Huaxiong Huang (York University)
Emmanuel Lorin (Carleton University)
Odile Marcotte (CRM et UQAM)

Comptes rendus du Quatrième atelier de résolution de problèmes industriels de Montréal Une activité CRM-Mprime

Proceedings of the Fourth Montreal Industrial Problem Solving Workshop A CRM-Mprime Event

CRM-3327

Preface

Le Quatrième atelier de résolution de problèmes industriels de Montréal eut lieu du 15 au 19 août 2011 et accueillit plus d'une cinquantaine de participants, incluant des représentants d'entreprises, des professeurs, des étudiants et des stagiaires postdoctoraux. Le lecteur pourra trouver les descriptions des problèmes sur le site de l'atelier (www.crm.umontreal.ca/probindustriels2011/). Le CRM est très reconnaissant aux entreprises qui lui ont fourni des problèmes pour l'atelier. Il s'agit des compagnies Acculogic, Hatch, Matrox et Turquoise Technology Solutions et de l'Institut national d'optique (qui a fourni deux problèmes). Le CRM est aussi très reconnaissant au réseau de centres d'excellence Mprime, qui a assumé les dépenses reliées à l'atelier. Pour la quatrième fois au CRM, cet atelier a servi d'incubateur à des collaborations fructueuses entre des chercheurs universitaires et des entreprises.

À titre personnel, j'aimerais remercier Huaxiong Huang et Emmanuel Lorin d'avoir accepté d'être corédacteurs de ces comptes rendus, et le personnel du CRM (en particulier Sakina Benhima, Louise Letendre, Suzette Paradis et Louis Pelletier) d'avoir contribué efficacement à l'organisation de l'atelier. André Montpetit me fournit une aide précieuse pour la mise en forme de ces comptes rendus. Finalement j'aimerais remercier de leurs conseils et de leur appui les coordonnateurs des équipes et les membres du comité d'organisation de l'atelier: Fabian Bastin, Eliot Fried, Michel Gendreau, Huaxiong Huang, Christian Léger, Emmanuel Lorin, Dominique Orban, Sylvain Perron, Jean-François Plante, Jean-Marc Rousseau, Louis-Martin Rousseau et Luc Vinet.

The Fourth Montreal Industrial Problem Solving Workshop took place from August 15 to 19, 2011, and was attended by more than fifty participants, including industry representatives, professors, students, and postdoctoral fellows. The reader may find a description of the problems on the workshop web site (www.crm.umontreal.ca/probindustriels2011/index_e.php). The CRM is very grateful to the companies that provided the problems examined during the workshop: Acculogic, Hatch, Matrox, Turquoise Technology Solutions, and INO (Institut national d'optique), which provided two problems. The CRM is also very grateful to

Mprime (a Network of Centres of Excellence), which defrayed the expenses related to the workshop. Once more the workshop served as an "incubator" of collaborative projects between companies and university researchers.

I am personally grateful to Huaxiong Huang and Emmanuel Lorin for having agreed to be coeditors of these proceedings, and to the CRM personnel (especially Sakina Benhima, Louise Letendre, Suzette Paradis and Louis Pelletier) for having helped me organize the workshop. André Montpetit's help in formatting the proceedings was invaluable. For their advice and support, I would like to thank the team coordinators and the members of the Workshop Organizing Committee: Fabian Bastin, Eliot Fried, Michel Gendreau, Huaxiong Huang, Christian Léger, Emmanuel Lorin, Dominique Orban, Sylvain Perron, Jean-François Plante, Jean-Marc Rousseau, Louis-Martin Rousseau, and Luc Vinet.

Montréal, le 31 mars 2013 *Odile Marcotte*

Contents

List of Contributors

Reynaldo Arteaga
Carleton University

Geneviève Autotte
Université de Montréal

Etienne Ayotte-Sauvé
Natural Resources Canada, Varennes (Québec)

Guillaume Blanchet
Université de Montréal

Gilles Boesch
Énergies renouvelables, Hatch

C. Sean Bohun
University of Ontario Institute of Technology

Michel Carreau
Énergies renouvelables, Hatch

Omar Cedrati
HEC Montréal

Stephen Y. Chen
School of Information Technology, York University

Marina Chugunova
Department of Mathematics, University of Toronto

Beatriz Cortes
Tecnològico de Monterrey (Mexico)

A. Cournoyer
Institut national d'optique

Ellis Cumberbatch
Claremont Graduate University

Simona Corina Curtean
Université du Québec à Montréal

Michel Denault
HEC Montréal

François Fillion-Gourdeau
CRM, Université de Montréal

Alistair Fitt
Oxford Brookes University

Ludovic Gagnon
Université Laval

Michel Gendreau
École Polytechnique de Montréal

Veronica Gheorghiade
University of Guelph

Ozgur Gurtuna
Turquoise Technology Solutions

Huaxiong Huang
York University

Christian Léger
Université de Montréal

Claude Le Bris
École Nationale des Ponts et Chaussées, Paris

Buyang Li
City University of Hong Kong

Arnaud Lina
Matrox, Montréal (Québec)

E. Lorin
Carleton University

Apala Majumdar
Mathematical Institute, University of Oxford

Shahrouz Mirzaalizadeh
École Polytechnique de Montréal

Mostafa Nasri
Université de Montréal

Dominique Orban
GERAD and École Polytechnique de Montréal

Jean-François Plante
HEC Montréal

Mason A. Porter
Oxford University

Chris Prior
Mathematical Institute, University of Oxford

Léonard Ryo Morin
Université de Montréal

Jean-Christophe Roy
École Polytechnique de Montréal

Suzanne Shontz
Pennsylvania State University

Babak Soltannia
McGill University

Hervé Tchouake Tchuiguep
étudiant, Université de Montréal

Rex Westbrook
University of Calgary

Vincent Zalzal
Matrox, Montréal (Québec)

Modelling of CO_2 Laser Polishing of Glass

C. Sean Bohun, Ellis Cumberbatch, Alistair Fitt, Huaxiong Huang, Buyang Li, Babak Soltannia, and Rex Westbrook

1 Problem Background

Ultrafast lasers have a huge potential as a powerful tool for material processing (milling and polishing) at the micro-scale, thanks to their increased controllability and the reduced residual damage to the workpiece (see the 2004 article by Stoian et al. entitled "Adaptive optimizatino in ultrafast laser material processing", in High-Power Laser Ablation V, *Proceedings of SPIE*, Vol. 5548, SPIE, Bellingham, WA). INO has developed a pulse-shaping fiber laser platform for designing complex micro-machining processes. The power of the laser must be carefully controlled to prevent vapourization of the materials (metal or glass). In addition, the workpiece surface often needs to be polished after milling (ablation). Alain Cournoyer from INO came to the Fourth Montreal Industrial Problem Solving Workshop, held at the CRM in August 2011, in order to ask whether mathematical models could be developed to investigate the ablation and polishing of a glass workpiece.

Huaxiong Huang
Coordinator, York University

C. Sean Bohun
University of Ontario Institute of Technology

Ellis Cumberbatch
Claremont Graduate University

Alistair Fitt
Oxford Brookes University

Buyang Li
Student, City University of Hong Kong

Babak Soltannia
Student, McGill University

Rex Westbrook
University of Calgary

With respect to the polishing process, a number of investigative priorities were set out by the industrial representative at the beginning of the workshop and these aims guided the efforts. In particular, the team members attempted to answer the following questions:

- Which size (length scale) of surface roughness could be smoothed out by the process?
- What are the main features of the heat transfer problem and how does the laser heat the material?
- Can one propose a predictive model for the smoothing of surface roughness?
- When might we expect cracking to occur and can we predict this?

In this report, we focus only on the surface polishing aspect; the ablation process is being addressed by another report within the current proceedings.

2 Model

We consider a workpiece of glass with surface dimensions of the order of a few centimeters and depth of a few tenths of cm. To polish the workpiece surface, a laser is used to heat the workpiece rapidly (from the top) until the workpiece reaches the melting temperature (at which point the molten material begins to flow, as displayed in Fig. 1).

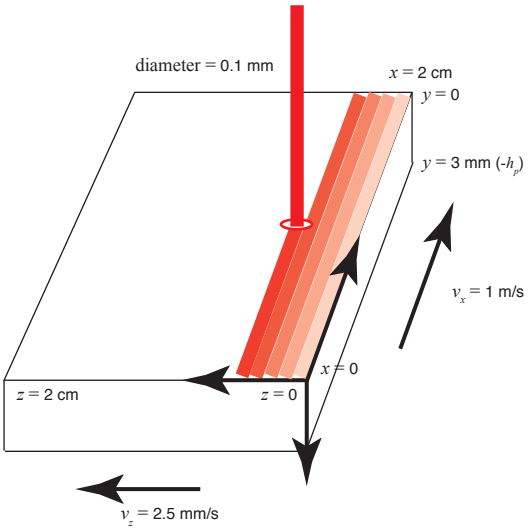

Fig. 1 Schematics of the surface polishing process and the workpiece dimensions.

We assume that the glass can be modelled as an incompressible viscous fluid with temperature-dependent viscosity. The governing equations are the continuity and Navier-Stokes equations:

$$\rho(\mathbf{v}_t + \mathbf{v}\cdot\nabla\mathbf{v}) = -\nabla\cdot\sigma + \mathbf{g}, \tag{1}$$

$$\nabla\cdot\mathbf{v} = 0, \tag{2}$$

where the viscous stress tensor σ is given by $\sigma = -p\mathbf{I} + \mu(\nabla\mathbf{v} + \nabla\mathbf{v}^T)/2$. Here ρ, $\mathbf{v}$, and p are the density, velocity, and pressure (respectively), and μ is the viscosity of the glass. Finally $\mathbf{g}$ is the gravity pointing in the positive direction of y.

The temperature of the glass (denoted by T) is governed by the heat equation

$$\rho c_p(T_t + \mathbf{v}\cdot\nabla T) = \nabla\cdot(k\nabla T) + \Phi, \tag{3}$$

where c_p and k are the specific heat capacity and thermal conductivity (respectively), and Φ is the volumetric heating (from the laser).

The relevant boundary conditions are the free-surface conditions for the Navier-Stokes equations at the workpiece surface $y = h(t,x,z)$.

$$h_t + \mathbf{v}\cdot\nabla(h-y) = 0, \tag{4}$$

$$\mathbf{n}\cdot\sigma\cdot\mathbf{n} = \gamma\kappa, \tag{5}$$

$$\tau\cdot\sigma\cdot\mathbf{n} = 0 \tag{6}$$

Here $\mathbf{n}$ and τ are the normal and tangential unit vectors of the surface, γ is the surface tension, and κ is the mean curvature of the surface. For the temperature equation, the boundary at the workpiece surface yields the radiation condition (when the effect of the air flow is small):

$$-k\mathbf{n}\cdot\nabla T = \sigma(T^4 - T_b^4), \tag{7}$$

where σ is the Stephan-Boltzmann constant and T_b is the background temperature. When the effect of radiation is small, we also use the no-flux condition.

3 Identification of the Dominant Processes

The team began by considering the subprocesses involved in the polishing process, essentially the modelling of a viscous fluid that is heated by a laser (volumetric heating) and allowed to flow. The competing mechanisms in this process are the surface tension, the viscosity, and the gravity. It was determined that surface roughness can be polished out provided that a deep region underneath the surface remains mobile and it takes longer for the heat to diffuse away. This is accomplished through the volumetric heating of the laser but can be frustrated by the decrease in the effective penetration depth of the laser as the material heats up. It was also determined that the smoothing process must be a result of the balance between surface tension and

viscosity, gravity playing an insignificant role. With a $h_p \sim 50\mu$m penetration depth and surface roughness size of volume h^2L with $h = L = 1\mu$m, the characteristic time for the heat to diffuse away is[1] $\tau_{\rm d} = h_p^2 \rho c_p/k \sim 10^{-3}$s, while the timescale for viscous flow over distance L is $\tau_{\rm v} = \mu L^4/\gamma h^3 \sim 3\times 10^{-1}$s. One can now expect reasonable smoothing with a number of passes over the workpiece.

With these time and length scales in mind, the problem was then broken into three interrelated subproblems and significant progress was made for all three.

1. The thermal problem is characterized by a rapid volumetric heating phase that reduces the viscosity of the glass and allows surface roughness to be smoothed away under the action of surface tension.
2. The flow problem that carries out the smoothing occurs in a very thin layer of viscous incompressible fluid with the viscosity suitably reduced through heating.
3. The cracking problem requires the modelling of the thermal stress generated in the heating phase and the stress generated in the cooling phase; it must compensate for a variable viscosity.

4 The thermal problem

Heating the workpiece is characterized by a short-duration heating pulse:

$$\Phi(t) = \frac{4\alpha(T)\beta}{\pi d^2} I_0(t) e^{-\alpha(T)y}, \tag{8}$$

where $I_0(t) = 100$ watts for $t \le t_1 := 10^{-4}$s. The other parameter values are the laser beam diameter $d = 1\times 10^{-4}$ m, the absorption coefficient $\beta = 0.85$, and the reciprocal penetration depth ($\alpha(T) \sim 2\times 10^4$ m^{-1}) near the surface. The following estimates based on a constant value of $\alpha(T)$ are on the conservative side, since the penetration depth is typically a decreasing function of T. Using the scalings $T \sim 4\alpha\beta I_0 t_1/\pi\rho c_p d^2$, $\mathbf{x} \sim 1/\alpha$, and $\mathbf{v} \sim \eta\alpha/\rho$, the heat equation can be nondimensionalized as

$$T_t + \frac{t_1\eta\alpha}{\rho}\mathbf{v}\cdot\nabla T = \frac{kt_1\alpha^2}{\rho c_p}\nabla^2 T + e^{-y}. \tag{9}$$

By looking at the heat balance, it is clear that diffusion dominates convection at these time and space scales, so that the convective term in the heat equation can be ignored up to first order. In other words, the heating phase is very rapid, so that during the heating phase there is no time for diffusion and the volumetric source term balances the time rate of change in the temperature. Therefore, during the heating phase, we obtain the following equation by going back to the dimensional form:

[1] The thermophysical parameters for glass are the following: $\rho = 2\times 10^3$kg m^{-3}, $c_p = 10^3$J kg^{-1}K^{-1}, $k = 5$W m^{-1}K^{-1}, $\mu = 10^5$Pa s, $\gamma = 0.3$N m^{-1}.

$$\rho c_p T_t = \frac{4\alpha(T)\beta}{\pi d^2} I_0(t) e^{-\alpha(T)y} \tag{10}$$

for $t \leq t_1$ with $T(0,\mathbf{x}) = T_0(\mathbf{x})$. Integrating the temperature equation with respect to t up to t_1, we obtain

$$T(t_1,\mathbf{x}) = T_0(\mathbf{x}) + \frac{4t_1\alpha(T)\beta}{\pi d^2 \rho c_p} I_0(t_1) e^{-\alpha(T)y}. \tag{11}$$

This sets up a depth-dependent temperature profile that is diffused away during the cooling phase (for $t > t_1$), when the heat source is switched off at $t = t_1$. The value of I_0 is chosen so that $T(t_1,\mathbf{x})$, as the initial temperature for the cooling phase, is above the melting temperature but below the vapourization temperature.

During the cooling phase, we use the following scalings: $T \sim T(t_1,\mathbf{x})$, $t \sim \rho c_p h^2/k$, $y \sim h$, and $(x,z) \sim d$. We drop the convective term and the nondimensional heat equation becomes

$$T_t = T_{yy} + \left(\frac{h}{d}\right)^2 \nabla_{xz}^2 T. \tag{12}$$

Because the surface roughness size h is quite small when compared to the diameter of the laser d, the variation in temperature is predominantly with respect to depth. These observations indicate that the temperature is primarily a function of time and distance from the surface, which is as follows (if we go back to the dimensional form):

$$\rho c_p T_t = k T_{yy}, \tag{13}$$

with $T(0,y) = T(t_1,\mathbf{x})$. Even if one takes into account the nonlinearities due to the dependence of the temperature upon the various thermophysical parameters, it can be computed rapidly and efficiently using standard techniques. Of interest is the boundary condition at the surface of the workpiece: it was not clear whether a radiative condition, a Newton cooling term, or a no-flux condition should be used. [2] We compared the two extremes (the no-flux condition and the radiative condition) and did not detect any effective difference in the thermal profile for a single pulse of heat, as shown in Fig. 2. After a number of pulses, however, the residual heating effect is significantly different under the two conditions, so that the simplification of the thermal problem achieved through the no-flux condition at the surface brings with it inaccuracies in the long-term temperature profile of the workpiece.

5 The Flow Problem

Modelling of the flow problem began with considering the Navier-Stokes equations for an incompressible viscous fluid. For a material such as glass, one of the key

[2] Far from the interface a no-flux condition was applied.

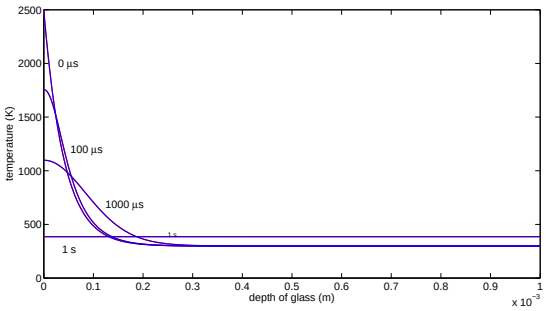

Fig. 2 Comparison between the temperature rise for a radiative condition and the rise for a no-flux boundary condition at the surface $y = 0$. The two solutions are indistinguishable.

properties is the large variation of viscosity as a function of temperature. In the current setup, the temperature varies (over the penetration depth) between the room temperature (at the bottom of the penetration depth $y = h_p := 1/\alpha$) and the melting temperature (at the workpiece surface $y = 0$). The viscosity variation over the same length scale is much greater, of the order of 10^5. Thus the typical length scale in the y direction is much smaller than that in the horizontal direction, which means that lubrication theory is applicable. In this report, we consider the case where the surface roughness h is much smaller than the penetration depth. In addition, we assume that h is independent of z. Using the scalings $y \sim h$ and $x \sim d$ and the continuity equation, we conclude that the ratio of the velocity scales U and V must satisfy $V/U \sim h/d$.

