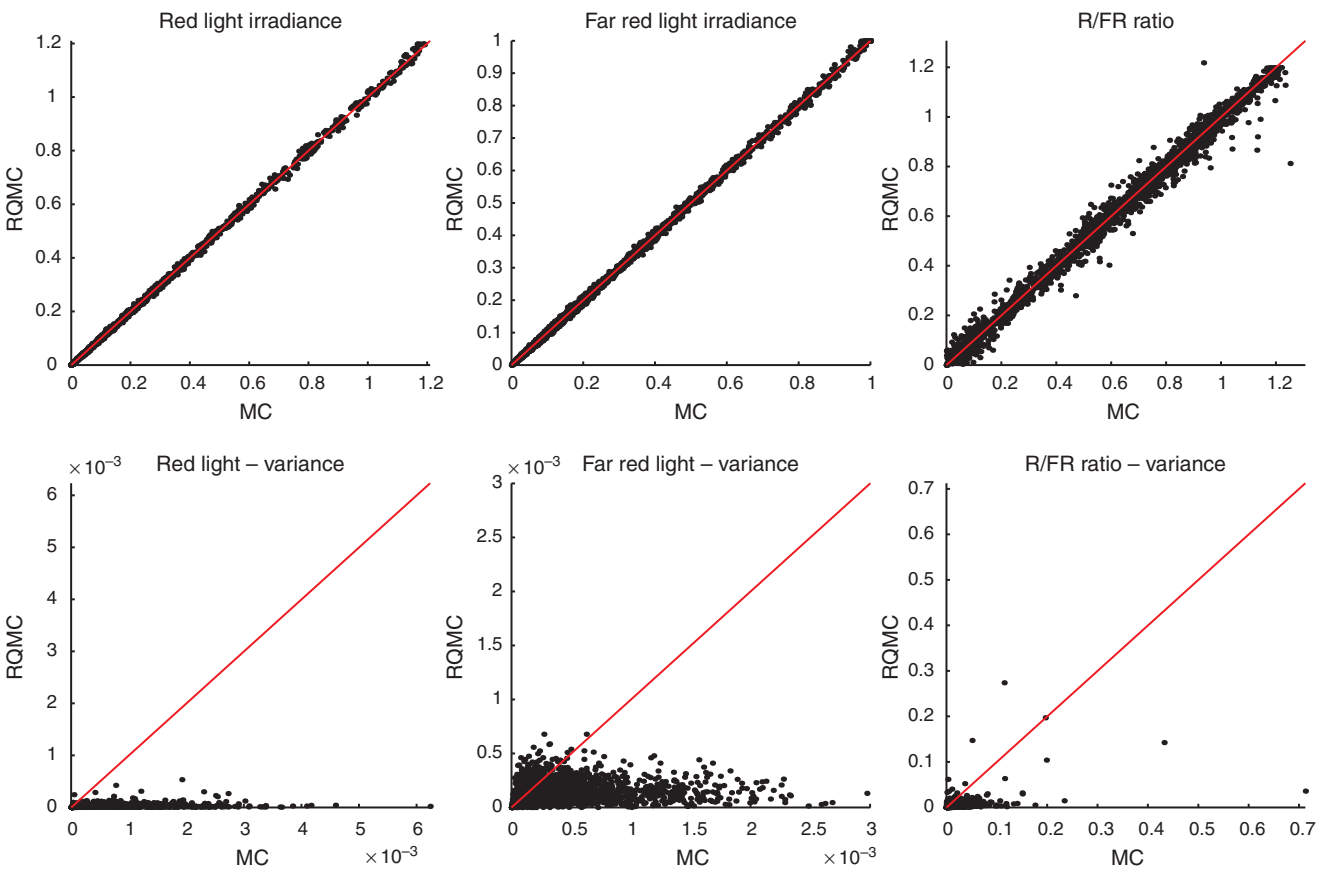

Fig. 5. The top row is a comparison of the RQMC method and the MC method for computing the incident irradiance per leaf for red and far red light, and for the R/FR ratio. The bottom row compares the sample variance of the two methods. In this case, the information of each wavelength is carried in a separate ray. Each point in the graph represents the relevant value for a single triangle.