The velocity scale U (or V) is determined by balancing the surface tension and viscous forces, i.e., we have $U \sim \gamma h/\mu d$. The Navier-Stokes equations become

$$\frac{\partial p}{\partial y} = 0, \tag{14}$$

$$\frac{\partial p}{\partial x} = \frac{\partial}{\partial y}\left(\mu \frac{\partial u}{\partial y}\right), \tag{15}$$

$$\frac{\partial u}{\partial x} + \frac{\partial v}{\partial y} = 0, \tag{16}$$

with boundary conditions

$$\frac{\partial h}{\partial t} + u\frac{\partial h}{\partial x} = v, \tag{17}$$

$$p = \gamma \frac{\partial^2 h}{\partial y^2}, \tag{18}$$

at $y = 0$ and $u = v = 0$ at $y = -\infty$. Eliminating u and v, we arrive at an integral equation

$$\frac{\partial h}{\partial t} + \frac{\partial}{\partial x}\left(\int_{-\infty}^{h} \frac{\gamma(y-h)^2}{2\mu} dy \frac{\partial^3 h}{\partial x^3}\right) = 0. \tag{19}$$

Numerical solutions of the resulting expression verify that surface roughness is smoothed after multiple passes. The actual amount of smoothing depends upon the choice of surface boundary condition, as shown in Fig. 3. The no-flux condition underestimates the number of passes that would be required while the radiative condition overestimates the number of passes. If we choose carefully the boundary condition at the surface by combining the effects of both radiative and convective cooling, we should be able to reproduce the rate of smoothing that is observed experimentally.

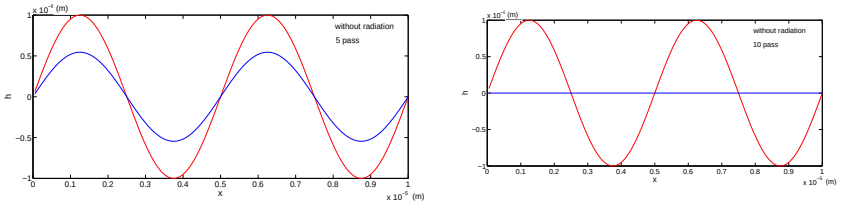

(5 passes, no flux) (10 passes, no flux)

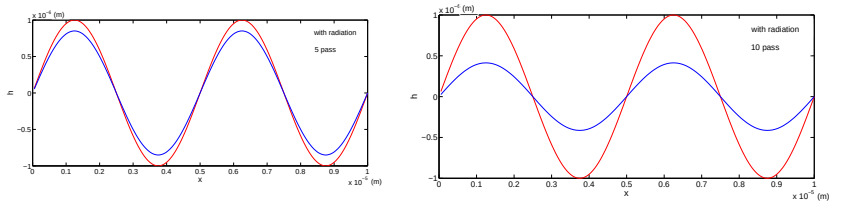

(5 passes, Radiation) (10 passes, Radiation)

Fig. 3 Evolution of the surface roughness.

6 The Cracking Problem

If the thermal regime is too severe in some fashion, then there may be cracking or "crazing" of the workpiece (which will make it useless). Examples of this effect are described in the literature, [3] which also describes a number of characteristics of the cracking. If cracking does happen, then the whole of the workpiece is affected and the cracks are typically 100μm in length and 15μm in depth. Experimentally it is noted that the cracking does not occur during heating but occurs later during cooling. There are two competing mechanisms with respect to the cracking

[3] See for example Allcock et al., J. Appl Phys. 78 (12): 7295-7303 (1995).

process: thermoelasticity (elastic expansion during heating) and viscosity (dissipation of stress by flowing). To see how these effects combine, it was noted that for small temperatures there is no stress due to expansion of the material and there is no flow since the material is cold. For material that is sufficiently heated, the stress due to the expansion of the material is large but the viscosity is reduced, which allows the material to flow (thereby dissipating the stress). What is required to analyze this problem is a thermoviscoelastic model that can resolve the "in between" regime, where the temperature is high enough to produce significant elastic stress due to the thermal expansion yet low enough to ensure that the material viscosity remains sufficiently large (thereby frustrating the flow and preventing the stress to dissipate by flowing).

The literature is very sparse in this field, with models that either ignore the effects of viscosity or include viscosity but ignore the thermal effects. In this report we consider a one-dimensional Maxwell model. We extend the Maxwell model in one-dimensional space ($0 \leq x \leq 1$) as follows:

$$\frac{\sigma}{E} + \int_0^t \frac{\sigma}{\mu} ds = \alpha(T - T_{ref}) - u_x, \tag{20}$$

$$\sigma_x = 0, \tag{21}$$

where σ is the stress, u the displacement, u_x the strain rate, and E Young's modulus. We assume the condition $u = 0$ at $x = 0$ and $x = 1$ and initially we have $T = T_{ref}$. Our model is more general than the one in Allcock et al. (1995) since T could depend on x and t and μ could depend on x and t (as a function of temperature).

From $\sigma_x = 0$ we can immediately conclude that $\sigma = \sigma(t)$ holds. We integrate the first equation in x and apply the condition $u = 0$ at $x = 0$ and $x = 1$ in order to obtain

$$\frac{\sigma}{E} + \int_0^t \frac{\sigma}{\overline{\mu}(t)} ds = \alpha \overline{T}, \tag{22}$$

where

$$\frac{1}{\overline{\mu}(t)} = \int_0^1 \frac{1}{\mu(t,x)} dx, \quad \overline{T}(t) = \int_0^1 (T(t,x) - T_{ref}) dx. \tag{23}$$

Equivalently, we have

$$\dot{\sigma} + \frac{E}{\overline{\mu}(t)} \sigma = \alpha E \dot{\overline{T}}(t). \tag{24}$$

This is an ODE that we can solve numerically in general; it is of the same form, when the temperature and viscosity in Allcock et al. (1995) are integrated, as the space average.

As an illustration, we consider a simple case where $\overline{\mu}$ is a constant. We find the solution

$$\sigma = \alpha E \int_0^t \dot{\overline{T}} e^{E/\overline{\mu}(s-t)} ds. \tag{25}$$

When $\overline{T} = (e^{-t/p} - 1)$ holds, we have $\sigma = \alpha E(e^{-t/p} - 1)$ for the pure elastic case and

$$\sigma(t) = \frac{\alpha E}{1 - \frac{pE}{\bar{\mu}}} \left(e^{-t/p} - e^{-\frac{Et}{\bar{\mu}}} \right) \tag{26}$$

for the viscoelastic case.

For a rapid cooling ($\tau \ll \bar{\mu}/E$) viscosity has little effect but for slow cooling ($\tau \gg \bar{\mu}/E$) viscosity dominates. Although the stress eventually dissipates with time, the most important factor is the maximal stress, which is expected to cause the crazing (if it takes long to cool down the workpiece). In particular, the time of maximal stress occurs at $t_{\max} = \tau\lambda \ln\lambda/(\lambda - 1)$ (where $\tau\lambda = \bar{\mu}/E$) and the corresponding maximal stress is $|\sigma_{\max}| = \alpha_0 \Delta T \lambda^{1/(1-\lambda)} \leq \alpha_0 \Delta T$ (with the upper bound corresponding to the purely thermoelastic result). Choosing a characteristic cooling time [4] of $\tau = \bar{\mu}/E \sim 10^{-5}$s generates a maximal stress at $t = \tau = 10^{-5}$s of $\alpha_0 E \Delta T e^{-1} \sim 37$MPa, roughly a third of the expected thermoelastic stress of $\alpha_0 E \Delta T \sim 100$MPa.

7 Conclusions and Recommendations

Surface roughness of the order of 1μm can be completely removed after about ten passes if there are no radiative losses and the polishing process relies on the heating of a depth of material significantly thicker than the size of the surface roughness to be polished out [5]. Modelling the residual heating and cooling rate (which is thought to be responsible for cracks in the material) may require modelling both convective and radiative cooling from the surface of the workpiece. If no heat leaves the surface, then after five passes the temperature rise is about 500K (less if there are radiative losses). Most of the heating in this process takes place within the workpiece and in comparison very little heat escapes through the surface.

A predictive model using lubrication theory and a simplified model for the volumetric heating have been effective at reproducing the smoothing out of the surface roughness through a coupling of a reduction in the viscosity and a reduction in surface tension. Heating a sufficient volume of the material is crucial to making this mechanism work.

Cracking is suspected to be a result of an excessive rate of material cooling (Allcock et al., 1995). Spatial variations in the viscosity are required to reduce the stress through viscous effects; cooling the material more rapidly than the character-

[4] Approximate viscoelastic properties of glass: $\bar{\mu} = 10^5$Pa s$, E = 10^{10}$Pa$, \alpha_0 = 10^{-5}$K$^{-1}, \Delta T = 10^3$K.

[5] This is based on a particular viscosity law obtained by fitting experimental data for fused silica glass. For other materials the numbers will likely vary but the qualitative conclusion should be similar.

istic time ($\bar{\mu}/E$) can result in significant reductions in stress and elimination of the cracking effect.

An examination of the current literature has revealed that there is a significant lack of robust models capturing the physics of this process. Significant advances in the quantitative predictability and control of the polishing process can be achieved by INO with very little capital investment.

Acknowledgement

We would like to acknowledge the financial support from the CRM and Mprime, and thank the CRM for providing a stimulating working environment during the workshop.

Évaluation de l'incertitude sur le calcul de production d'énergie éolienne

Christian Léger, Geneviève Autotte, Gilles Boesch, Michel Carreau, Stephen Y. Chen et Hervé Tchouake Tchuiguep

1 Introduction

La production d'électricité à partir de l'énergie éolienne est l'une des industries énergétiques qui connaissent la croissance la plus rapide au monde. En raison de sa géographie, le Canada se prête parfaitement à l'exploitation à grande échelle de l'énergie éolienne et pourrait en tirer de grands avantages, notamment des réductions d'émissions de gaz à effet de serre et des économies d'énergie obtenues en intégrant l'énergie éolienne aux modes de production d'énergie actuels. Selon l'Association canadienne de l'énergie éolienne (CanWEA) [2], les ressources en énergie éolienne du Canada pourraient combler 20% des besoins en matière d'électricité. Or, en septembre 2011, l'énergie éolienne fournissait 4 708 mégawatts (MW) annuels à seulement 1,2 million de foyers canadiens ; ces mégawatts étaient répartis comme le montre la Fig. 1. Même si d'énormes avancées ont été réalisées pendant les dernières décennies, de nombreuses entreprises se penchent sur le développement et le perfectionnement de la production de cette énergie renouvelable.

Christian Léger
coordonnateur, Université de Montréal

Geneviève Autotte
rédactrice et étudiante, Université de Montréal

Gilles Boesch
représentant industriel, Énergies renouvelables, Hatch

Michel Carreau
représentant industriel, Énergies renouvelables, Hatch

Stephen Y. Chen
School of Information Technology, York University

Hervé Tchouake Tchuiguep
étudiant, Université de Montréal

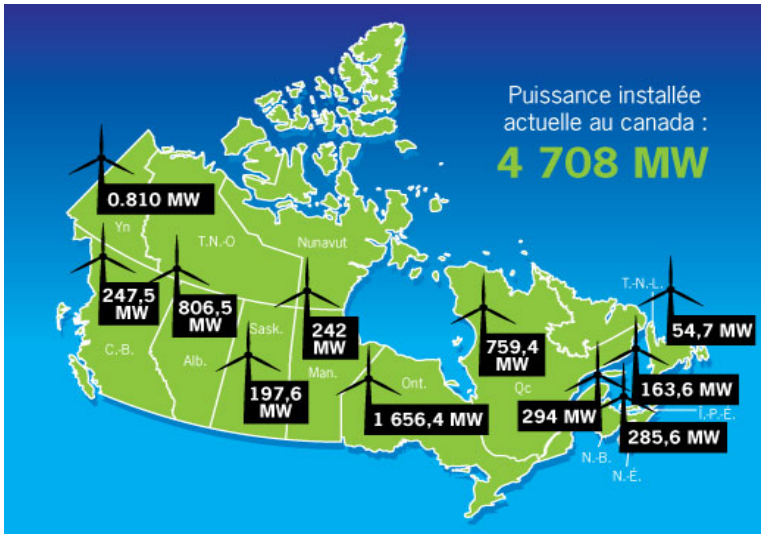

Fig. 1 Répartition de la puissance éolienne au Canada (septembre 2011) [2]

Avant de pouvoir utiliser un site pour son potentiel éolien, il faut estimer la ressource et surtout la production énergétique qui lui sera associée. La précision de cette dernière est cruciale puisque les institutions financières étudient cette estimation de la puissance énergétique minimale que le projet fournira avec une probabilité égale ou supérieure à 99% (notée P99), c'est-à-dire ce que les statisticiens appellent le quantile 1%. Or l'étude de la faisabilité du projet comporte des sources diversifiées d'incertitudes qu'il faut évidemment quantifier. De nombreuses étapes aboutissent au calcul de la puissance et de son incertitude :

1. définition du lieu de développement du futur parc éolien,
2. installation de mâts pour recueillir des données météorologiques,
3. collecte de données pendant au moins un an,
4. ajustement des données pour l'obtention d'un ensemble représentatif du long terme (généralement 10 ans),
5. création d'une carte des vents,
6. optimisation de la disposition des éoliennes en considérant l'effet de sillage créé par les autres turbines,
7. calcul de la production annuelle d'énergie, exprimée en mégawatts-heure par année (MWh/ année),
8. évaluation des pertes, et
9. évaluation de l'incertitude associée à l'énergie produite et calcul du P99.

En soumettant son problème aux participants de l'atelier, l'entreprise Hatch désire réviser et améliorer la méthodologie mathématique utilisée actuellement afin de réduire l'incertitude de la production annuelle d'énergie de l'étape 9 en se concentrant sur l'incertitude créée aux étapes 2 et 3. L'idéal serait de fournir une estima-

tion quantitative de l'incertitude du modèle numérique. De plus, l'entreprise désire connaître l'influence de dispositifs de mesure à distance tels que le SODAR ou le LIDAR.

Ainsi, dans la première partie, un exposé détaillé du problème et des concepts nécessaires à sa compréhension sont présentés. Suivent la description de quelques sources d'incertitude, les solutions proposées par l'équipe et l'influence de ces propositions sur le calcul de la puissance estimée. Par la suite, une liste d'autres sujets d'études sera donnée.

2 Description de la problématique

Dans cette section seront présentés les concepts utiles à la compréhension de la problématique ainsi que les données étudiées.

2.1 Fonctionnement d'une éolienne

De façon générale, les éoliennes commerciales de grande capacité possèdent un mât ayant de 100 à 120 mètres de haut et trois pales rattachées à un rotor pouvant avoir jusqu'à 90 mètres de diamètre. La puissance de chacune de ces éoliennes varie entre 1 et 3 mégawatts, selon le type d'éolienne. Le manufacturier de celles-ci fournit et garantit une courbe de puissance permettant la conversion de la vitesse du vent capté par l'éolienne en énergie (kW) dans des conditions standards. Pour que les éoliennes fournissent un maximum d'énergie, le vent doit être perpendiculaire aux pales. Pour cette raison, les éoliennes et leurs pales peuvent effectuer des mouvements rotationnels tenant compte de la direction du vent. De plus, les éoliennes ont besoin d'un apport minimal de 12 à 14 km/h de vent pour commencer à tourner, mais lorsque le vent dépasse 90 km/h, la production doit être interrompue pour éviter les bris d'équipement. Une distance de 300 mètres est nécessaire entre deux éoliennes, puisque chacune d'elles absorbe l'énergie du vent et le rend turbulent, ce qui influence les autres éoliennes en créant un effet de sillage derrière elle.

2.2 Instruments utilisés et données étudiées

Lors de la récolte des données météorologiques, des mâts sont installés à des emplacements bien précis sur le site du futur parc éolien. Sur ces mâts, trois anémomètres sont installés à des hauteurs différentes pour mesurer la vitesse du vent. Les mesures sont prises toutes les deux secondes, mais l'anémomètre fournit la moyenne de toutes les mesures prises pendant une période de 10 minutes. Lorsque les vents ne sont pas favorables à une prise de mesure, les données ne sont pas fournies. Les

contraintes économiques associées aux études de faisabilité et les complications techniques empêchent d'installer des anémomètres à des hauteurs du même ordre de grandeur que la hauteur des pales de l'éolienne (entre 100 mètres et 120 mètres).

En parallèle, une mesure expérimentale de la vitesse du vent a été prise par un SODAR. Encore ici, les mesures sont données à des intervalles de 10 minutes, si possible. Les SODAR (Sonic Detection And Ranging) sont utilisés depuis plus de 30 ans et sont des instruments utilisant les ondes sonores pour mesurer la vitesse des vents [4]. Le SODAR permet d'obtenir la vitesse du vent à des hauteurs semblables à celle du rotor de l'éolienne. Par contre, sa mesure n'est pas aussi précise que la mesure de l'anémomètre, et c'est pourquoi l'équipe cherchera à améliorer la prévision à l'aide des anémomètres tout en gardant à l'esprit la vitesse cible mesurée par le SODAR, en supposant celle-ci valide.

Les données fournies par Hatch provenaient de trois emplacements différents sur un même site. Le Tableau 1 résume les informations fournies pour chacun des jeux de données. L'équipe s'est concentrée sur l'analyse complète du premier emplacement, mais certains résultats ont été vérifiés pour les autres emplacements.

Tableau 1 Description des jeux de données

		Emplacement 1	Emplacement 2	Emplacement 3
Début de la prise de données		21-10-2009	15-10-2009	08-09-2010
Fin de la prise de données		05-01-2010	28-11-2009	05-06-2011
Nombre de mesures de vitesse		3879	2721	10964
Hauteur des anémomètres	1	27,9 m	29,9 m	35,0 m
	2	38,0 m	39,8 m	50,0 m
	3	48,0 m	48,8 m	80,0 m
Hauteur des mesures du SODAR	1	40,0 m	30,0 m	30,0 m
	2	50,0 m	40,0 m	40,0 m
	3	60,0 m	50,0 m	50,0 m
	4	80,0 m	60,0 m	60,0 m
	5	100,0 m	70,0 m	70,0 m
	6		80,0 m	80,0 m
	7		90,0 m	90,0 m
	8		100,0 m	100,0 m
	9		110,0 m	110,0 m
	10		120,0 m	120,0 m

2.3 Physique du vent

2.3.1 Loi de puissance

Les différents calculs utilisés dans cette étude préliminaire reposent sur la loi de puissance du vent donnée par

$$\left(\frac{V_1}{V_{ref}}\right) = \left(\frac{h_1}{h_{ref}}\right)^{\alpha}, \tag{1}$$

où V_{ref} est la vitesse de référence du vent mesurée à la hauteur h_{ref}, V_1 est la vitesse du vent à la hauteur h_1 et α est le coefficient de cisaillement.

La valeur du coefficient de cisaillement varie selon la période de la journée et de l'année et selon la géographie de l'emplacement étudié. Si nous connaissons α et la vitesse de référence à une hauteur h_{ref}, il est possible d'extrapoler la vitesse pour une hauteur donnée h_1. En effet, en prenant le logarithme de chaque côté de l'équation (1), nous obtenons les relations suivantes :

$$\log\left(\frac{V_1}{V_{ref}}\right) = \alpha \log\left(\frac{h_1}{h_{ref}}\right) \tag{2}$$

$$\log(V_1) = \log(V_{ref}) + \alpha(\log(h_1) - \log(h_{ref}))$$

$$V_1 = V_{ref} \exp(\alpha(\log(h_1) - \log(h_{ref}))) . \tag{3}$$

2.3.2 Densité Weibull

La densité Weibull est généralement utilisée comme approximation de la densité de la vitesse du vent [3]. Cette densité est donnée par la formule

$$f_{k,c}(x) = \frac{c}{k^c} x^{c-1} \exp\left[-\left(\frac{x}{k}\right)^c\right], \tag{4}$$

où x est la vitesse du vent, k le paramètre d'échelle et c le paramètre de forme.

3 Sources d'incertitudes et améliorations suggérées

L'incertitude de la production énergétique est due, en partie, à l'incertitude affectant l'estimation de la vitesse du vent à une hauteur comprise entre 100 et 120 mètres. Deux améliorations seront présentées dans cette section.

3.1 Estimation du coefficient de cisaillement α

Une manière de prédire la vitesse du vent à une hauteur désirée est d'obtenir une estimation du coefficient de cisaillement α puis d'utiliser les équations présentées dans la Section 2.3.1. Une étude détaillée de cette méthodologie sera accomplie dans cette section, où une description et une critique de la méthodologie seront effectuées. Des améliorations seront ensuite proposées.

3.1.1 Méthodologie actuelle

La méthodologie employée actuellement pour prédire la vitesse à une hauteur désirée consiste en les trois étapes ci-dessous.

1. Calculer les moyennes des vitesses obtenues avec les anémomètres pour les trois hauteurs pendant toute la cueillette des données. Ces trois moyennes seront notées $(\overline{V}_1, \overline{V}_2, \overline{V}_3)$.
2. Estimer le coefficient α en régressant le logarithme des trois moyennes $(\overline{V}_1, \overline{V}_2, \overline{V}_3)$ sur le logarithme des trois hauteurs (h_1, h_2, h_3). Rappelons que la régression est une technique statistique permettant de modéliser la relation linéaire entre une ou plusieurs variables explicatives (dans notre cas, le logarithme de la hauteur de la mesure) et une variable à expliquer (dans notre cas, le logarithme de la moyenne des vitesses du vent). L'estimation sera notée $\hat{\alpha}_{moy,ane}$. Un seul $\hat{\alpha}_{moy,ane}$ est donc obtenu et il correspond à la pente de la droite de régression.
3. Calculer les prévisions pour chaque période de temps à l'aide de l'équation (3), des vitesses de référence à la hauteur la plus élevée mesurée par l'anémomètre (V_{ref}), de la hauteur de référence (h_{ref}), de la hauteur pour laquelle la vitesse est désirée (h_1) et le coefficient $\hat{\alpha}_{moy,ane}$. Une moyenne de ces prévisions fournit une prévision globale de la vitesse du vent à une hauteur désirée.

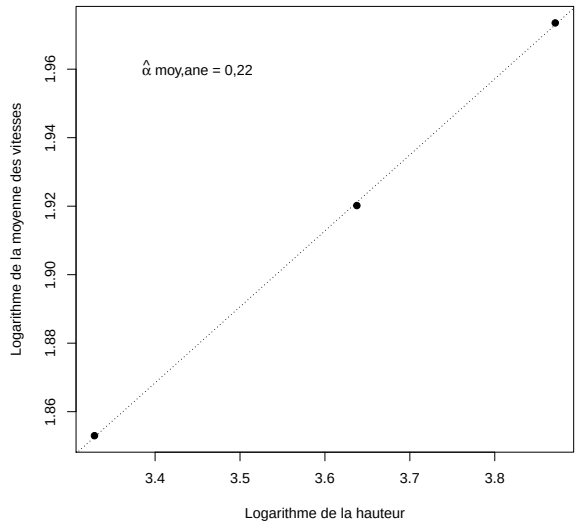

Fig. 2 Loi de puissance pour la moyenne des vitesses des anémomètres (Emplacement no 1) : les points et la droite de régression du logarithme des moyennes des vitesses obtenues par l'anémomètre en fonction du logarithme de la hauteur

La Fig. 2 présente la droite de régression obtenue pour les données de l'emplacement no 1 ainsi que les points représentant le logarithme des moyennes des vitesses obtenues par l'anémomètre en fonction du logarithme de la hauteur qui est associée à un point. La valeur estimée du coefficient $\hat{\alpha}_{moy,ane}$ pour cet emplacement est 0,22.

3.1.2 Critiques de la méthodologie actuelle

Certaines critiques de la méthodologie actuelle seront maintenant présentées. Premièrement, l'équipe a observé une variabilité considérable de l'estimation du coefficient de cisaillement (α) dans le temps. Plutôt que d'ajuster une droite de régression sur le logarithme des vitesses moyennes pendant la durée totale des mesures en fonction du logarithme de la hauteur, il est possible d'ajuster une droite de régression du logarithme de la vitesse, et ce pour chaque période de 10 minutes. On peut faire cet ajustement tant pour les données d'anémomètres (menant à une estimation $\hat{\alpha}_{10min,ane}$) que pour les données SODAR (menant à une estimation $\hat{\alpha}_{SODAR}$).

La Fig. 3 représente les variations du coefficient $\hat{\alpha}_{10min,ane}$ et du coefficient $\hat{\alpha}_{SODAR}$ pour quatre périodes successives de 10 minutes. Pour chaque période, une droite de régression du logarithme des trois vitesses du vent mesurées par les anémomètres en fonction du logarithme des trois hauteurs associées est dessinée en pointillé. On observe graphiquement que la pente de cette droite n'est pas la même pour les quatre périodes, comme en témoigne la valeur de la pente ($\hat{\alpha}_{10min,ane}$). De plus, pour chaque période, une droite de régression du logarithme des trois vitesses du vent obtenues par le SODAR en fonction du logarithme des trois hauteurs associées est dessinée en trait plein. La pente des différentes droites de régression, qui est aussi le coefficient $\hat{\alpha}_{SODAR}$, n'est pas identique pour les quatre périodes. On observe aussi une différence entre les pentes de régression pour les vitesses correspondant aux anémomètres et au SODAR. La Fig. 4 présente les variations du coefficient $\hat{\alpha}_{10min,ane}$ sur une période de 10 jours (calculé comme dans le cas de la Fig. 3). Sur ce graphique aussi on observe une variabilité non négligeable du coefficient $\hat{\alpha}_{10min,ane}$ dans le temps.

Deuxièmement, à l'aide de la Fig. 5, on peut voir que la relation entre les vitesses moyennes (m/s) et la hauteur des anémomètres est presque identique à la relation entre les vitesses moyennes du SODAR et la hauteur des mesures pour l'emplacement no 2. Cela signifie que si on considère les vitesses obtenues par le SODAR comme valeurs cibles, c'est-à-dire qu'on suppose que le SODAR fournit la *vraie* vitesse du vent, la méthode d'estimation par les moyennes fournit des vitesses équivalentes et la méthodologie semble efficace. Notons que dans ce cas, l'échelle originale est utilisée.

Par contre, lorsqu'on effectue la même analyse sur un emplacement différent (ici l'emplacement no 1), la Fig. 6 montre une certaine différence entre les vitesses moyennes du vent obtenues par les anémomètres et celles qui sont obtenues par le SODAR. Ceci jette une ombre sur la méthodologie actuellement utilisée, en admet-

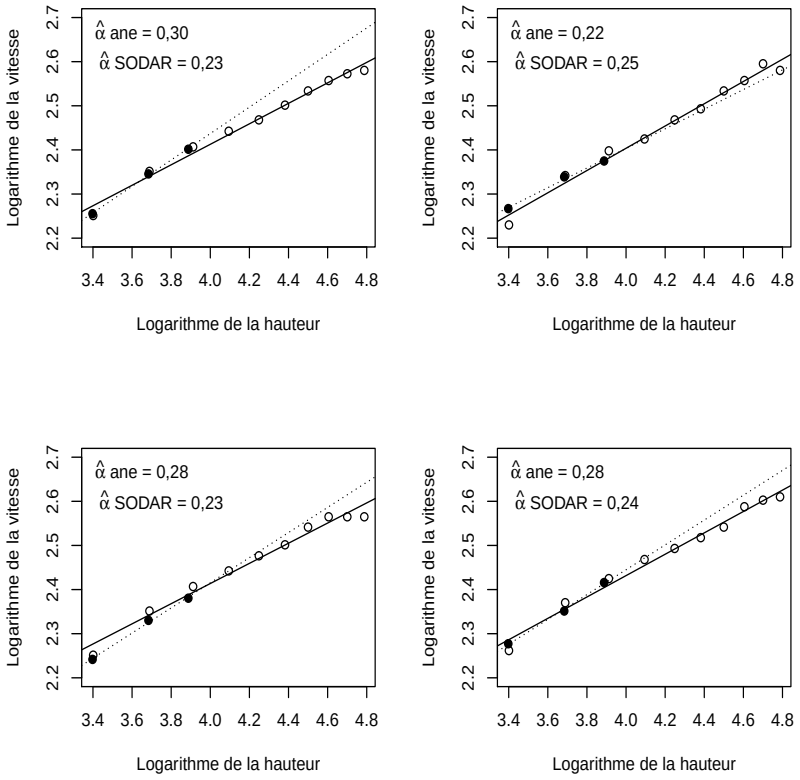

Fig. 3 Variation du coefficient $\hat{\alpha}$ pour quatre périodes consécutives de 10 minutes (Emplacement no 2) : les droites de régression du logarithme des vitesses du vent mesurées par les anémomètres (en trait pointillé) et par le SODAR (en trait plein), ainsi que les points représentant le logarithme des vitesses mesurées par les anémomètres (en trait pointillé) et par le SODAR (en trait plein)

tant que les données fournies par le SODAR sont les bonnes. Notons, encore ici, que l'échelle originale est conservée.

Troisièmement, l'équation (3) montre que l'estimation de la vitesse à la hauteur h_1 est la constante $\exp(\alpha(\log(h_1)-\log(h_{ref})))$ multipliée par la vitesse de référence (généralement la plus grande hauteur observée). Si on considère que la vitesse du vent suit une densité Weibull($k;c$) telle que fournie à l'équation (4) et qu'on multiplie cette vitesse par une constante, seul le paramètre d'échelle k devrait être modifié et non le paramètre de forme c ; en effet la densité Weibull a la propriété que si on a $X \sim$ Weibull($k;c$), alors on a aussi $bX \sim$ Weibull($bk;c$) pour $b \in \mathbb{R}$). Or, en examinant les données, l'équipe a observé que le paramètre de forme c semble augmenter selon la hauteur. En effet, pour chaque hauteur, des estimations du maximum de vraisemblance pour k et c ont été obtenues et le meilleur ajustement de la densité Weibull a été effectué (il est illustré à la Fig. 7). En ce qui concerne la densité des

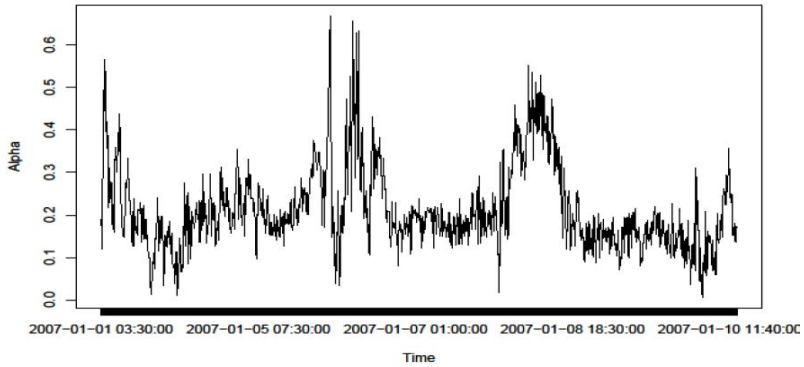

Fig. 4 Variation du coefficient $\hat{\alpha}_{10min,ane}$ calculé en effectuant une régression du logarithme des vitesses du vent mesurées par les anémomètres en fonction du logarithme des hauteurs pour une période de 10 jours (Emplacement no 3)

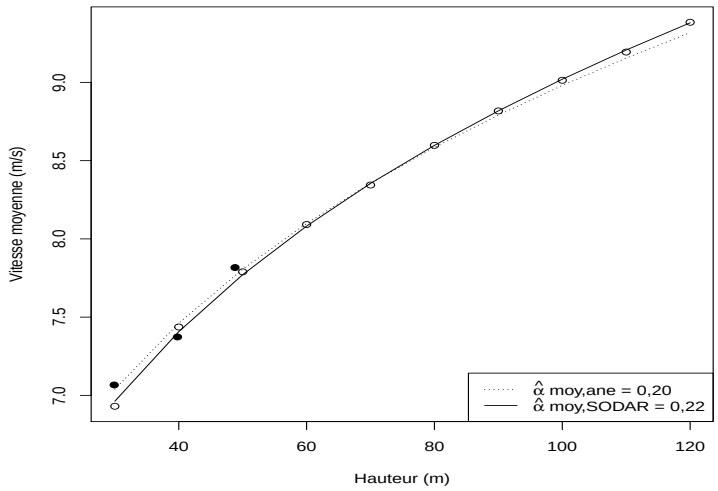

Fig. 5 Comparaison des moyennes des vitesses selon les hauteurs pour les anémomètres et pour le SODAR (Emplacement no 2)

prévisions de vitesses du vent à 100 mètres calculées par la méthode des moyennes, la valeur $c = 3{,}02$ (c'est-à-dire l'estimation à la hauteur la plus élevée) est utilisée, et on applique l'équation (3) afin d'estimer le paramètre k. Celui-ci est estimé en multipliant l'estimation de k obtenue pour la hauteur la plus élevée ($k = 8{,}05$) par

$$\exp(\hat{\alpha}_{moy,ane}(\log(h_1) - \log(h_{ref}))) = \exp(0{,}22(\log(100) - \log(48))) = 1{,}18\ .$$

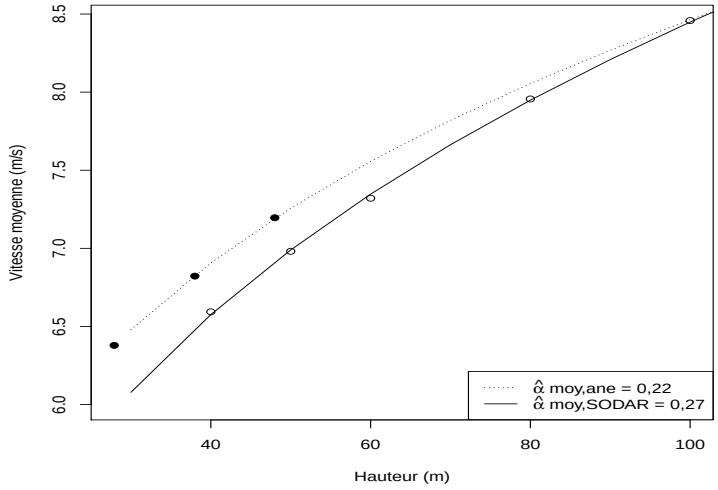

Fig. 6 Comparaison des moyennes des vitesses pour les anémomètres et pour le SODAR (Emplacement no 1)

Ainsi la densité des estimations de vitesses du vent à 100 mètres calculées par la méthode des moyennes est la densité Weibull($k = 1{,}18 \times 8{,}05; c = 3{,}02$). Dans la figure, on voit que l'estimation du paramètre c augmente avec la hauteur de la mesure. Il est donc légitime de penser que l'estimation du paramètre c à 100 mètres de hauteur n'est pas équivalente à celle faite à la hauteur la plus élevée. Ceci implique que la distribution du vent change selon la hauteur et la distribution à 100 mètres ne semble pas la même que celle à la hauteur de référence. L'utilisation de la vitesse de référence et de la hauteur de référence pour obtenir les prévisions de la vitesse du vent à une hauteur élevée ne semble pas une méthode idéale. On observe aussi sur la figure que la vitesse à 100 mètres est décalée vers la droite par rapport aux autres densités, témoignant de la modification de l'estimation de k expliquée précédemment.

3.1.3 Amélioration suggérée

Pour tenir compte des différentes critiques sur l'estimation du coefficient de cisaillement, l'équipe a proposé d'estimer la pente de la droite de régression pour la vitesse du vent mesurée par l'anémomètre ($\hat{\alpha}_{10min,ane}$) pour chaque période de 10 minutes et ainsi faire une prévision à une hauteur donnée pour une période de 10 minutes précise à partir des vitesses durant cette même période. Ainsi la méthodologie proposée se détaille comme suit.

1. Pour chaque période de temps, on effectue une régression du logarithme de (V_1, V_2, V_3) en fonction du logarithme des hauteurs associées (h_1, h_2, h_3). Les

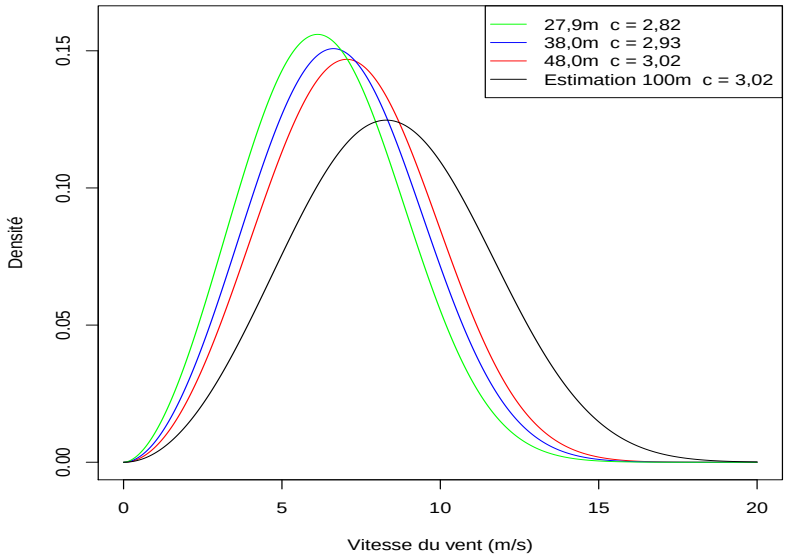

Fig. 7 Ajustements de la Weibull (Emplacement no 1) pour les vitesses des anémomètres à 27,9, 38,0 et 48,0 mètres et densité des prévisions de la vitesse à 100 mètres calculées avec un seul coefficient $\hat{\alpha}_{moy,ane}$

pentes de ces régressions donnent des valeurs de $\hat{\alpha}_{10min,ane}$ pour chaque période de temps. Nous noterons $\bar{\hat{\alpha}}_{10min,ane}$ la moyenne de ces estimations et $\hat{\beta}_{10min,ane}$ l'estimation de l'ordonnée à l'origine mesurée toutes les 10 minutes.