Table 1. Statistical analysis of incident irradiance for red and far red light, and R/FR ratio per triangle calculated using MC path tracing and RQMC path tracing

The means and sample variances are given, as well as the coefficient of determination (r^2) and the linear regression with slope (a) and intercept (b) for separate rays and one ray carrying information for both wavelengths

	Mean irradiance MC	Mean variance MC	Mean irradiance RQMC	Mean variance RQMC	a	b	r^2
Red light (separate/one ray)	0.0983	1.79e-4	0.0985	7.22e-6	0.9975	6.88e-5	0.9998
	0.0983	1.82e-4	0.0984	7.01e-6	0.9977	8.94e-5	0.9998
Far red light (separate/one ray)	0.1572	2.12e-4	0.1574	8.81e-5	0.9978	2.22e-4	0.9996
	0.1572	2.13e-4	0.1574	8.86e-5	0.9975	2.31e-4	0.9997
R/FR ratio (separate/one ray)	0.2679	4.53e-3	0.2670	1.23e-3	1.0026	2.39e-4	0.9976
	0.2640	1.15e-3	0.2650	5.59e-4	0.9982	–4.87e-4	0.9992

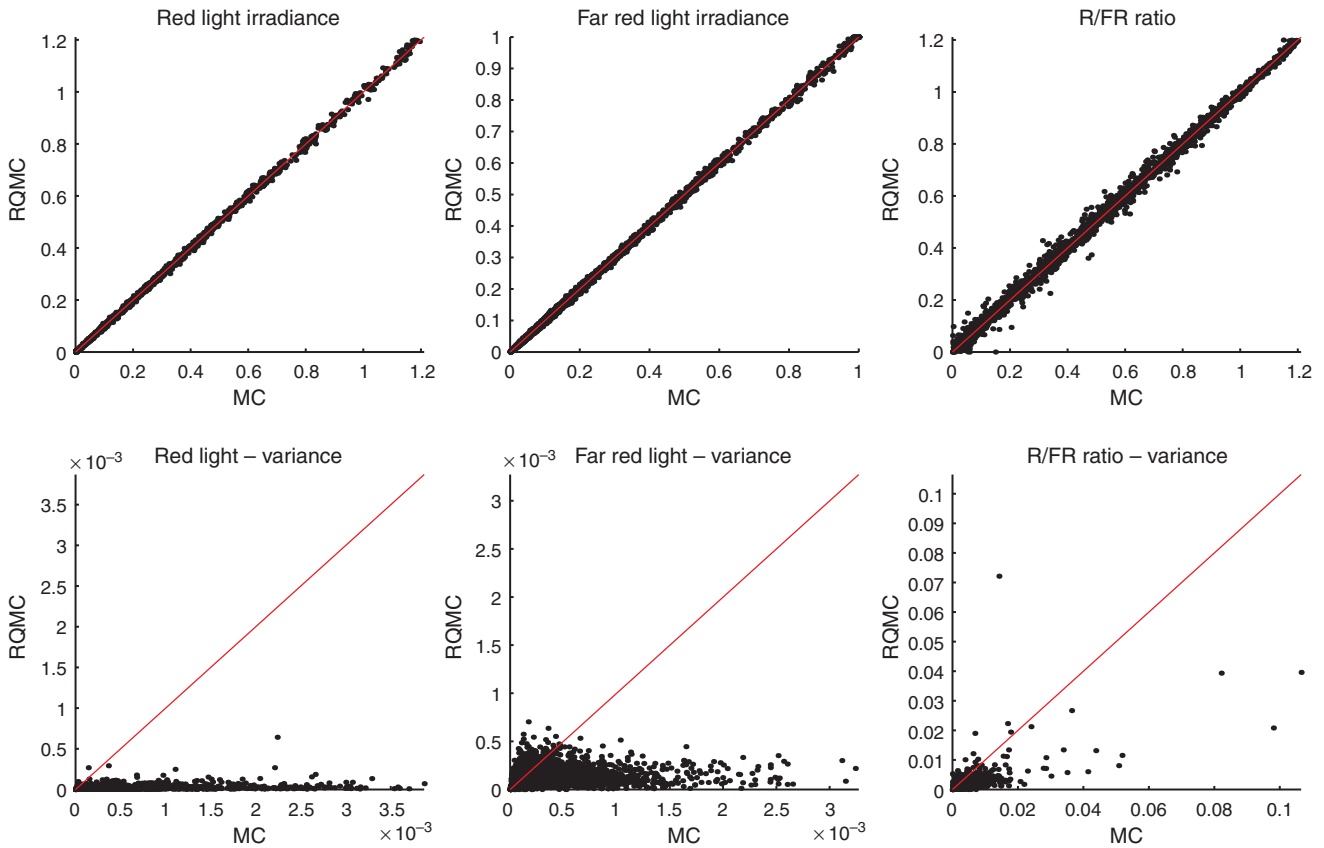

Fig. 6. A similar computation as shown in Fig. 5 but using correlated sampling, where one ray carries information for both red and far red light.