2. Pour chaque période, on obtient une estimation de la vitesse du vent à la hauteur désirée à l'aide de la relation linéaire

$$log(\hat{V}_{autre}) = \hat{\beta}_{10min,ane} + \hat{\alpha}_{10min,ane} \times log(h_{autre}) \,, \tag{5}$$

équivalente à

$$\hat{V}_{autre} = \exp\left(\hat{\beta}_{10min,ane} + \hat{\alpha}_{10min,ane} \times log(h_{autre})\right) . \tag{6}$$

Notons qu'aucune valeur de référence à la hauteur la plus élevée n'est nécessaire.

Avec cette méthodologie, une équation différente sera utilisée toutes les 10 minutes pour prédire la vitesse à une hauteur désirée. La Fig. 8 compare la méthode d'estimation aux 10 minutes décrite précédemment avec la méthode d'estimation par les moyennes des vitesses aux trois hauteurs décrite à la Section 3.1.1 en opposant les distributions des prévisions de la vitesse du vent à 100 mètres. Seules de légères variations sont observées dans les fréquences, mais il faut se rappeler qu'un changement dans l'estimation de la vitesse, aussi petit soit-il, peut modifier

considérablement la puissance énergétique. Pour l'emplacement no 1, nous avons $\hat{\alpha}_{10min,ane} = 0{,}24$, alors qu'on avait obtenu précédemment $\hat{\alpha}_{moy,ane} = 0{,}22$. La transformation de la vitesse du vent en puissance sera effectuée plus loin.

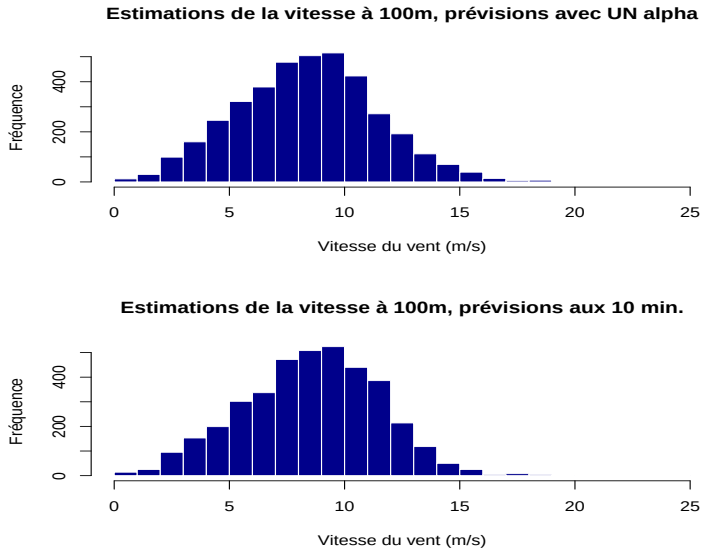

Fig. 8 Histogrammes des estimations de la vitesse du vent à 100 mètres pour la méthode actuellement utilisée, consistant à estimer un $\hat{\alpha}_{moy,ane}$ global et à utiliser les valeurs de référence afin d'obtenir les prévisions (en haut), et pour la méthode proposée, consistant à estimer $\hat{\alpha}_{10min,ane}$ toutes les 10 minutes (en bas) (Emplacement no 1)

On peut maintenant se demander si une régression basée sur seulement 3 points, telle que celle effectuée toutes les 10 minutes, représente fidèlement le modèle. Le coefficient de détermination R^2 permettant de mesurer la qualité de l'ajustement des points par la droite de régression sera calculé pour toutes les régressions à 10 minutes d'intervalle. Le Tableau 2 contient la distribution de ces R^2. On voit que plus de 75% des régressions effectuées possèdent un R^2 supérieur à 97,9%. Ainsi, comme dans ce cas le R^2 est très près de 1, la vitesse du vent semble très bien expliquée par la hauteur et ce, même si l'ajustement est effectué à partir de trois points seulement. Le Tableau 2 résume aussi la distribution de $\hat{\alpha}_{10min,ane}$, l'estimation du coefficient de cisaillement pour l'anémomètre lorsqu'il est obtenu par une régression pour chaque intervalle de 10 minutes.

La Fig. 9 est la même figure que la Fig. 7, où la densité Weibull ajustée aux vitesses prédites par la régression toutes les 10 minutes ainsi que la densité Weibull ajustée aux vitesses obtenues par le SODAR sont ajoutées. La première a été obtenue en effectuant le meilleur ajustement de la densité Weibull sur les vitesses prédites par la régression toutes les 10 minutes et la dernière sur les vitesses mesurées par le SODAR. On peut observer que la densité de la vitesse prédite par la régression faite

Tableau 2 Distribution des R^2 et $\hat{\alpha}_{10min,ane}$ pour la méthode proposée, consistant à prédire une vitesse différente toutes les 10 minutes (Emplacement no 1)

	R^2	$\hat{\alpha}_{10min,ane}$
Minimum	0,000	-0,544
1^{er} quartile	0,979	0,145
Médiane	0,995	0,206
Moyenne	0,959	0,239
3^{e} quartile	0,999	0,302
Maximum	1,000	3,108

toutes les 10 minutes est plus proche de la densité cible de la vitesse mesurée par le SODAR à 100 mètres. De plus, les densités des vitesses à 100 mètres, prédites ou observées, sont décalées vers la droite par rapport aux autres densités en raison du changement dans l'estimation du paramètre d'échelle k. La vitesse du vent est en moyenne plus élevée à 100 mètres qu'elle ne l'est à 48 mètres ou moins.

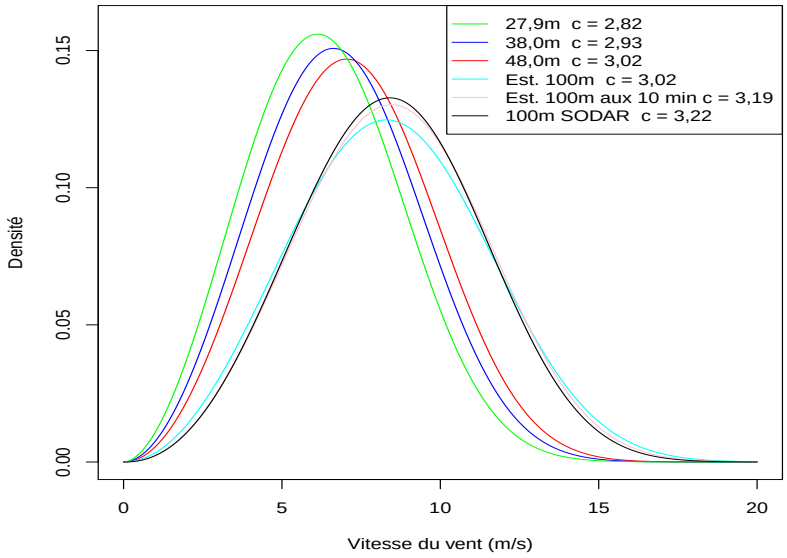

Fig. 9 Ajustements de la Weibull (Emplacement no 1) pour les vitesses des anémomètres à 27,9, 38,0 et 48,0 mètres ainsi que pour les vitesses du SODAR à 100 mètres et densités de la prévision de la vitesse à 100 mètres calculée avec un seul $\hat{\alpha}_{moy,ane}$ ou avec des $\hat{\alpha}_{10min,ane}$ toutes les 10 minutes

3.1.4 Influence de cette suggestion sur la puissance

Des modifications ont été proposées dans la section précédente pour la prévision de la vitesse du vent à une hauteur voulue. Il est maintenant intéressant d'évaluer l'effet de ces changements sur l'estimation de la puissance électrique. La Fig. 10 présente la courbe de puissance fournie par le manufacturier (permettant la conversion des vitesses du vent (m/s) en puissance (kW)) ainsi que la distribution des prévisions obtenues par des régressions toutes les 10 minutes des vitesses du vent. On observe que la puissance est très faible lorsque le vent est inférieur à 5 m/s, qu'à partir d'environ 12 m/s la puissance créée par la rotation des pales est maximale et qu'à 25 m/s la production d'énergie est arrêtée.

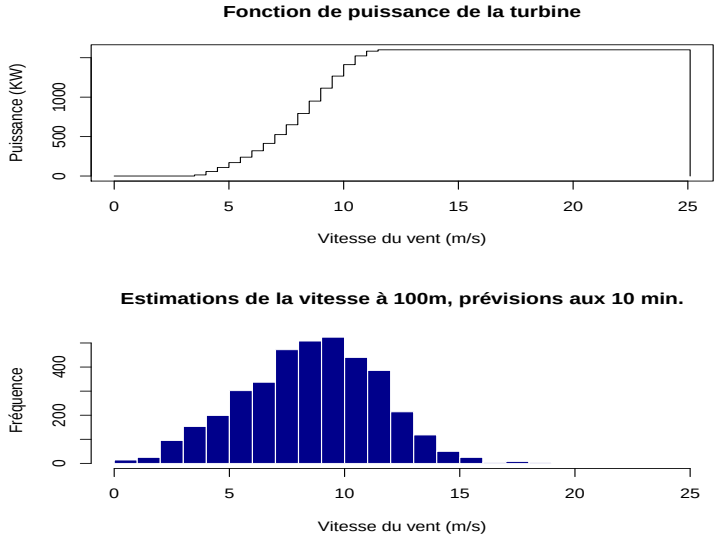

Fig. 10 Fonction de puissance (en haut) et distribution des estimations de la vitesse du vent à 100 mètres pour la méthode consistant à prédire une nouvelle vitesse toutes les 10 minutes (en bas) (Emplacement no 1)

La Fig. 11 présente une comparaison des prévisions de la puissance associée aux prévisions de la vitesse du vent à 100 mètres obtenues avec les moyennes des vitesses (méthodologie actuelle), des prévisions toutes les 10 minutes (méthodologie proposée) et des vitesses mesurées par le SODAR (puissance cible). Chaque prévision de la puissance est détaillée avec un diagramme à moustache où la borne inférieure de la boîte représente le premier quartile, la ligne blanche représente la médiane, la borne supérieure de la boîte représente le troisième quartile, et le point noir représente la moyenne.

La moyenne de la puissance estimée pour la méthodologie actuelle est de 913,73 kW alors qu'elle est de 942,98 kW pour la méthodologie proposée. Un changement de

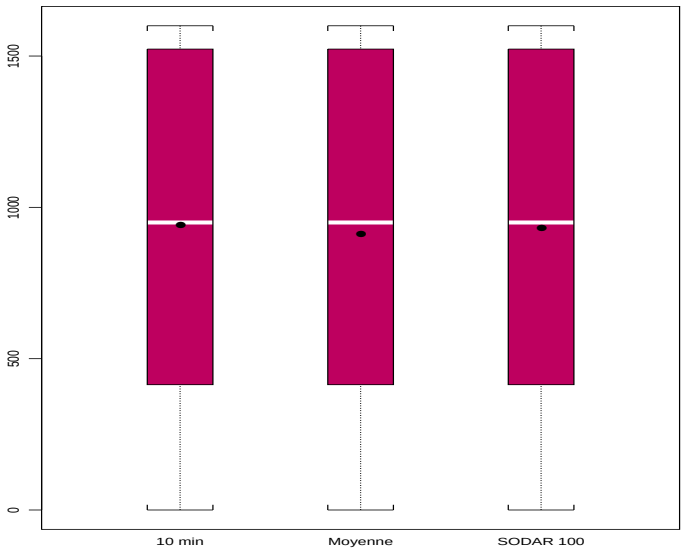

Fig. 11 Prévision de la puissance à 100 mètres : diagrammes à moustache et moyennes des prévisions de la puissance (en kW) obtenues avec les prévisions toutes les 10 minutes (méthodologie proposée), avec les prévisions par les moyennes des vitesses (méthodologie actuelle) et avec les vitesses mesurées par le SODAR à 100 mètres (puissance cible) (Emplacement no 1)

$$\frac{(942{,}98 - 913{,}73)}{913{,}73} \times 100\% = 3{,}20\%$$

pour la prévision de la puissance est noté lorsqu'on calcule avec les prévisions toutes les 10 minutes au lieu des prévisions obtenues par la méthode utilisée actuellement par l'entreprise. On remarque aussi que la prévision de la puissance pour des estimations toutes les 10 minutes est plus proche de la prévision obtenue par le SODAR qui est de 933,25 kW. Si l'hypothèse selon laquelle le SODAR fournit les vitesses exactes du vent pour des hauteurs élevées est vérifiée, on peut affirmer que la méthodologie proposée fournit une meilleure estimation de la puissance en plus d'augmenter de $3,20\%$ l'estimation de la puissance (en comparaison avec la méthodologie actuelle). Par contre, si on regarde la Fig. 6 pour l'emplacement no 1, on remarque que la droite de vitesses du SODAR et celle de l'anémomètre diffèrent alors que leurs prévisions de puissance à 100 mètres sont semblables.

Ainsi, de nouvelles méthodologies diminuant l'incertitude sur le calcul de production d'énergie éolienne mènent aussi à une modification de la prévision de la puissance.

3.2 Ajouter deux anémomètres

Pour l'instant, trois anémomètres sont installés à des hauteurs différentes afin de prendre des mesures de la vitesse du vent sur un site potentiellement intéressant pour l'industrie éolienne. Comme l'utilisation de l'anémomètre n'est pas très coûteuse, il est intéressant de poser la question suivante : l'ajout d'un ou de plusieurs anémomètres à des hauteurs différentes des autres ne pourrait-il pas réduire l'incertitude sur le calcul de production d'énergie éolienne ?

3.2.1 Méthodologie actuelle

Actuellement, trois anémomètres sont installés sur chaque mât. Des intervalles de confiance à 95% sont utilisés afin de quantifier l'incertitude associée à l'estimation du coefficient de cisaillement $\hat{\alpha}_{ane}$ ou $\hat{\alpha}_{SODAR}$. On peut calculer l'intervalle de confiance à 95% de la valeur du coefficient de cisaillement de la manière suivante :

$$\hat{\alpha} \pm t_{n-2}^{1-\alpha} \sqrt{\frac{s^2}{S_{hh}}}\,, \tag{7}$$

où

- $\hat{\alpha}$ est l'estimation du coefficient de cisaillement pour lequel on désire un intervalle de confiance,
- n est le nombre de mesures de la vitesse du vent prises à chaque période,
- $t_{n-2}^{1-\alpha}$ est le quantile bilatéral de niveau $1-\alpha$ de la loi de Student à $n-2$ degrés de liberté,
- s^2 est l'estimation de la variance de la distance de chaque point à la droite de régression,
- S_{hh} est défini comme $\sum_{i=1}^{n} (\log(h_i) - \overline{\log(h)})^2$, h_i étant la $i$ième hauteur à chaque période, pour $i = 1, \ldots, n$.

Prenons le cas simple où les hauteurs des anémomètres sont 30, 40 et 50 mètres. Alors nous avons $n = 3$, $t_{3-2}^{95\%} = t_1^{95\%} = 12{,}71$, $\overline{\log(h)} = 3{,}67$ et

$$S_{hh} = \sum_{i=1}^{n} (\log(h_i) - \overline{\log(h)})^2 = 0{,}13\,.$$

Lorsque trois mesures de vitesse sont prises ($n = 3$), la marge d'erreur est donnée par la formule

$$t_{n-2}^{1-\alpha} \sqrt{\frac{s^2}{S_{hh}}} = 12{,}71 \times \frac{s}{\sqrt{0{,}13}} = 35{,}25 \times s\,. \tag{8}$$

3.2.2 Amélioration suggérée

Si on ajoutait, par exemple, deux anémomètres supplémentaires entre les anémomètres déjà présents, c'est-à-dire si on plaçait des anémomètres à des hauteurs de 30, 35, 40, 45 et 50 mètres, alors on aurait $n = 5$, $t_{5-2}^{95\%} = t_3^{95\%} = 3{,}18$, $\overline{\log(h)} = 3{,}67$ et

$$S_{hh} = \sum_{i=1}^{n} (\log(h_i) - \overline{\log(h)})^2 = 0{,}16 \, .$$

Ainsi, lorsque cinq mesures de vitesse sont prises ($n = 5$), la marge d'erreur devient égale à

$$t_{n-2}^{1-\alpha} \sqrt{\frac{s^2}{S_{hh}}} = 3{,}18 \times \frac{s}{\sqrt{0{,}16}} = 7{,}95 \times s \, .$$

En supposant que l'estimation de la variance de la distance des points à la droite de régression s ne sera pas modifiée en moyenne par l'ajout de deux anémomètres, la marge d'erreur associée à l'estimation du coefficient α sera diminuée d'un facteur de $\frac{35{,}25}{7{,}95} = 4{,}43$. Il est très important de noter que les deux nouveaux anémomètres peuvent être installés entre les anémomètres actuels, évitant ainsi des augmentations de coût reliées à l'ajout d'un mât plus élevé. C'est donc une piste d'amélioration non négligeable de l'incertitude associée à l'estimation de la vitesse du vent.

3.3 Autres sujets de discussion éventuels

- Contrainte par le temps, l'équipe s'est concentrée sur trois emplacements d'observation de données météorologiques, mais il reste à étudier les données d'une dizaine d'autres mâts sur le même site potentiel. Pour chacun de ces mâts, des estimations de la vitesse du vent peuvent être obtenues et une analyse statistique plus poussée sera donc nécessaire.
- Pour ajuster les données à long terme (étape no 4 de l'étude de faisabilité du projet), la méthodologie actuelle consiste à obtenir une corrélation entre la vitesse du vent à l'anémomètre le plus élevé (autour de 60 mètres) et la vitesse du vent à 10 mètres provenant de la station météorologique d'Environnement Canada la plus proche (à une cinquantaine de kilomètres du site). En connaissant la corrélation entre la station et le mât étudié, la moyenne à long terme de la vitesse du vent au mât est simplement obtenue en appliquant une relation proportionnelle à la moyenne à long terme de la vitesse du vent à la station. Cette méthode mériterait d'être examinée de plus près puisqu'elle pourrait avoir une grande influence sur l'estimation de la puissance énergétique globale.
- Il est important de se rappeler que les éoliennes peuvent faire des rotations sur elles-mêmes afin de capter le vent de la manière la plus perpendiculaire possible. Dans les études exploratoires, les anémomètres n'ont pas cette capacité de rota-

tion et la vitesse du vent ne les frappe pas toujours perpendiculairement. Dans ce projet, la direction du vent n'a pas été prise en compte et il faudrait déterminer si elle a un impact considérable sur l'estimation de l'énergie éolienne (ou non).
- Tout au long de cette étude préliminaire, la vitesse obtenue par le SODAR est considérée comme la valeur cible que l'on cherche à obtenir. Si cette vitesse est biaisée, l'analyse est aussi biaisée. Comme l'a montré la Fig. 6, il faudrait étudier de façon détaillée les vitesses obtenues par le SODAR.