method, and is less efficient when computing light in finite canopies with the classic radiosity method. Infinite canopies, however, are not supported by the current *QuasiMC* implementation, which makes it impossible to compare nested-radiosity with RQMC.

Application example: light distribution in virtual kiwifruit

We applied *QuasiMC* to calculate the light environment in a model of the annual growth cycle of a managed mature kiwifruit vine (Cieslak *et al*. 2007). Calculations were conducted for a plant 100 days after budbreak. Each leaf was represented as a single-patch Bezier surface (Watt 2000), approximated using 72 triangles (as a compromise between the accuracy and complexity of representation). There were 1063 leaves on the vine, amounting to the total of 76 536 triangles, plus one rhombus representing the ground. Leaves were assumed to have a preference for approximately horizontal orientations, but no experimental data were used to quantify these orientations. In each simulation, 1 048 576 rays were traced using the Korobov method. The simulations were randomised and repeated 10 times to estimate variance.

Greer and Laing (1992) determined that the absorbance ratio of kiwifruit leaves is ~78% for the PAR domain. Based on this

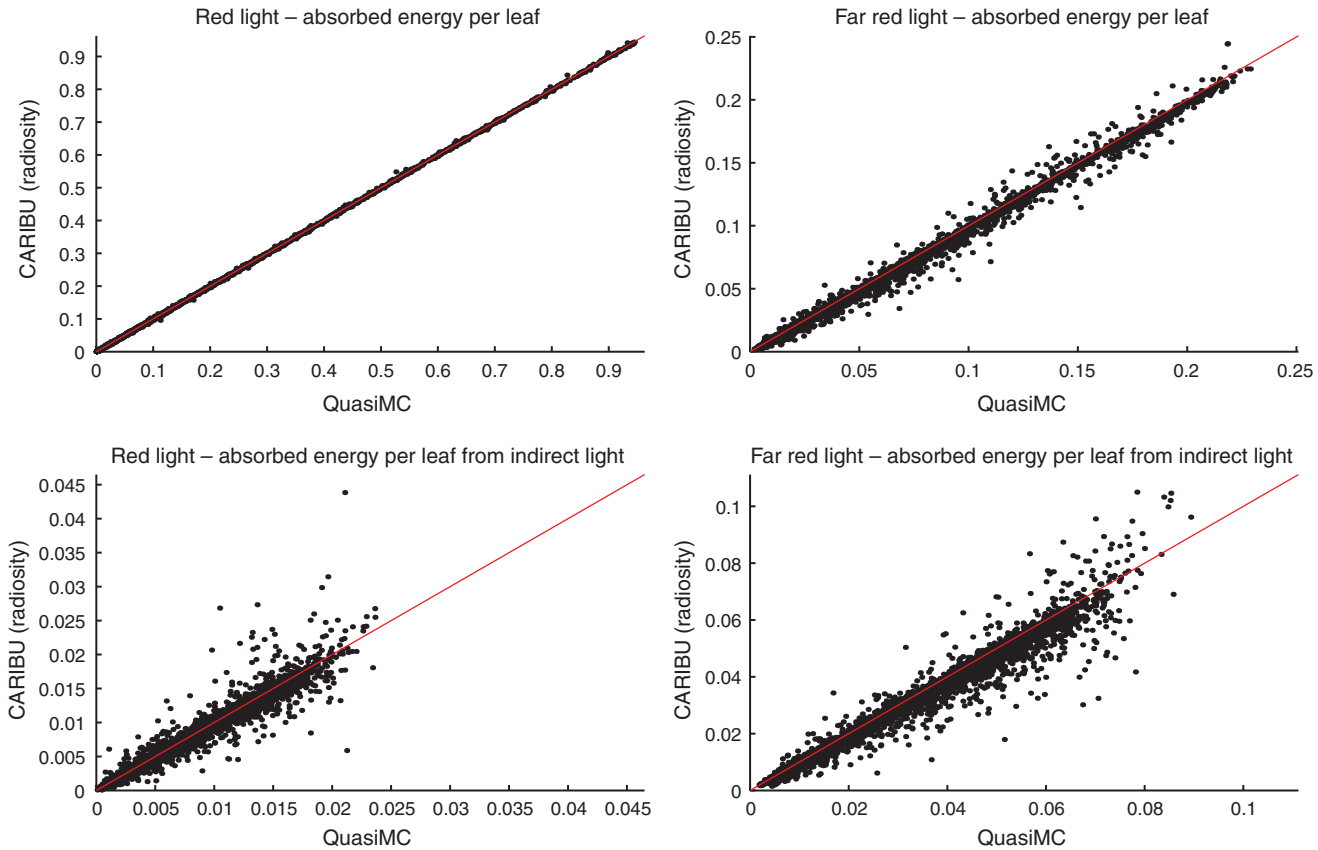

Fig. 7. Comparison of red and far red absorbed radiant energy per leaf computed by the RQMC path tracing algorithm and the classic radiosity algorithm. The two lower graphs are a comparison of the energy contribution from indirect light only.

Table 2. Statistical analysis of absorbed red and far red energy per triangle, with contribution from indirect light, as computed by the *QuasiMC* and CARIBU programs

The variables are mean absorbed energy per triangle, mean variance from the *QuasiMC* estimate, the slope (*a*) and intercept (*b*) of the linear regression, and the coefficient of determination (r^2)

	Mean absorbed energy (*QuasiMC*)	Mean sample variance (*QuasiMC*)	Mean absorbed energy (CARIBU)	*a*	*b*	r^2