4 Conclusion

Une approche statistique du problème a permis de détailler quelques sources d'incertitude associées à la prévision de la puissance. Ainsi, les suggestions de l'équipe à l'entreprise Hatch sont d'estimer une droite de régression pour chaque période de 10 minutes afin de prédire la vitesse à différentes hauteurs (ce qui entraîne une modification de 3,2% de la prévision de la puissance énergétique pour une éolienne si les vitesses mesurées par le SODAR sont exactes) et d'ajouter deux anémomètres supplémentaires entre les anémomètres déjà présents (ce qui fait diminuer d'un facteur de 4,43 la marge d'erreur pour l'estimation du coefficient de cisaillement). Par ces suggestions, l'équipe a conçu un modèle statistique pouvant tenir compte de la variation quotidienne et annuelle de la vitesse du vent et du coefficient de cisaillement.

Remerciements. Nous aimerions remercier les représentants des *Énergies renouvelables de Hatch* qui ont proposé ce problème tout en nous fournissant gracieusement les données utilisées dans ce travail préliminaire.

Références

[1] W. Al-Nassar, S. Alhajraf, A. Al-Enizi, and L. Al-Awadhi. Potential wind power generation in the State of Kuwait. *Renewable Energy*, (30)14 :2149-2161, 2005.
[2] Association Canadienne de l'Énergie Éolienne. 2011. *CanWEA : L'Association Canadienne de l'Énergie Éolienne*. En ligne. <www.canwea.ca>. Consulté le 4 novembre 2011.
[3] Matthew A. Lackner, Anthony L. Rogers, and James F. Manwell. Uncertainty Analysis in MCP-Based Wind Resource Assessment and Energy Production Estimation. *Journal of Solar Energy Engineering*, 130 (3), 2008.
[4] Isidro A. Pérez, M. Luisa Sánchez, M. Ángeles García, and Beatriz de Torre. Description and distribution fitting of transformed sodar wind observations. *Journal of Atmospheric and Solar-Terrestrial Physics*, 70(1) :89-100, 2008.

Optimization of the Temporal Pulse Shape of Laser Pulses for Ablation

Emmanuel Lorin, Reynaldo Arteaga, Guillaume Blanchet, Alain Cournoyer, François Fillion-Gourdeau, Ludovic Gagnon, Claude Le Bris, Mason A. Porter, and Suzanne Shontz

Abstract In this report we describe a numerical model for the optimization of a laser-ablation process. The industrial problem was proposed by the Institut National d'Optique (based in Québec City) at the Fourth Montreal Industrial Problem Solving Workshop in August 2011.

1 Introduction

During the workshop we were planning to model, simulate, and optimize a laser-ablation process that was described to us by the Institut National d'Optique (INO).

Emmanuel Lorin
Coordinator, Carleton University

Reynaldo Arteaga
Student, Carleton University

Guillaume Blanchet
Student, Université de Montréal

Alain Cournoyer
Industrial representative, Institut national d'optique

François Fillion-Gourdeau
Student, CRM and Université de Montréal

Ludovic Gagnon
Student, Université Laval

Claude Le Bris
École Nationale des Ponts et Chaussées, Paris

Mason A. Porter
Oxford University

Suzanne Shontz
Pennsylvania State University

For details on the original problem, we refer the reader to the web site of the workshop (www.crm.umontreal.ca/probindustriels2011). In summary, we optimize the temporal shape of a laser pulse in order to obtain the maximal ablation depth of a material. In our optimization problem, the efficiency is represented by the following function:

$$J = \frac{\text{depth}}{\text{energy of the pulse}}\,.$$

The temporal pulse intensity of the laser pulse (to be optimized) is decomposed into 256 time bins, as follows:

$$I(t) = \sum_{i=1}^{256} a_i \chi_{\omega_i}\,,$$

where a_i denotes the amplitude and χ_{ω_i} the characteristic function of the bin ω_i. In addition

$$T_{bin} = |\omega_i|$$

holds, where T_{bin} is ~ 2.5 ns. The total energy of the pulse is given by

$$E_p(\mathrm{J{\cdot}m^{-2}}) = T_{bin} \sum_{i=1}^{256} a_i$$

and this relation is actually a constraint of the optimization problem. The problem is then to choose a sequence $\{a_i\}$ that optimizes the ablation depth.

We expect an optimal pulse shape to depend upon the material considered. In particular the optical absorption coefficient and the latent heat of vaporization will play a key role. In this report we focus on silicon.

2 Model

2.1 Heat Transfer Model

A Stefan model with moving boundary (obtained by a change of variables) is used to model the laser-material interaction (see (8) for details). The model is as follows:

$$\frac{\partial H}{\partial t} + u_v \frac{\partial H}{\partial x} = \frac{\partial}{\partial x} k \frac{\partial T}{\partial x} + \frac{\partial I}{\partial x}, \tag{1}$$

where H is the enthalpy ($\mathrm{J{\cdot}m^{-3}}$), u_v the vaporization velocity ($\mathrm{m{\cdot}s^{-1}}$), T the temperature (K), k the thermal conductivity of the material ($\mathrm{W{\cdot}mK^{-1}}$), and I the laser beam intensity ($\mathrm{W{\cdot}m^{-2}}$). The target is initially at the room temperature. The boundary

condition at the origin $x = 0$ is a Neumann condition, i.e.,

$$-k\frac{\partial T}{\partial x} = u_v H_v \rho \ ,$$

where ρ is the material density and H_v the latent heat of vaporization ($\mathrm{J{\cdot}Kg^{-1}}$). The numerical values of these parameters (for silicon) can be found in (8). According to the Hertz-Knudsen equation the vaporization velocity is modeled by:

$$u_v = \frac{1}{\rho}\beta p_0(T_s)\Big(\frac{M}{2\pi k_B T_s}\Big)^{1/2} \tag{2}$$

with T_s the target surface temperature, β the evaporation coefficient, M the molecular mass, and k_B the Boltzmann constant. To model the equilibrium vapor pressure at the surface we take

$$p_0(T_s) = p_\infty \exp\Big(\frac{H_v M}{T_v k_b}\big(1-\frac{T_v}{T_s}\big)\Big) \ ,$$

where p_∞ is the equilibrium vapor pressure at T_∞. The laser beam intensity is modeled by

$$\frac{\partial I}{\partial x}(x,t) = \alpha I(x,t), \qquad I(0,t) = (1-R)I_0(t) \ ,$$

where α is the absorption coefficient of the material, R the reflectivity, and I_0 the laser intensity.

From this computation we can deduce the ablation depth, which is given by the formula below.

$$\Delta z = \int_0^{t_{final}} u_v(s)ds$$

2.2 Discussion on the Physical Parameters

During the week of the workshop, we concentrated on silicon as an example, since it is the most interesting material for INO. INO has also worked with aluminum and steel, and in general they are interested in materials that can be readily purchased rather than materials that one might have to synthesize (but have physical parameters that seem to give good results based on numerical simulations). Once the methodology has been shown to work, however, there is a small finite set of parameter values for which the experiment should be repeated.

As illustrated in experiments using silicon, aluminum, and steel (6; 7)), one should expect to find different optimal pulse shapes for different materials. Hence we need to concentrate on a single material in order to develop a methodology. In

consultation with INO's representative, we chose to consider silicon. Reference (10) gives values for many of the physical parameter values that we need to know for silicon, but it does not give all of them and some of those it gives involve significant complications. This required us to do a literature search for various physical parameters. The search was complicated by the fact that different experiments have different setups and employ different variants of the materials.

INO believes that the most important physical parameters are the *latent heat of vaporization* and the *optical absorption coefficient* (denoted by α). (Indeed, the results of our numerical simulations depended heavily on α.) Based on our literature search, it seems that the parameter α has never been measured for silicon beyond the melting temperature (about 1600 K), but as an input for the PDE simulations we need the value of α up to the vaporization temperature (about 3500 K). This suggests that experiments to measure α for silicon past the melting point would be desirable, but recent papers (see (5; 7)) still extrapolate values of α based on ideas (or variants thereof) from a 1982 paper (3). This suggests that the desirable experiments are in fact very difficult (or else they would have been carried out already, given their importance).

Hence in practice, one is forced to extrapolate values of α above the melting temperature for silicon from those values that are below the melting point (7; 6). (Of course this is rather dubious, since there is no reason to expect the physical properties of the material to be the same after a phase transition as before!) It is likely that similar issues will arise for the α values beyond the melting point of other materials, but we did not check this. Given the importance of the parameter α, a potentially useful idea might be for the company to study the pulse-shape problem in materials for which the values of α are more trustworthy.

The estimation formula for α from (3) included some physical ideas, but it also included a lot of curve-fitting. The main behaviour was that α depended roughly exponentially on temperature. Hence we decided to fit $\alpha(T)$ to an exponential curve.[1] Given the many uncertainties about α (see the discussion above), we concluded that using the much more complicated (and only semi-physical) form in (3) was highly unlikely to improve results and would be difficult to implement during the workshop. Thus we used the (non-physical) exponential form

$$\alpha(T) = \alpha_0 \exp\{p_1(T-300)/k_b\}\,, \tag{3}$$

where k_b is the Boltzmann constant and α_0 equals $\alpha(300)$ (300 K was used for room temperature). The values $\alpha(300) \approx 1000$, $\alpha(600) \approx 24000$, and $\alpha(1000) \approx 200000$ are based on experimental measurements (7; 6).

We solved equation (3) directly using only the value corresponding to $T = 1000$, which yielded a coefficient of $p_1 = 1.05 \times 10^{-25}$. We examined the other choice as well to ascertain the difference in the values of p_1 for the two choices (it was reasonably small). We made our choice essentially arbitrarily. Note that the strategy

[1] Discussion with the INO representative confirmed that this was a reasonable choice.

of ignoring the semi-physical form for α in (3) in favor of curve fitting was also employed in (5), but it was carried out in a way that was different from ours. Because we needed estimates of $\alpha(T)$ for T above the melting point, we extrapolated (3) past the phase transition, even though there is no reason to believe such an extrapolation to be valid. We had to use *some* value of α as an input, but we hope that eventually a way will be found to measure values of α above the melting point of silicon.

3 Numerical Method

We discuss broadly the general numerical approach that was proposed and implemented in the `matlab` code.

3.1 General Principle

The overall algorithm to solved the optimization problem can be summarized as follows.

- Initialization of the initial intensity I_0.
- Optimization loop: step k
  1. Solving of the Partial Differential Equation from I_k to obtain the depth of ablated material.
  2. Computation of the efficiency J_k. The computed depth is used to update the objective function for the optimization.
  3. Update I_{k+1}.
- The best pulse is obtained.

The two components of the numerical method are the optimization loop and the PDE solver.

- The PDE is solver is based on a finite difference scheme. An unconditionally stable, order 2 Crank-Nicolson scheme is used, with semi-implicit source term discretization. This scheme necessitates the solving of a linear system (it is the price to pay for unconditional stability).

$$A^n T^{n+1} = B^n T^n + F^{n,n+1}(BC, source)$$

  In this system, for all n, A^n and B^n are matrices and T^n and F^n are vectors.
- Simulated annealing (for obtaining global minima) as well as multistart (for local minima) have been implemented for the optimization loop.

The different numerical parameters are defined in the code.

3.2 Some Details about the Multistart Method for Local Minima

About the Importance of Considering Local Minima. Current experiments have tried a very limited set of pulse shapes as their inputs. In most instances the pulses considered contained only single humps (7; 6), ranging from ones with most of the intensity in the back of the pulse to those with most of the intensity in the front. Seven specific shapes have been the major focus of the experiments, but of course there are numerous other pulse shapes that one could try. As discussed above, there are considerable uncertainties in both the values of the physical parameters and the PDE model itself. Changes in physical parameters, models, etc. (all of which are uncertain) may influence the relative ordering of the values of the objective function. Accordingly, one must be extremely cautious when presenting results from an optimization procedure. We observed that different choices of $\alpha(T)$ had a significant effect on the optimization results, and obviously we expect the choice of model to have a similar effect.

For these reasons it is likely that a global optimum achieved using one specific set of choices will no longer be a global optimum for another set of choices. Relying exclusively on global optima is therefore not a good idea, and finding local optima allows one to find qualitatively different pulse shapes from different parts of the objective-function landscape and hence different parts of the solution space. As discussed with the INO representative, after the initial setup, each experiment with a single initial pulse shape takes roughly one hour. (We have the impression that an experiment is also reasonably cheap in terms of direct monetary costs.) Hence it is not very costly to try (for instance) a couple of dozen qualitatively different pulse shapes from different parts of the solution space. Local minima from different parts of the objective-function landscape can give very different bin coefficients values and hence qualitatively different pulse shapes. Our goal in the modelling and computation process is to provide such candidate shapes, whose efficacy must ultimately be tested in laboratory experiments.

When looking for local minima, we specifically want to find ensembles of local minima rather than just one local minimum. This is why the Multistart idea should be a good one. This approach gives qualitatively different pulses that will be used as inputs in experiments. There are, however, additional issues to consider. For example, when calculating local optima from a complicated landscape, it is desirable to determine which ones are qualitatively the same (from the physical point of view) and which ones suggest genuinely different pulse shapes (to be used in experiments). A naive way to do this (which could be sufficient initially) would be to use some distance measure on the space of pulse shapes (as given by the bin coefficients) in order to ensure that two suggested pulse shapes are different enough from each other to be used in experiments. Then for any shape that yields promising results, one can try small deviations from it in order to see whether there is any improvement in performance. Given all the uncertainties involved in the formulation of the model, this methodology is required in any case. The results yielded by the optimization procedure should only be used as potential starting points in the solu-

tion space (these starting points should be "reasonable" from a qualitative point of view).

Another idea that we did not implement (or indeed discuss much) is that one should expect a relatively strong correlation between the optimal values in the *i*th bin and those in nearby bins. One hopes to observe some of this in the candidate pulses obtained through optimization, but one might want to refine the optimization procedure by including into the model some correlations chosen in advance.

Finding Local and Global Minima. It is interesting to consider both local and global minima of the efficiency function (under constraints), as various minima will likely correspond to very different temporal pulse shapes. Therefore each of the minima corresponds to a temporal pulse shape that an experimentalist may "try" as an input to the laser.

In order to find an ensemble of local and global minima for our industrial problem, we used the multistart method proposed by Ugray *et al.* (9), as implemented in the Matlab global optimization toolbox. The multistart method can be used to find local and global minima of nonlinear programming problems. It allows one to explore the energy landscape by starting from numerous starting points (either provided by the user or computed by the algorithm). The multistart method carries out two phases, namely a global phase and a local phase. The global phase consists of a scatter search population-based metaheuristic that searches the domain "intelligently" for constrained global minima. The local phase seeks to converge to a constrained local minimum by applying a gradient-based nonlinear programming technique such as sequential quadratic programming or an interior-point method (4).

Our Matlab script uses the MultiStart function to determine an ensemble of local and global minima for the constrained optimization problem. Once the minima have been determined, Matlab is used again to convert the real-valued coordinates of the minima into integral values, as we seek a discrete temporal pulse for the laser. Our script also returns information on the optimal values, the numbers of function evaluations and iterations performed, and convergence information.

Numerical Experiments. We tested the performance of the multistart method on a few global optimization problems of increasing difficulty. The first problem on which we tested the multistart method is known as the Six-hump Camel Back problem (1). The constrained optimization problem is the problem of minimizing

$$\left(4-2.1x_1^2+(x_1^4)/3\right)x_1^2+x_1x_2+\left(-4+4x_2^2\right)x_2^2$$

subject to the constraints $-2 \leq x_1 \leq 2$ and $-2 \leq x_2 \leq 2$. This problem has six minima: two global minima and four local minima, all of which were found by the multistart method.

The second problem on which we tested the multistart method is a problem arising in atomic physics. In particular, the following Morse potential energy function for n atoms (see (2)) was minimized:

$$f(x_1,\ldots,x_n,\rho)=\sum_{i=1}^{n-1}\sum_{j=i+1}^{n}\left(e^{\rho(1-norm(x_i-x_j))}-1\right)^2-1.$$

For our tests we chose $\rho=6$ and $n=85$. Here the variables x_i belong to $\mathbb{R}^3$. Simple bounds were added so as to constrain the problem. For this problem the execution of the multistart method was not completed because of a lack of time. Numerous local minima, however, were found by the method. It is hoped that the method will find more local minima as well as the global minima if it is allowed to run for a longer period and additional starting points are provided to the multistart algorithm.

The third problem on which we tested the multistart method is the industrial problem under consideration. Unfortunately, the multistart method was not able to determine local or global minima for the problem in its current form. In particular, we experienced numerical instabilities in the calculations of the Hessian via finite differences. This is something that could be addressed by INO either through the use of analytic gradients and Hessians or the use of automatic differentiation to compute them.

Conclusion. We employed a multistart method to find local and global minima of constrained, global optimization problems. In particular, our main goal was to determine the optimal discrete temporal laser pulse for the ablation experiments. We also considered the determination of local minima as a secondary goal, since local minima suggest new experimental inputs. The method was able to find local and global minima for problems of low to medium difficulty. The success of this method has to be demonstrated for large global optimization problems with expensive function evaluations.

There are several additional challenges with the proposed optimization methodology. First, it is important to distingush the relevant (i.e., physical or analytical) minima from the local minima that are the result of the discretization process. Second, one drawback of this technique is that nearby minima may be qualitatively similar (or symmetric) while being indistinguishable from an experimental point of view. Third, it will be important to address the numerical instability issue described earlier. Finally, one must make the multistart method as fast as possible, possibly by introducing a faster nonlinear programming technique for finding local minima; another option would be to parallelize the method based on the various starting points.

4 Preliminary results and conclusion

Some Preliminary Results. We have attempted to find an optimal (or at least near-optimal) solution for the problem submitted by INO through a simulated annealing algorithm. To use simulated annealing, we need initial solutions (i.e., pulses). Fig. 1 displays two examples of randomly chosen initial pulses; the first (resp. second) one has 20 (resp. 50) bins. An example of the successive values taken by the objective function may be found in Fig. 2.

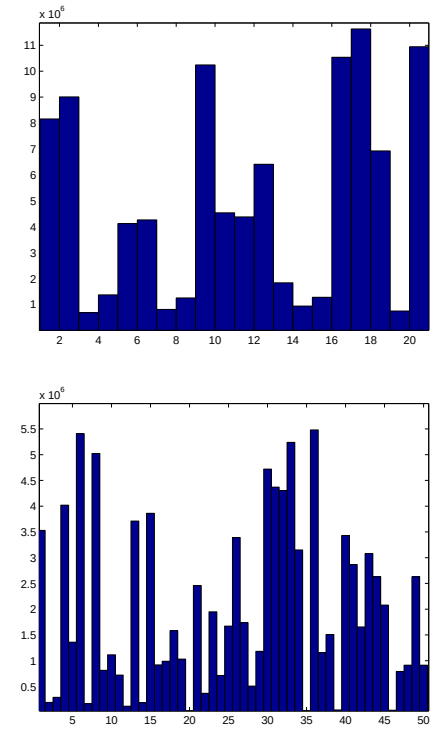

Fig. 1 Initial pulse - for 20 and 50 bins. Intensities are randomly chosen.

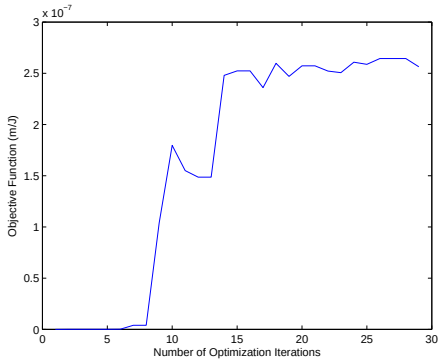

Fig. 2 Dependence of the objective function value upon the number of iterations of the optimization algorithm.

Note that at this point, the results can only show that the overall behaviour of the code is good (see Fig. 3). Realistic results will require additional tests and improvements. In particular numerical results are very sensitive to the choice of α.

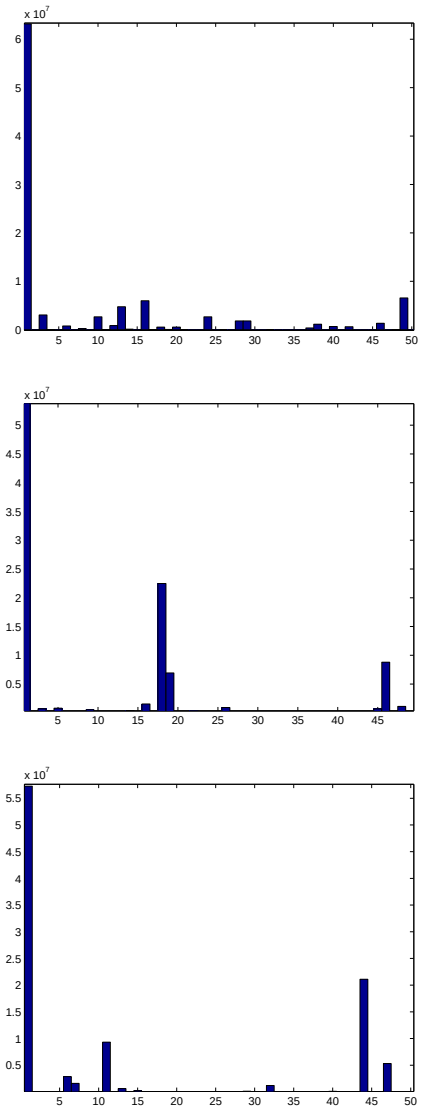

Fig. 3 Three optimized temporal pulse shapes for a similar set of data.

Conclusion. A numerical code has been written in order to optimize the laser ablation process. Preliminary results have demonstrated the good behaviour of this numerical tool. Here are the results for the company.

- A preliminary matlab code, including a PDE simulator, and two optimization methods (simulated annealing for global optima and multisart for local optima).

- The outcomes of optimization could suggest experiments with specific pulse shapes for silicon.
- A few percentage points of efficiency improvement can result in monetary benefits.

Important improvements, however, are necessary for obtaining relevant and trustworthy optimization results. Here are some suggestions.

- The numerical PDE solver could be improved.
- One could consider the optimization of multipulse laser ablation (repetition rate, scanning speed).
- The choice of the objective function could take into account additional physical parameters.
- One could carry out multidimensional tests.
- The physical model could be improved.

References

[1] L.C.W. Dixon and G.P. Szego (eds.). The optimization problem: An introduction. In *Towards Global Optimization II*. 1978.

[2] Y.G. Evtushenko, V.U. Malkova, and A.A. Stanevichyus. Parallel global optimization of a function of several variables. *Computational Mathematics and Mathematical Physics*, 49(2):246–260, 2009.

[3] Jellison Jr., G.E., and D.H. Lowndes. Optical absorption coefficient of silicon at 1.152 μ at elevated temperatures. *Applied Physics Letters*, 41(7):594–596, 1982.

[4] J. Nocedal and S. Wright. Numerical optimization. 2006.

[5] Y. Shuto, S. Yanagi, S. Asakawa, M. Kobayashi, and R. Nagase. Evaluation of high-temperature absorption coefficients of optical fibers. *IEEE Photonics Technology Letters*, 16(4):1008–1010, 2004.

[6] R Stoian. Adaptive techniques for ultrafast laser material processing. *PhD Thesis, Jean Monnet University*, 2009.

[7] R. Stoian, A. Mermillod-Blondin, A. Rosenfeld, I.V. Hertel, M. Spyridaki, E. Koudoumas, C. Fotakis, I.M. Burakov, and N.M. Bulgakova. Adaptive optimization in ultrafast laser material processing. volume 5448, pages 73–83, 2004.

[8] S. Tao, B. Wu, Y. Zhou, and Y. Gao. Thermal modeling and experimental study of infrared nanosecond laser ablation of silicon. *Journal of Applied Physics*, 106(12), 2009.

[9] Z. Ugray, L. Lasdon, J. Plummer, F. Glover, J. Kelly, and R. Martí. Scatter search and local NLP solvers: A multistart framework for global optimization. *INFORMS Journal on Computing*, 19(3):328–340, 2007.

[10] D.A. Willis and X. Xu. Heat transfer and phase change during picosecond laser ablation of nickel. *International Journal of Heat and Mass Transfer*, 45(19):3911–3918, 2002.

On Equidistant Points on a Curve

Dominique Orban, Etienne Ayotte-Sauvé, Marina Chugunova, Beatriz Cortes, Arnaud Lina, Apala Majumdar, Chris Prior, and Vincent Zalzal

Abstract We consider two problems of applied imaging science posed by Matrox on the occasion of the Fourth Montréal Industrial Problem Solving Workshop. The first problem consists in the optimal reduction of a rectangular region that intersects a given rectangular domain. The second problem consists in determining a set of N equidistant points on a curve in space. We provide theoretical results as well as efficient numerical models to solve those problems.

1 Introduction

On the occasion of the Fourth Montréal Indutrial Problem Solving Workshop, Matrox, a software and hardware design company for graphics, imaging, and computer

Dominique Orban
Coordinator, GERAD and École Polytechnique de Montréal

Etienne Ayotte-Sauvé
Natural Resources Canada, Varennes (Québec)

Marina Chugunova
Department of Mathematics, University of Toronto

Beatriz Cortes
Tecnològico de Monterrey (Mexico)

Arnaud Lina
Industrial representative, Matrox, Montréal (Québec)

Apala Majumdar
Mathematical Institute, University of Oxford

Chris Prior
Mathematical Institute, University of Oxford

Vincent Zalzal
Industrial representative, Matrox, Montréal (Québec)

vision applications, provided us with two problems originating in computer vision. The first problem consists in the optimal reduction of a rectangular region that intersects a given rectangular domain. The second problem consists in determining a set of N equidistant points on a curve in space. The first problem turned out to have been already be well studied and understood and our only contribution was to provide an optimization model that yields numerical solutions efficiently. In contrast the second problem called for existence results *and* efficient numerical solution methods. Section 2 describes the first problem along with our optimization formulation. Section 3 describes the second problem and our theoretical and numerical approaches. Numerical results are presented in §4. Conclusions and description of future work appear in §5.

2 Intersecting Rectangles

Many image processing tasks must be performed within a given region of an image, either because one wants to remove the influence of the information contained in the image as a whole, or one wants to perform the task quickly by restricting it to a domain smaller than the original one. The rectangular region is undoubtedly the most widespread geometrical shape, on account of its simplicity and ease of implementation within the computer memory. The demand for more complex shapes, however, is increasing, because of some image analyses that require a greater flexibility. A tilted rectangular region is defined by the position (i, j) of its upper left corner, its height h, its width l, and the angle a between the "wide" side of the rectangle and the x-axis of the image. The tilted rectangular region may have any dimensions or angle (with respect to the x-axis). It may be partially contained within the image or contain the entire image (among other configurations). All possible intersections between the region and the image are allowed. The region reduction consists of determining the best rectangle of angle a included in the original rectangle and the image—see Figure 1. Here, "best" is understood as maximizing the area, the height, or the width.

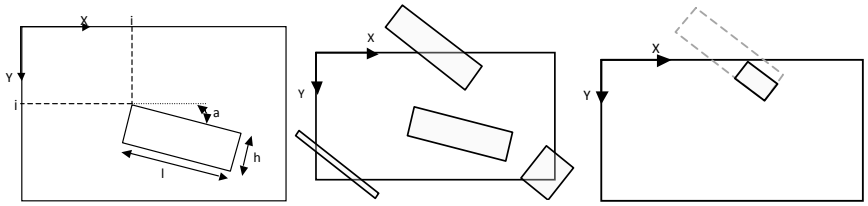

Fig. 1 The rectangle inscription problem.

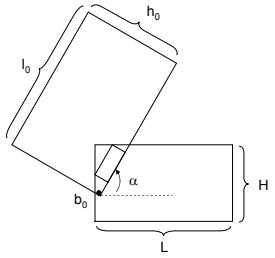

Fig. 2 Notation for problem (1).

It turns out that this is a well-known and previously studied problem. Alt et al. (1) solved the maximal area problem for the case where the intersection of the two rectangles consists in a convex polygon. They devised a nested binary search algorithm that converges in $O(\log(n))$ time, where n is the number of sides of the polygon. Daniels et al. (4) generalized previous results to nonconvex polygonal intersections.

Figure 2 illustrates our notation. Given two rectangles in the plane, find a maximal-area (resp. maximal-length, maximal-height) rectangle contained in their intersection with angle α. We formulated a mathematical optimization model for this problem and solved it numerically. Here is our formulation.

$$
\begin{aligned}
\underset{l,h,b_1,b_2}{\text{maximize}} \quad & f(l,h,b_1,b_2) \\
\text{such that} \quad & (l,h) \geq 0, \\
& 0 \leq b_1 \leq L, \\
& 0 \leq b_2 \leq H, \\
& 0 \leq b_1 + l\cos(\alpha) \leq L, \\
& 0 \leq b_2 + l\sin(\alpha) \leq H, \\
& 0 \leq b_1 - h\sin(\alpha) \leq L, \\
& 0 \leq b_2 + h\cos(\alpha) \leq H, \\
& 0 \leq b_1 - h\sin(\alpha) + l\cos(\alpha) \leq L, \\
& 0 \leq b_2 + h\cos(\alpha) + l\sin(\alpha) \leq H, \\
& 0 \leq (b_1 - b_{0,1})\cos(\alpha) + (b_2 - b_{0,2})\sin(\alpha) \leq l_0, \\
& 0 \leq (b_2 - b_{0,2})\cos(\alpha) - (b_1 - b_{0,1})\sin(\alpha) \leq h_0, \\
& 0 \leq (b_1 - b_{0,1})\cos(\alpha) + (b_2 - b_{0,2})\sin(\alpha) + l \leq l_0, \\
& 0 \leq (b_2 - b_{0,2})\cos(\alpha) - (b_1 - b_{0,1})\sin(\alpha) + h \leq h_0,
\end{aligned}
\tag{1}
$$

Note that $f(l,h,b_1,b_2)$ represents the area, height, or length of the inscribed rectangle and that (1) is a linear or quadratic program.

3 Equidistant Points on a Curve

In this problem, one studies the sampling (or selection) of points lying on a given curve but that must be spaced in a regular fashion in the plane. If one does not impose any additional constraint, it is easily shown that there exist many solutions (and indeed an infinity of them). Matrox proposed several objective functions (to be maximized or minimized) for the case of a *polyline*, i.e., a piecewise linear curve. The constraints are as follows: the solution must contain exactly $N > 2$ points on the curve, the endpoints of the curve must be part of the solution, and the Euclidean distance between two consecutive points on the curve (not the distance computed along the curve!) must be constant—see Figure 3.

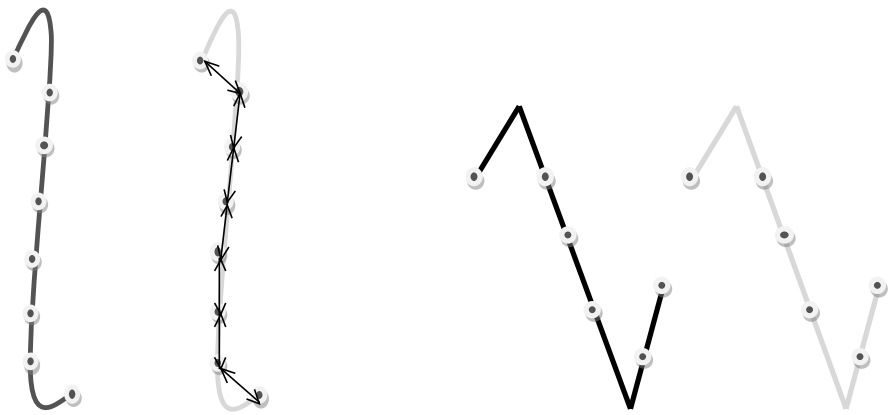

Fig. 3 Equidistant points on a smooth curve and on a polyline.

It was initially believed that solving the problem for a polyline would be easier than for a general smooth curve. In either case, the new polyline obtained by joining the equidistant points should verify one of the following criteria:

- the new polyline has minimal length,
- the new polyline has maximal length,
- the difference between the length of the new polyline and that of the original curve is minimized (in absolute value),
- the square of the difference between the length of the new polyline and that of the original curve is minimized.

3.1 Geometric Approach

In this section we consider a curve $\gamma(t) : [0,L] \to \mathbb{R}^2$ that is open-ended ($\gamma(0) \neq \gamma(L)$) and continuous. The problem consists of placing N points on the curve's

interior, at parameters $t = t_1, t_2, \ldots t_N, t_{i+1} > t_i$, $t_1 > 0$, $t_n < L$, which are ordered in t such that their arc lengths $s(t)$ are ordered on γ ($s(t_{i+1}) > s(t_i)$), and whose nearest-neighbour Euclidean distance $d(t_n, t_{n+1}) = |\gamma(t_{n+1}) - \gamma(t_n)|$ is the same for all pairs (t_n, t_{n+1}). As originally presented the problem concerned piecewise linear curves; in what follows, however, $\gamma(t)$ could also be a differentiable curve.

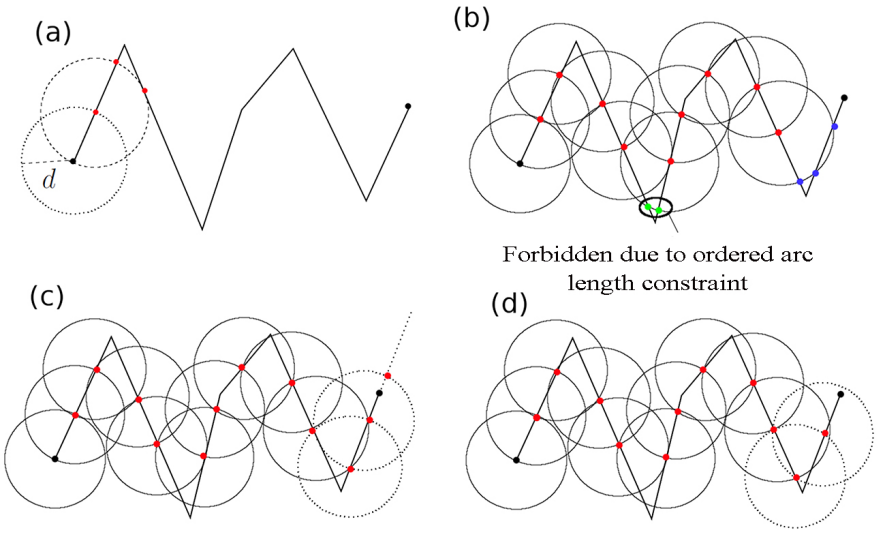

Fig. 4 Figures depicting the geometrical approach to solving the problem. This example has $N = 12$. In (a) we depict the process by which next neighbour points are chosen. In (b) we show the continuation of this process. The circled green points are points the algorithm cannot choose because they do not respect the arc length ordering. The blue points depict the possibility of multiple next neighbour points that could occur at any stage. In (c) and (d) (respectively) we depict the sets of points corresponding to two possible choices among the blue points in (b) ((c) overshoots, (d) undershoots).

To find a solution we first fix a value $d_f > 0$, which represents the radius of a circle centred at $\gamma(0)$ (see (a) in Figure 4). The subsequent intersections of the circle and γ give i possible choices for $\gamma(t_1)_i$ (see e.g. (b) in Figure 4 and (a) in Figure 5). The same process is then performed at each $\gamma(t_1)_i$ in order to obtain possible points $\gamma(t_2)_{ij}$, creating a tree structure (see Figure 5). This process is continued until either

1. the N points have been chosen (e.g (d) in Figure 4), or
2. we can no longer place points on the curve (e.g. (c) in Figure 4).

In Case 1 we must check whether $d - \varepsilon \leq d(t_N, L) \leq d + \varepsilon$ holds for some predefined accuracy ε. If this relation is satisfied we have a solution. In Case 2 we have no solution and have overshot. This check is performed for each possible solution branch.

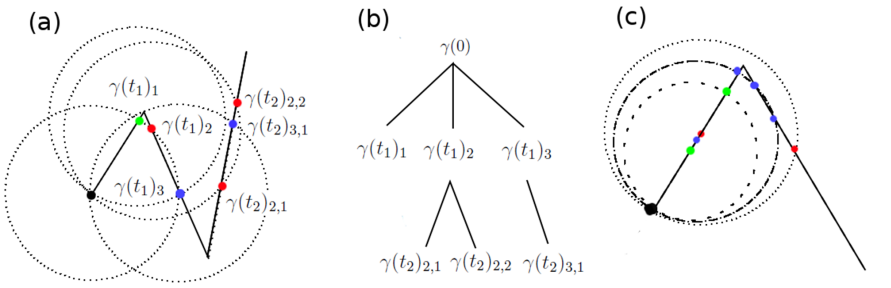

Fig. 5 In (a) and (b) we depict the tree structure of possible solution paths that must be calculated for each choice of d_f. In (c) we depict the process by which possible solution paths can appear and disappear as the radius is altered in the geometrical algorithm. As the radius decreases, the number of possible next neighbours goes from 1 to 3 then back to 1.

This process is then repeated for a different choice of d_f if no solution was found for the first choice (or if the user seeks alternative solutions). The set of possible solutions is, in general, different for each d_f, as the number of possible next neighbours can change for any given point (see (c) in Figure 5). One may continue altering d_f until a solution, or a satisfactory solution, is found. There is at least one solution to the problem (a statement proved by the group after the workshop ended).

The shortest possible path between $\gamma(0)$ and $\gamma(L)$ is $D_c = |\gamma(L) - \gamma(0)|$. We use this for defining the minimum possible distance as $d_c = D_c/(N+1)$. A consequence of our requirement that the points are ordered by arc lengths is that the arc length of γ (denoted by D_a) represents the maximum possible distance defined by a piecewise linear line connecting our points (it is a strict upper bound if γ is a differentiable curve). We can define the maximum radius as $d_a = D_a/(N+1)$. The solution (i.e., the value of d_f that satisfies our requirements) can be found amongst the possible solution trees defined for each $d \in [D_c, D_a]$, although $d[D_c, D_a]$ is an uncountable set and must in practice be sampled finitely.