Red light	0.1545	9.16e-6	0.1542	0.9987	–1.00e-4	0.9999
Far red light	0.0568	6.02e-6	0.0552	0.9845	–6.60e-4	0.9977
Indirect red light	0.0060	6.86e-7	0.0059	0.9731	4.52e-5	0.9754
Indirect far red light	0.0297	5.72e-6	0.0282	0.9481	3.86e-5	0.9819

value, we assumed leaf reflectance of 11% and transmittance of 11%. Furthermore, we set the ground reflectance to 20%. We considered four light conditions: directional light source placed above the canopy, clear sky over the entire day, overcast sky, and midday sun (average between 1100–1300 hours). For the three sky light conditions, we assumed the latitude of 38°S (Bay of Plenty, New Zealand). In all cases, one wavelength (450 nm) was used to represent the PAR domain.

Distribution of the radiant energy absorbed by the leaves for two of the above conditions is visualised in Fig. 8. Fig. 9 shows the frequency distributions of the leaves absorbing different amounts of radiant energy for all four light conditions. The time needed to estimate the light distribution and sample variance for each light condition was about 8 min on a 3.0 GHz computer.

The frequency distribution for absorbed energy per leaf under direct light and the approximation of the sky at midday are not similar, because the direction of incoming light is fixed under the first case but not the second. The mean and standard deviation of radiant energy absorbed by the leaves are 0.35 ± 0.24 for direct light and 0.24 ± 0.1 for midday. The difference in means is related to the direction of incoming rays and leaf orientation. In this case, since the leaves are mostly planar, they absorb more light from a light source directly above the plant canopy than from a source

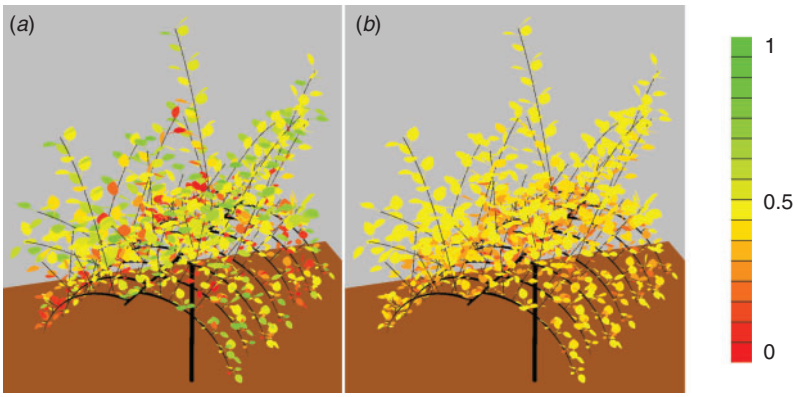

Fig. 8. Virtual kiwifruit vine with leaves coloured according to the amount of absorbed radiant energy for each leaf, normalised to the range [0,1]. (*a*) The light source is directly above the vine. (*b*) The light source is an approximation of the sky on a clear day.

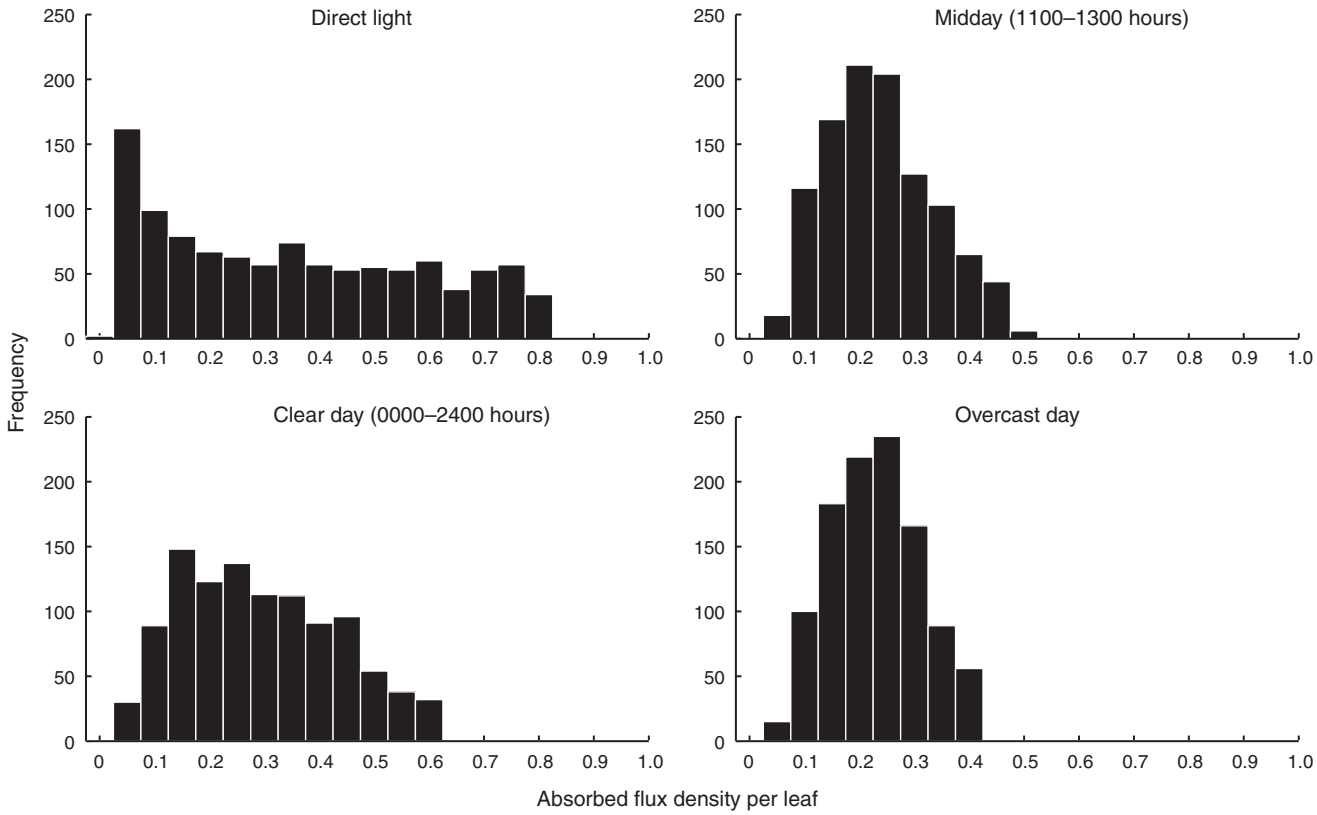

Fig. 9. The frequency distributions of the absorbed radiant energy per leaf of a virtual kiwifruit vine 100 days after budbreak under four light conditions: a light source directly above the vine, a midday sun (1100–1300 hours), a day with clear skies and an overcast day.

away from zenith. The difference in standard deviations is likely due to the more uniform distribution of incoming light when the sky model is used.