The continuity of γ means that each solution can be thought of as the limit of a convergent subsequence in d_f. That is, given an exact solution d_e, there is always a solution in some ε-neighbourhood $d_e \pm \varepsilon$. If we simply sampled the domain $[D_c, D_a]$ with an equal spacing of step $\varepsilon' < \varepsilon$, where ε' is the user-defined tolerance, we would find a solution. Two major problems with this approach are the following:

- we currently do not know any way of estimating the value of ε,
- we could overcome this difficulty by making our sampling granularity very small, but this would be time-consuming.

It is difficult to get an estimate of the value of ε because many solution paths only exist over a finite range of d_f values (as depicted in Figure 5 (c)). Thus a solution that appears to be converging under a gradual change in d_f could suddenly vanish. This was probably the source of the difficulty with the first algorithm used by Matrox. The length of a particular solution path depends on all of its points and hence

the global structure of the curve, so that it is hard to generalize. With this in mind the group suggested a constructive approach to solving the problem.

The idea is quite simple. We sample the domain $[D_c, D_a]$ at $D_c + (D_a - D_c)/2$, trying all paths. If a solution is found the algorithm terminates. If not, the number of overshoots and the number of undershoots are counted. If the number of overshoots is greater (resp. less) than the number of undershoots, we choose the midpoint of $[D_c, D_c + (D_a - D_c)/2]$ (resp. $[D_c + (D_a - D_c)/2, D_a]$). If the number of overshoots equals the number of undershoots, we try the midpoints of both subintervals. This procedure is iterated until a solution is found. During the week Matrox coded a version of this algorithm and reported that it worked very well. We then tackled the question of existence: does a solution exist in every instance? We now report on the work we did during the workshop in order to answer this question.

3.2 Existence in the two-segment case

We begin with the simplest non-trivial case of a non-closed curve, that is, a polyline formed by two distinct segments P_1, P_2 of lengths l_1, l_2 (respectively); we assume that the angle between the two segments is denoted by θ and comprised in $(0, \pi)$ (see Figure 6 (a)). We consider the problem of placing N points (excluding the endpoints) on the two P_i (splitting the polyline $P_1 \cup P_2$ into $N+1$ segments). Our initial configuration is generated using arc length spacing, all consecutive points being separated by a distance $s = (l_1 + l_2)/(N+1)$. In general two points will not be separated by a Euclidean distance s, but rather a distance d satisfying the strict inequality $d < s$ (see Figure 6 (a)). Without loss of generality we can assume that there are m segments on P_1 and $N - m$ on P_2; the remaining segment is the one belonging to both P_1 and P_2. From Figure 6 (a) we have

$$d = e_1^2 + e_2^2 - 2e_1e_2\cos(\theta), \tag{2}$$

$$e_1 = l_1 - ms, \tag{3}$$

$$e_2 = l_2 - (N-m)s. \tag{4}$$

We now decrease the length of any of the m segments on P_1 and the $N - m$ segments on P_2 to $s - \delta$, where $0 < \delta < s$ (see Figure 6 (b)), in such a way that d is now a continous function of δ.

$$d(\delta) = (l_1 - ms + m\delta)^2 + (l_2 - (N-m)s + (N-m)\delta)^2 \tag{5}$$

$$- 2(l_1 - ms + m\delta)(l_2 - (N-m)s + (N-m)\delta)\cos(\theta) \tag{6}$$

Equation (5) is quadratic in δ:

$$
\begin{aligned}
d(\delta) &= a\delta^2 + b\delta + c, \\
a &= m^2 + (N-m)^2 - 2m(N-m)\cos\theta, \\
b &= 2(m(l_1 - ms) + (N-m)(l_2 - (N-m)s)) \\
&\quad - 2\left[(N-m)(l_1 - ms) + m(l_2 - (N-m)s)\right], \\
c &= d(0),
\end{aligned}
\tag{7}
$$

where $d(0)$ is given by (2)–(4).

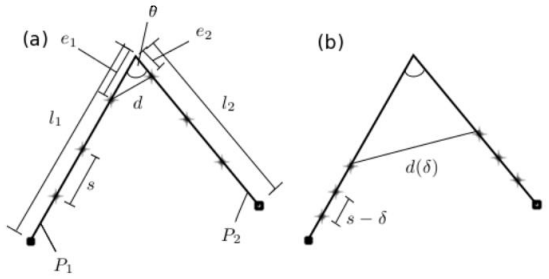

Fig. 6 Depiction of the setup used to solve the two-segment problem.

Proposition 0.1. *On the domain $\delta \in [0,s)$, there is always at least one point verifying*

$$s - \delta = d(\delta) \tag{8}$$

that solves the problem.

Proof. From the above considerations we know that $d(0) \leq s$ holds. If $d(0) = s$ holds we have found our solution. If not, as δ tends to s, there must be a positive quantity ε such that $d(\delta) > \varepsilon$ in this limit, i.e., $d(\delta) > (s-\delta)$ as $\delta \to s$. Finally, as both $s-\delta$ and $d(\delta)$ are continuous with respect to δ, (8) must be satisfied for some $\delta \in [0,s)$.

Following this work, some team members extended the result to the general (smooth) open-curve case. In addition a second constructive approach was followed. In this case explicit solutions for the polyline version of the problem with two and three segments were calculated.

3.3 A self-contained approach to the polyline problem

Consider a polyline consisting of two straight segments. The first segment has a fixed length denoted by R and makes an angle α (where $0 < \alpha < \pi/2$ holds) with

respect to the x-axis. The vertices of the first segment are located at $A = (0,0)$ and $B = R(\cos\alpha, \sin\alpha)$, respectively. There are N_1 equally spaced points on the first segment with spacing h. The second segment of length L makes an angle $\pi - \beta$ with respect to the x-axis such that $\alpha + \beta = \frac{\pi}{2}$. There are N_2 equally spaced points on the second segment with spacing h. The vertices of the second segment are located at the points $B = R(\cos\alpha, \sin\alpha)$ and $C = (R\cos\alpha + L\cos\beta, R\sin\alpha - L\sin\beta)$, respectively.

We denote the set of points on the first segment by

$$\mathscr{S}_1 = \{a_0, \ldots, a_{N_1}\} \tag{9}$$

and the set of points on the second segment by

$$\mathscr{S}_2 = \{b_0, \ldots, b_{N_2}\}. \tag{10}$$

We require that

$$a_0 = A, \quad b_{N_2} = C$$

and

$$|a_{N_1} - b_0| = h.$$

In what follows, we take $R, N_1, N_2, \alpha, \beta$ to be given quantities, subject to the constraint $\alpha + \beta = \frac{\pi}{2}$, and compute the corresponding h (the spacing) and L (the length of the second segment). This effectively answers the question: how can we construct a polyline consisting of two straight segments and containing a total of N equally spaced points with spacing h (with respect to the Euclidean norm)?

One admissible solution is the following: given N_1 and N_2 (see equations (9) and (10)), we have

$$h = \frac{R}{N_1 + \delta} \tag{11}$$

and

$$L = R\frac{N_2 + \sqrt{1 - \delta^2}}{N_1 + \delta}, \tag{12}$$

where $0 < \delta < 1$ is arbitrary. A direct computation shows that

$$|a_{N_1} - b_0|^2 = (R - N_1 h)^2 + (L - N_2 h)^2 \tag{13}$$

since $\alpha + \beta = \frac{\pi}{2}$ holds. Using the condition

$$|a_{N_1} - b_0| = h,$$

we obtain

$$h = \frac{R}{N_1 + \delta} = \frac{L}{N_2 + \sqrt{(1 - \delta^2)}},$$

where $0 < \delta < 1$ is arbitrary. Then the lengths R and L are related by

$$\frac{R}{L} = \frac{N_1 + \delta}{N_2 + \sqrt{(1-\delta^2)}},$$

as claimed in (12).

In the explicit solution (11) and (12), we can let N_1 and N_2 vary (subject to the constraint that they are positive integers and $N_1 + N_2$ is fixed) and let δ vary continuously in the range $(0,1)$ till we achieve the required geometry specified by $\{(R,\alpha),(L,\beta)\}$.

3.4 The three-segment problem

We are given a polyline consisting of three straight segments of respective lengths L_1, L_2, and L_3, oriented at the respective angles α, β, γ with respect to the x-axis. We assume that $0 < \alpha, \beta, \gamma < \pi$ holds.

The vertices of the first segment (which is of length L_1) are located at $\mathbf{A} = (0,0)$ and $\mathbf{B} = L_1(\cos\alpha, \sin\alpha)$. We have N_1 equally spaced points on this segment and we denote the distribution of points by

$$\{a_0, \ldots, a_{N_1}\}, \tag{14}$$

such that $a_0 = \mathbf{A} = (0,0)$ holds. The vertices of the second segment (which is of length L_2) are located at

$$\begin{aligned} \mathbf{B} &= L_1(\cos\alpha, \sin\alpha), \quad \text{and} \\ \mathbf{C} &= (L_1\cos\alpha + L_2\cos\beta, L_1\sin\alpha - L_2\sin\beta). \end{aligned}$$

We have N_2 equally spaced points on this segment and we denote the distribution of points by

$$\{b_0, \ldots, b_{N_2}\}. \tag{15}$$

We note that the coordinates of b_0 and b_{N_2} have to be explicitly computed. The vertices of the third segment (which is of length L_3) are located at

$$\mathbf{C} = (L_1\cos\alpha + L_2\cos\beta, L_1\sin\alpha - L_2\sin\beta)$$

and

$$\mathbf{D} = (L_1\cos\alpha + L_2\cos\beta + L_3\cos\gamma, L_1\sin\alpha - L_2\sin\beta + L_3\sin\gamma).$$

We have N_3 equally spaced points on this segment and we denote the distribution of points by

$$\{c_0, \ldots, c_{N_3}\}, \tag{16}$$

where $c_{N_3} = \mathbf{D} = (L_1\cos\alpha + L_2\cos\beta + L_3\cos\gamma, L_1\sin\alpha - L_2\sin\beta + L_3\sin\gamma)$.

We require that all consecutive points be equally spaced with respect to the Euclidean norm, with spacing h. The non-trivial part of the problem comes from the following requirement:

$$|b_0 - a_{N_1}| = |b_{N_2} - c_0|. \tag{17}$$

In coordinate form, we have

$$b_0 = (d_1, d_2),$$

and simple geometry yields

$$b_{N_2} = (d_1 + N_2 h\cos\beta, d_2 - N_2 h\sin\beta).$$

Similarly we obtain

$$a_{N_1} = N_1 h(\cos\alpha, \sin\alpha)$$

and

$$c_0 = (L_1\cos\alpha + L_2\cos\beta + (L_3 - N_3 h)\cos\gamma, L_1\sin\alpha - L_2\sin\beta + (L_3 - N_3 h)\sin\gamma).$$

The unknowns in the problem are h, d_1, and d_2. From (17) we obtain

$$\begin{aligned}&(d_1 - N_1 h\cos\alpha)^2 + (d_2 - N_1 h\sin\alpha)^2 = \\ &(d_1 + (N_2 h - L_2)\cos\beta - L_1\cos\alpha - (L_3 - N_3 h)\cos\gamma)^2 + \\ &(d_2 - (N_2 h - L_2)\sin\beta - L_1\sin\alpha - (L_3 - N_3 h)\sin\gamma)^2.\end{aligned} \tag{18}$$

Equation (18) specifies one condition for the three unknowns. We need two additional conditions. We have

$$(d_1 - N_1 h\cos\alpha)^2 + (d_2 - N_1 h\sin\alpha)^2 = h^2$$

for some fixed constant h to be worked out. Hence we have

$$(d_1 - N_1 h\cos\alpha)^2 = \delta^2 h^2$$

for some fixed $0 < \delta < 1$ and

$$(d_2 - N_1 h\sin\alpha)^2 = (1 - \delta^2)h^2.$$

Finally, for each $0 < \delta < 1$, we have three equations in the three unknowns h, d_1, and d_2.

$$\begin{aligned}&(d_1 - N_1 h\cos\alpha)^2 + (d_2 - N_1 h\sin\alpha)^2 \\ &= (d_1 + (N_2 h - L_2)\cos\beta - L_1\cos\alpha - (L_3 - N_3 h)\cos\gamma)^2 \\ &+ (d_2 - (N_2 h - L_2)\sin\beta - L_1\sin\alpha - (L_3 - N_3 h)\sin\gamma)^2 \\ &(d_1 - N_1 h\cos\alpha)^2 = \delta^2 h^2\end{aligned}$$

$$(d_2 - N_1 h \sin\alpha)^2 = (1-\delta^2)h^2$$

We solve the three above equations for each $0 < \delta < 1$. This algorithm will find a solution to the three-segment problem if it exists and will also yield explicit analytic formulas for the spacing h. If multiple solutions exist, i.e., there are multiple values of δ (in the range $(0,1)$) for which the above system can be solved, then we need to select an optimality criterion.

3.5 Mathematical Optimization Approach

3.5.1 Smooth Global Parameterization

Throughout this section we consider a curve $C \subseteq \mathbb{R}^2$ for which a continuous parameterization $p : [a,b] \to \mathbb{R}^2$ is known. Given a positive integer $N > 1$, there may exist more than one set of N regularly spaced points on C (with respect to the Euclidean metric) containing the endpoints $p(a)$ and $p(b)$. From a mathematical programming point of view, we aim to find a solution such that the Euclidean distance between consecutive points is maximal. This leads naturally to the following optimization problem.

$$\begin{aligned} \underset{t_1,\ldots,t_N}{\text{maximize}} \quad & \sum_{i=1}^{N-1} \|p(t_{i+1}) - p(t_i)\|^2 \\ \text{such that} \quad & t_1 = a,\ t_N = b \\ & t_i \le t_{i+1},\ i = 1\ldots,N-1, \\ & \|p(t_{i+1}) - p(t_i)\|^2 = \|p(t_2) - p(t_1)\|^2,\ i = 1,\ldots,N-1 \end{aligned} \tag{19}$$

We distinguish three cases:

1. the continuous parameterization p is given globally as a function that is at least twice continuously differentiable,
2. the parameterization is given in a piecewise manner over a partition of the interval $[a,b]$ via $M > 1$ local parameterizations that are at least twice continuously differentiable,
3. the parameterization is described in a piecewise linear manner.

These three situations will be respectively referred to as the global, piecewise smooth, and piecewise linear cases. For the global case, the above nonlinear program (NLP), which may be nonconvex, can be tackled with open-source (e.g. IPOPT) or commercial (e.g. CONOPT, KNITRO, SNOPT) optimization algorithms.

For the piecewise smooth or piecewise linear case, if the global parameterization p is at least twice continuously differentiable, then the values $p(t_i)$ $(i = 1,\ldots,N)$ and their derivatives may be evaluated by an external program linked to the chosen

optimization algorithm (i.e., this is the "black-box" approach). In such a context, however, optimization algorithms requiring the existence and continuity of derivatives up to the second order may not be suitable if the parameterization p is not continuously differentiable in the first or second order at breakpoints linking local parameterizations. One way to overcome this limitation is to modify (19) in order to describe explicitly all local parameterizations. This could be achieved via the introduction of binary variables (leading to a mixed integer nonlinear program formulation), but in the next section we present instead the use of complementarity constraints, leading to a mathematical program with complementarity constraints (MPCC).

3.5.2 Piecewise Parameterization

Given a positive integer $N > 1$, assume that the continuous parameterization p is described by $M > 1$ local parameterizations. Let $a = a0 < a_1 < \cdots < a_{M-1} < a_M = b$ be a partition of $[a,b]$ and let $p_j : [a_{j-1}, a_j] \to \mathbb{R}^2$ be the parameterization local to the j-th subinterval of $[a,b]$. The global parameterization is defined piecewise as

$$p(t) := p_j(t) \quad \text{for } t \in [a_{j-1}, a_j] \text{ and for all } j = 1, \ldots, M.$$

For the piecewise smooth or piecewise linear cases, the functions p_j are assumed to be twice continuously differentiable or affine (respectively).

In a mathematical program, discrete decisions that are not mutually exclusive may also be modelled without invoking binary variables through the use of complementarity constraints. Recall that such constraints take the form $u^T v = 0$ with $u, v \in \mathbb{R}^n_+$. Ragunanthan (6) introduced a reformulation method to express a real-valued piecewise smooth function in a NLP as a convex inner minimization problem, thus obtaining a bilevel NLP. In order to obtain a single-level NLP, this inner problem may be reformulated using its first-order optimality conditions (i.e., KKT conditions) as in (2). As a consequence, complementarity constraints involving Lagrange multipliers and constraint functions arise from this reformulation. In the present study, we use such a reformulation in a coordinate-wise manner, thus giving rise to the mathematical program (20). Note that in the formulation below, the variables μ_i can take negative values.

$$
\begin{aligned}
\underset{t_1,\ldots,t_N}{\text{maximize}} \quad & \sum_{i=1}^{N-1} \|x_{i+1} - x_i\|^2 \\
\text{such that} \quad & t_1 = a,\ t_N = b, \\
& t_i \leq t_{i+1}, i = 1,2,\ldots,N-1 \\
& \|x_{i+1} - x_i\|^2 = \|x_2 - x_1\|^2, \quad i = 1,\ldots,N-1 \\
& x_i = \sum_{j=1}^{M} y_{i,j} p_j(t_i), \quad i = 1,\ldots,N \\
& \sum_{j=1}^{M} y_{i,j} = 1, \quad i = 1,\ldots,N \\
& y_{i,j}\lambda_{i,j} = 0, \quad i = 1,\ldots,N,\ j = 1,\ldots,M \\
& (t_i - a_{j-1})(t_i - a_j) = \lambda_{i,j} + \mu_i, \quad i = 1,\ldots,N,\ j = 1,\ldots,M \\
& (y_{i,j}, \lambda_{i,j}) \geq 0, \quad i = 1,\ldots,N,\ j = 1,\ldots,M
\end{aligned}
\tag{20}
$$

In practical terms, with such a reformulation, the potential difficulty of solving a MINLP is transformed and not totally eliminated. To be more specific, a wide class of optimization algorithms (e.g., those considered in the present work) exploit first-order optimality conditions to construct iterates. These conditions require that the constraint gradients satisfy certain regularity assumptions (e.g., LICQ, MFCQ) in order for Lagrange multipliers to exist and be bounded, i.e., in order for the dual optimization problem to be well posed. The presence of complementarity constraints, however, ensures that LICQ and MFCQ constraint qualifications are not met. One way to surmount this difficulty (which has the advantage of enabling the use of standard optimization algorithms) is to reformulate complementarity constraints using relaxation or penalization schemes. Such approaches were implemented and successfully tested in the context of this work — see §4.