The frequency distribution of absorbed energy under the CIE overcast sky model resembles the distribution for the midday sky, because the overcast sky model has the brightest region at zenith. The mean and standard deviation of the absorbed energy for the overcast sky is 0.23 ± 0.08, which is similar to that observed under the midday sky. The means under these two conditions have nearly the same value, because *QuasiMC* normalises the absorbed energy so that the flux density of a flat surface above the canopy is 1 $W.m^{-2}$.

Finally, the mean and standard deviation of the absorbed radiant energy per leaf for the sky on a clear day is 0.29 ± 0.14.

There is an increase in the variance when compared with midday, because the time over which the direction of incoming rays is averaged is increased (24 *v.* 2 h).

Conclusion

We introduced randomised quasi-Monte Carlo path tracing as a method for computing light distribution in a plant canopy. The method makes it possible to effectively simulate both the distribution of monochromatic light and the distribution of light composed of different wavelengths; the latter is important, for example, when calculating R/FR ratios. The randomised quasi-Monte Carlo and the original Monte Carlo path tracing have been implemented in our light simulation program *QuasiMC*. The program makes it possible to simulate plant–light interaction at the level of individual organs, and can return the incident irradiance and the amount of radiant energy absorbed by these organs in the course of a plant's development. Both the computation of the distribution of monochromatic light and of light of different wavelengths are supported. *QuasiMC* has been designed to work with the plant simulators *cpfg* and *lpfg* within the L-studio/Virtual Laboratory modelling platforms (Prusinkiewicz 2004), but can also be used with other plant simulators, as long as they employ the same protocol for communicating with the plant environment. The organisation of communication between *QuasiMC* and *cpfg*/*lpfg* has been presented using sample L-system code. The operation of *QuasiMC* was illustrated and analysed using an abstract virtual canopy (a triangle mix) and a model of kiwifruit as examples.

Comparing randomised quasi-Monte Carlo path tracing from *QuasiMC* with the Monte Carlo path tracing from *MonteCarlo* (Měch 1997), we found that RQMC makes it possible to reduce the number of rays, and therefore the computation time, as much as four times. This reduction applies to both the computation of the absorbed radiant energy and the incident irradiance of monochromatic light, and the computation of R/FR ratios. We also confirmed that correlated sampling is an effective variance reduction technique for computing the R/FR ratio, regardless of whether MC or RQMC method is used. To additionally test *QuasiMC*, we compared it with CARIBU, an independently developed program for calculating light distribution using the radiosity method (Chelle *et al*. 2004). The comparison showed an agreement of the results, with *QuasiMC* using less computation time.

To illustrate the operation of *QuasiMC* in a practical modelling setting, we applied it to compare distributions of absorbed energy in a virtual kiwifruit canopy under various light conditions. This comparison captured the effect of weather and time of day on the incident radiation of leaves.

Acknowledgements

We thank Michael Chelle for help with operation of his radiosity program and useful discussions, and Alla Seleznyova for help with the construction of the kiwifruit model. We also gratefully acknowledge the support of this research by the Natural Sciences and Engineering Research Council of Canada (MC, CL, and PP), the Horticulture and Food Research Institute of New Zealand Limited (MC) and the ARC Centre for Complex Systems at the University of Queensland (MC).

References

Arvo J, Kirk D (1990) Particle transport and image synthesis. *Computer Graphics* **24**, 63–66. doi: 10.1145/97880.97886

Bousquet L, Lachérade S, Jacquemoud S, Moya I (2005) Leaf BRDF measurements and model for specular and diffuse components differentiation. *Remote Sensing of Environment* **98**, 201–211. doi: 10.1016/j.rse.2005.07.005

Breece H, Holmes R (1971) Bidirectional scattering characteristics of healthy green soybean and corn leaves *in vivo*. *Applied Optics* **10**, 119–127.

Chelle M, Andrieu B (1998) The nested radiosity model for the distribution of light within plant canopies. *Ecological Modelling* **111**, 75–91. doi: 10.1016/S0304-3800(98)00100-8

Chelle M, Andrieu B (2007) Modelling the light environment of virtual crop canopies. In 'Functional–structural plant modelling in crop production'. (Eds J Vos, L Marcelis, P de Visser, P Struik, J Evers) pp. 75–89. (Springer-Verlag: Berlin)

Chelle M, Andrieu B, Bouatouch K (1998) Nested radiosity for plant canopies. *The Visual Computer* **14**, 109–125. doi: 10.1007/s003710050127

Chelle M, Hanan J, Autret H (2004) Lighting virtual crops: the CARIBU solution for open L-Systems. In 'Proceedings of the 4th international workshop on functional–structural plant models'. p. 194. (UMR AMAP: Montpellier, France)

Chelle M, Evers J, Combes D, Varlet-Grancher C, Vos J, Andrieu B (2007) Simulation of the three-dimensional distribution of the red : far-red ratio within crop canopies. *New Phytologist* **176**, 223–234. doi: 10.1111/j.1469-8137.2007.02161.x

Cici S, Adkins S, Hanan J (2008) A canopy architectural model to study the competitive ability of chickpea and sowthistle. *Annals of Botany* **101**, 1311–1318. doi: 10.1093/aob/mcn040

CIE-110 (1994) 'Spatial distribution of daylight – luminance distributions of various reference skies.' (Commission Internationale de l'Eclairage: Vienna)

Cieslak M (2004) 'QuasiMC user's manual.' (University of Calgary: Calgary, Canada)

Cieslak M, Seleznyova AN, Hanan J (2007) Virtual kiwifruit: modelling annual growth cycle and light distribution. In 'Proceedings of the 5th international workshop on functional–structural plant models' pp. 46-1–46-2. (Print Solutions Hawke's Bay Limited: Napier, New Zealand)

Cranley R, Patterson T (1976) Randomization of number theoretic methods for multiple integration. *SIAM Journal on Numerical Analysis* **13**, 904–914. doi: 10.1137/0713071

Disney M, Lewis P, North P (2000) Monte Carlo ray tracing in optical canopy reflectance modelling. *Remote Sensing Reviews* **18**, 163–196.

Evans G, McCool M (1999) Stratified wavelength clusters for efficient spectral Monte Carlo rendering. In 'Proceedings of the graphics interface 1999 conference'. pp. 42–49. (Canadian Human-Computer Communications Society: Kingston, CA)

Faure H (1982) Discrépance de suites associées à un système de numération. *Acta Arithmetica* **61**, 337–351.

Gautier H, Měch R, Prusinkiewicz P, Varlet-Grancher C (2000) 3D Architectural modelling of aerial photomorphogenesis in white clover (*Trifolium repens* L.) using L-systems. *Annals of Botany* **85**, 359–370. doi: 10.1006/anbo.1999.1069

Goral C, Torrance K, Greenberg D, Battaile B (1984) Modeling the interaction of light between diffuse surfaces. *Computer Graphics* **18**, 213–222. doi: 10.1145/964965.808601

Govaerts Y (1996) A model of light scattering in three-dimensional plant canopies: a Monte Carlo ray tracing approach. PhD thesis, Université Catholique de Louvain, Louvain-la-Neuve, Belgium.

Greer D, Laing W (1992) Photoinhibition of photosynthesis in intact kiwifruit (*Actinidia deliciosa*) leaves: changes in susceptibility to photoinhibition and recovery during the growth season. *Planta* **186**, 418–425. doi: 10.1007/BF00195323

Halton JH (1960) On the efficiency of certain quasi-random sequences of points in evaluating multi-dimensional integrals. *Numerische Mathematik* **2**, 84–90. doi: 10.1007/BF01386213

Jacquemoud S, Ustin S (2001) Leaf optical properties: a state of the art. In 'Proceedings 8th international symposium physical measurements and signatures in remote sensing'. pp. 223–232. (CNES: Aussois, France)

Jensen H (2001) 'Realistic image synthesis using photon mapping.' (A K Peters: Natick, MA, USA)

Kajiya JT (1986) The rendering equation. *Computer graphics (SIGGRAPH '86 conference proceedings)* **20**, 143–150.

Karwowski R, Prusinkiewicz P (2004) The L-system-based plant-modelling environment L-studio 4.0. In 'Proceedings of the 4th international workshop on functional–structural plant models'. pp. 403–405. (UMR AMAP: Montpellier, France)

Karwowski R, Lane B (2007) 'LPFG User's manual.' (University of Calgary: Calgary, Canada)

Keller A (1996) Quasi-Monte Carlo methods in computer graphics: the global illumination problem. *Lectures in Applied Mathematics* **32**, 455–469.

Korobov NM (1959) The approximate computation of multiple integrals. *Doklady Akademii Nauk SSSR* **124**, 1207–1210.

L'Ecuyer P (1999) Good parameters and implementations for combined multiple recursive random number generators. *Operations Research* **47**, 159–164.

L'Ecuyer P, Lemieux C (2002) Recent advances in randomized quasi-Monte Carlo methods. In 'Modeling uncertainty: an examination of stochastic theory, methods, and applications'. (Eds M Dror, P L'Ecuyer, F Szidarovszky) pp. 419–474. (Kluwer Academic Publishers: Boston)

Měch R (1997) Modeling and simulation of the interaction of plants with the environment using L-systems and their extensions. PhD thesis, University of Calgary, Calgary, Canada.

Měch R, Prusinkiewicz P (1996) Visual models of plants interacting with their environment. In 'SIGGRAPH '96: Proceedings of the 23rd annual conference on computer graphics and interactive techniques'. pp. 397–410.

Měch R, James M, Hammel M, Hanan J, Prusinkiewicz P, Karwowski R, Lane B (2005) 'CPFG Vesion 4.0 user's manual.' (University of Calgary: Calgary, Canada)

Owen AB (2003) Quasi-Monte Carlo sampling. In 'Monte Carlo tay tracing: SIGGRAPH 2003 Course 44'. (Ed. HW Jensen) pp. 69–88. (SIGGRAPH)

Prusinkiewicz P (2004) Art and science for life: designing and growing virtual plants with L-systems. *Acta Horticulturae* **630**, 15–28.

Prusinkiewicz P, Hanan J, Hammel M, Měch R (1997) L-systems: from the theory to visual models of plants. In 'Plants to ecosystems. Advances in computational life sciences'. (Ed. MT Michalewicz) pp. 1–27. (CSIRO Publishing: Melbourne)

Prusinkiewicz P, Karwowski R, Měch R, Hanan J (2000) L-studio/*cpfg*: a software system for modeling plants. In 'Proceedings of applications of graph transformations with industrial relevance'. pp. 457–464. (Springer-Verlag: Berlin)

Room P, Hanan J, Prusinkiewicz P (1996) Virtual plants: new perspectives for ecologists, pathologists and agricultural scientists. *Trends in Plant Science* **1**, 33–38. doi: 10.1016/S1360-1385(96)80021-5

Ross J, Marshak A (1988) Calculation of canopy bidirectional reflectance using the Monte Carlo method. *Remote Sensing of Environment* **24**, 213–225. doi: 10.1016/0034-4257(88)90026-0

Rubinstein R (1981) 'Simulation and the Monte Carlo method.' (John Wiley & Sons Inc.: New York)

Shirley P, Wang Y (1992) Distribution ray tracing: theory and practice. In 'Proceedings of the 3rd Eurographics rendering workshop'. pp. 33–44. (Consolidation Express: Bristol, England)

Shirley P, Ashikhmin M, Gleicher M, Marschner S, Reinhard E, Sung K, Thompson W, Willemsen P (2005) 'Fundamentals of computer graphics.' (A K Peters: Wellesley, MA, USA)

Sobol' IM (1967) The distribution of points in a cube and the approximate evaluation of integrals. *U.S.S.R. Computational Mathematics and Mathematical Physics* **7**, 86–112. doi: 10.1016/0041-5553(67)90144-9

Soler C, Sillion F, Blaise F, de Reffye P (2003) An efficient instantiation algorithm for simulating radiant energy transfer in plant models. *ACM Transactions on Graphics* **22**, 204–233. doi: 10.1145/636886.636890

Spanier J, Gelbard E (1969) 'Monte Carlo principles and neutron transport problems.' (Addison-Wesley: Reading, MA, USA)

Thornley JHM, Johnson IR (1990) 'Plant and crop modelling.' (Oxford University Press: Oxford)

Veach E (1997) Robust Monte Carlo methods for light transport simulation. PhD thesis, Stanford University, Stanford, CA, USA.

Watt A (2000) '3D Computer graphics.' (Addison-Wesley: Reading, MA, USA)

Manuscript received 17 March 2008, accepted 22 September 2008