4 Numerical Experiments

4.1 Intersecting Rectangles

The mathematical program (1) was implemented in Mathematica version 8. All numerical tests were carried out with the built-in NMaximize function with default algorithmic settings (including initialization). For all problem instances considered, the computing times were less than 0.5 seconds (on an AMD Athlon 64X2 Dual Core Processor 4600+ 2.40 GHz with 2.00 GB of RAM). One could expect to improve this performance by using a compiled language. Figure 7 shows sample results.

Note that if one wishes to maximize the height h or length l of the inscribed rectangle, a positive lower bound should be imposed on l or h, respectively, so as to avoid degenerate solutions (rectangles that reduce to line segments).

4.2 Equidistant Points on a Curve

From a mathematical programming perspective, a total of five examples were considered: a circle, a spiral, a fuzzy circle, a 4-cusp hypocycloid, and a curve based on industrial case study data provided by Matrox. For all examples studied, unless otherwise stated, all optimization runs were initialized by taking equidistant parameter values in the interval $[a,b]$. The first four examples were formulated using a global smooth parameterization and implemented in both the AMPL (5) and GAMS[1] modelling languages. For these examples, the solver GAMS-IPOPT successfully converged for the following values of N (where N is the number of equidistant points): $N = 10$, 50, 100, 1000.

For the last example, two approaches were chosen on the basis of the set of data points provided: a piecewise smooth parameterization via spline interpolation of the data and a piecewise linear parameterization. The piecewise smooth parameterization was used in two ways: (i) by considering the constraints $x_i = p(t_i)$ (for $i = 1,2,\ldots,N$) as "black-boxes" during the optimization; (ii) by considering the local parameterization explicitly via the MPCC formulation. The black-box and MPCC formulations of the piecewise smooth case, implemented with Mathematica-NMaximize and GAMS-IPOPT (respectively), led to the same solution (up to 4 decimals for the coordinates of equidistant points).

The piecewise linear parameterization was formulated as a MPCC and solved with GAMS-IPOPT. As was previously observed, complementarity constraints lead to non-unique and unbounded Lagrange multipliers. For the piecewise smooth and piecewise linear parameterizations, the original MPCC formulation lead IPOPT to converge to infeasible points. Thus we considered various reformulation strategies for the complementarity constraints: $Reg(\varepsilon)$, $RegComp(\varepsilon)$, $RegEq(\varepsilon)$, and $PF(\rho)$ — see Biegler (3, Chapter 11). The $Reg(\varepsilon)$ reformulation leads to convergence for both the piecewise smooth and piecewise linear cases through solving successive optimization problems corresponding to $\varepsilon = 0.1$, 0.01, 0.001, 0.0 (we use the solution of a run to initialize the next one). In-depth testing of other reformulation approaches constitutes an avenue for future work.

[1] www.gams.com

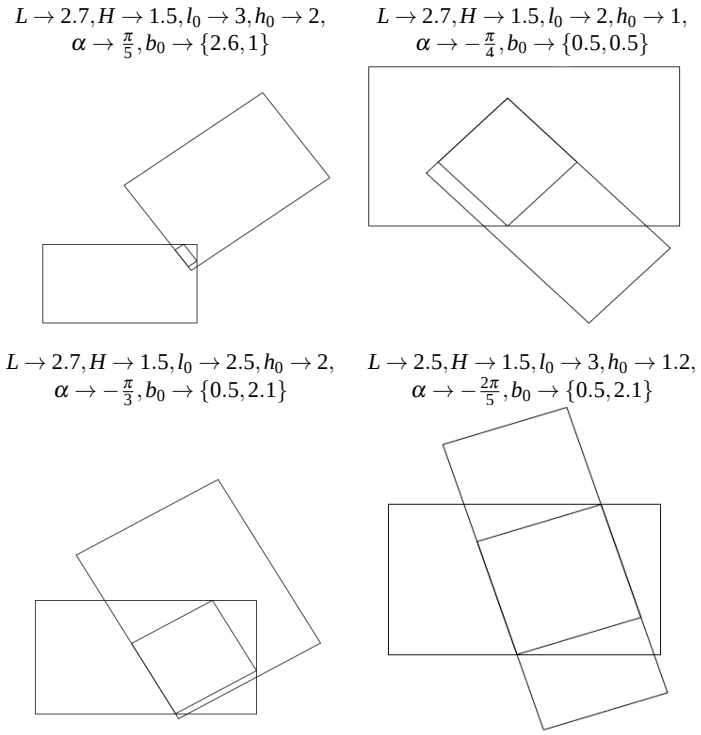

Fig. 7 Examples of a maximal area rectangle in the intersection of two rectangles.

5 Conclusions and Future Work

In this report, we summarized our achievements after four days of team work on two problems submitted by Matrox. It turned out that one of the problems had been well studied but we nevertheless presented an elegant optimization formulation that reduces to a linear or quadratic program. Highly efficient numerical methods are available for both classes of problems. The second problem, which is more challenging, calls for both existence results and efficient numerical approaches. We made substantial progress in both areas by providing partial existence results under some restrictions on the curve and preliminary numerical methods (essentially based on bisection ideas and optimization models amenable to solving high-dimensional variants of the problem).

After discussion with the industrial partner, it appears that, in the problem of equidistant points on a curve, the equidistance condition can be relaxed. If this is so, one could attempt to minimize the "defect" (i.e., the extent to which the points are not equidistant) subject to the constraint that all the points lie on the given curve.

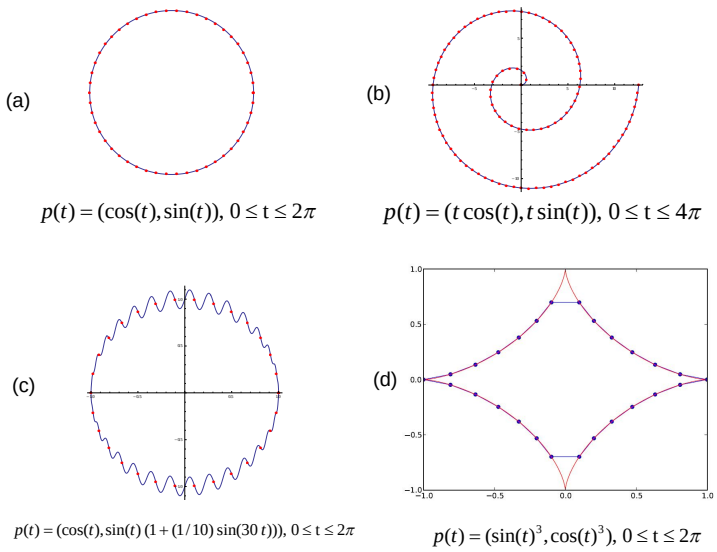

Fig. 8 Numerical results for a global parameterization of the circle (a), spiral (b), fuzzy circle (c), and 4-cusp hypocycloid (d).

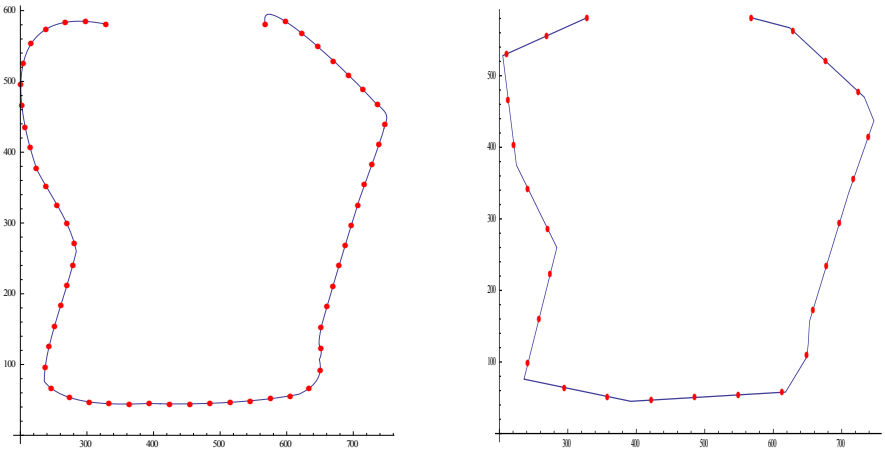

Fig. 9 Numerical results for a piecewise smooth (left) and piecewise linear (right) parameterization based on Matrox data (for the black-box and MPCC formulations).

References

[1] H. Alt, D. Hsu, and J. Snoeyink. Computing the largest inscribed isothetic rectangle. In *Proceedings of the Seventh Canadian Conference on Computational Geometry (CCCG 1995)*, pages 67–72, 1995.

[2] B. T. Baumrucker, J. G. Renfro, and L. T. Biegler. MPEC problem formulations and solution strategies with chemical engineering applications. *Computers and Chemical Engineering*, 32(12):2903–2913, 12 2008.

[3] L. T. Biegler. *Nonlinear Programming: Concepts, Algorithms and Applications to Chemical Processes*. Number 10 in MOS-SIAM Series on Optimization. SIAM, Philadelphia, PA, 2010.

[4] Karen Daniels, Victor Milenkovic, and Dan Roth. Finding the largest area axis-parallel rectangle in a polygon. *Computational Geometry*, 7(1-2):125 – 148, 1 1997. ISSN 0925-7721. doi: 10.1016/0925-7721(95)00041-0. URL http://www.sciencedirect.com/science/article/pii/0925772195000410.

[5] R. Fourer, D. M. Gay, and B. W. Kernighan. *AMPL: A Modeling Language for Mathematical Programming*. Duxbury Press/Brooks/Cole Publishing Company, 2 edition, 2002.

[6] A. U. Ragunathan. *Mathematical programs with equilibrium constraints (MPECs) in process engineering*. PhD thesis, Carnegie Mellon University, Pittsburgh, PA, 2004.

Solar Energy Portfolio Analysis

Jean-François Plante, Omar Cedrati, Simona Corina Curtean, Michel Denault, Michel Gendreau, Veronica Gheorghiade, Ozgur Gurtuna, Shahrouz Mirzaalizadeh, Mostafa Nasri, Jean-Christophe Roy, and Léonard Ryo Morin

Abstract Investments in solar energy often involve the development of solar plants in many locations. Determining the level of investment at each location can be seen as a portfolio problem whose solution requires an understanding of the interaction between all assets. The financial return from a solar energy investment is affected by multiple factors such as the fluctuation of solar radiation over time and the price of electricity. In Ontario, the FIT program guarantees a fixed price over 20 years to new eligible projects. Thus two main sources of uncertainty remain: the technical risks

Jean-François Plante
Coordinator, HEC Montréal

Omar Cedrati
Student, HEC Montréal

Simona Corina Curtean
Student, Université du Québec à Montréal

Michel Denault
HEC Montréal

Michel Gendreau
École Polytechnique de Montréal

Veronica Gheorghiade
Student, University of Guelph

Ozgur Gurtuna
Industrial representative, Turquoise Technology Solutions

Shahrouz Mirzaalizadeh
Student, École Polytechnique de Montréal

Mostafa Nasri
Student, Université de Montréal

Jean-Christophe Roy
Student, École Polytechnique de Montréal

Léonard Ryo Morin
Student, Université de Montréal

stemming from design and engineering of the plant and the variation of the solar radiation received at the site. For the latter, we use historical solar radiation data for different locations in Ontario in order to determine a representative distribution for the fluctuations in solar radiation. We use this distribution to perform a portfolio analysis where the risk of an investment is minimized, subject to constraints on the rate of return and on annual profits being sufficient to cover the debt payment obligations. The problem is solved numerically for a hypothetical investment project in Ontario.

1 Context

The planning of investments, including those in solar energy, must take into consideration the expected yield and its associated risk. For a solar plant, revenues are generated by selling the electricity produced. This production is directly related to the amount of solar radiation that the site receives.

The intensity of solar radiation at a given site depends on many geographical and climatic conditions. Atmospheric processes, such as ozone, water vapour, and clouds have a big impact on solar radiation and consequently, they are one of the main sources of uncertainty in the output of a solar installation. Under clear sky conditions, a large fraction of the solar radiation reaches the ground and can be transformed into usable power. On a cloudy day, however, a significant portion of the radiation is absorbed or reflected by the clouds and a smaller amount of energy reaches the ground. Rapidly changing weather conditions, such as the formation of clouds over a project site, can drastically alter the output of a solar power plant from one minute to the next.

Solar investment opportunities often offer the possibility to develop or buy plants from different sites. To assess the financial risk of such investments, it is important to model the variations in the production. Turquoise Technology Solutions Inc., a Montréal-based company specialized in the development of decision support tools for solar energy development, has the technology to infer the radiation of many locations in Canada from data based on satellite imagery. For the purposes of this workshop, however, we have used radiation measurements that were determined by Environment Canada for 14 sites over a period of 50 years.

The behaviour of solar radiation at various time scales has been studied by a number of authors. While Zaharim et al. (2009) propose a time series model for the daily radiation in the Tropics, Boland (2008) proposes time series models for hourly and daily radiation using Australian climatic data. Davies & Binyamin (2008) focus on Canadian data with an emphasis on UV-B irradiance. Many others propose different approaches to capture the stochastic behaviour of the radiation in one location, but we did not find research papers where a joint model is developed for multiple locations. Depending on geographical distance and relief, some locations may experience similar radiation levels while others may be uncorrelated. This property has

to be taken into consideration when assessing the variability of the total production from multiple sites.

Large investment projects often consider developing more than one site for solar energy production. The business plan of such large projects requires evaluating the variations in production, so that the aggregate risk of the development may be assessed. This risk depends on the geographical correlation of irradiation.

The proposed problem is thus to optimize a solar portfolio: given a level of investment and data on some sites that could be developed, what is the best allocation of funds to develop those sites? First introduced by Markowitz (1952), portfolio theory addresses the problem of allocating funds amongst different investment opportunities, with the goal of minimizing the risk given a lower bound on return. Historically, the variance has been used as a measure of risk, but we will consider the value-at-risk as an alternative.

Portfolio theory is used in energy policy planning to optimize the cost of electricity by choosing the proportion of each source of energy in the generation mix (see e.g. Bazilian & Roques 2008). Drake & Hubacek (2007), however, used portfolio theory to evaluate the variability of the amount of energy produced by wind farms at different locations. Our problem is very similar to theirs, except that we focus on solar (instead of wind) power.

While Section 2 presents the assumptions that we can make for this problem in the context of an energy investment in Ontario, Section 3 describes the data that we use, and more importantly, their stochastic properties. The optimization problem to be solved is presented in Section 4 and a solved example is described in Section 5.

2 Simplifying Assumptions

We assume that the investment project takes place in Ontario, Canada, in the near future. Currently Ontario has a government subsidy program to foster the development of solar energy. The FIT (Feed-in Tariff) program guarantees a fixed rate for all the electricity produced by eligible new projects for a period of 20 years.

In many markets, including Ontario, the electricity produced is usually sold on the spot market, where its price is determined daily, or at a high frequency (every five minutes or more often) on the basis of supply and demand. Estimating the value of the produced electricity in this context would require a model that can predict how much power is generated from the plant every five minutes, and a pricing model that estimates how much the electricity produced is worth at that time.

The output of a solar power plant can change drastically as the prevailing weather conditions change. If the price of electricity is fixed, an important amount of uncertainty is removed, and one does not need to know when the electricity is produced, but only how much is produced.

It is a well-known fact in Canada that the length of days and the amount of solar radiation reaching the ground vary throughout the year. These variations affect the production of solar plants as well as the demand for electricity. Since the price is fixed, however, these seasonal variations can be ignored, and it suffices to consider the yearly amount of radiation measured at the sites of interest.

The design of a solar plant will also impact its capacity to transform radiation into electricity. For instance the inclination of the panels will change the amount of electricity generated. To simplify the problem, we assume that each plant has a fixed capacity factor that we use to convert the radiation into an estimation of the amount produced. If a better technology is chosen for any given site, then its capacity factor will be higher. We further assume that a site can be scaled: a site can be replaced by a site that is twice as big and produces twice the amount produced by the original site (assuming that the efficiency does not change). For every site we define a variable called "size" that represents the area of the project.

The investment decision must take into consideration costs that are relatively independent of the site size: legal and environmental evaluations, infrastructures such as roads and buildings, connection to the grid, etc. Although we are aware that the scaling function may not be simple in practice, we assume that the fixed cost of developing a site is negligible when compared to the cost per unit of size; in other words, the cost of developing a site is directly proportional to its size.

Additional assumptions are made on the financial aspects of the project. An investment horizon of 20 years is assumed and the annual profits are invested at a fixed risk-free interest rate of j. It is assumed that 70% of the investment is financed with a loan that bears an annual interest rate of u for a term of n years with annual payments. It is further assumed that the plant has no depreciation and that its value at the end of the 20-year period will be the same as the construction cost (unadjusted for inflation).

3 Data

Turquoise Technology Solutions Inc. compiled data from Environment Canada in order to provide the workshop participants with yearly solar radiation measured or modelled at 23 sites in Ontario. Some locations had data from 1953 to 2004, but at some locations measurements did not start until later. We decided to focus on the 14 sites that had complete data for the 1955 to 2004 period. Figure 1 shows the annual average irradiation in watts per square meter for these 14 locations over the 50-year period.

To assess investment risk, we need to determine the nature of the stochastic variations in the solar radiation. We thus look at certain characteristics of the data in greater detail. Considering yearly data removes the seasonal effect, but there could be longer-term cycles. Weather is known to have multi-year cycles (e.g. el niño

and el niña) and the sun is known to have cycles as well (e.g. sun spots). For every location, therefore, we look at the autocorrelation functions depicted in Figure 2. On these figures, every vertical line shows the correlation between the yearly solar radiation at that location and a lagged version of itself. The horizontal bars show the critical values for a test that those correlations are null. If autocorrelation is present, some of the vertical bars will exceed the horizontal bars, but as we can see, the autocorrelation is typically low for every location and every lag. Generally speaking, it appears reasonable to assume that the data is independent from year to year.

The next task is to find a distribution for the yearly solar radiation. Figure 3 shows the histograms of the data for every location. With a sample of 50 years, we cannot expect very smooth histograms. Nonetheless they look rather symmetric, which is in agreement with the bell shape of the normal distribution. If the normal distribution were a good fit, its mathematical properties would simplify many calculations. To investigate further whether this assumption holds (or not), Figure 4 shows the QQ-plots of the data at every location. A QQ-plot is a diagnostic plot that indicates whether data comes from a hypothesized distribution or not. If the data comes from the specified model (the normal distribution, in our case), the QQ-plot is a straight line. Otherwise it displays curvature. The fit is not perfect (see e.g. the general shape in Ottawa_NRC and the outliers in Timmins), but globally, the assumption of a normal distribution seems excellent for some locations and acceptable for the rest.

Finally we need to determine whether there are correlations between the different sites. Table 1 displays the correlation matrix calculated from the data. Given

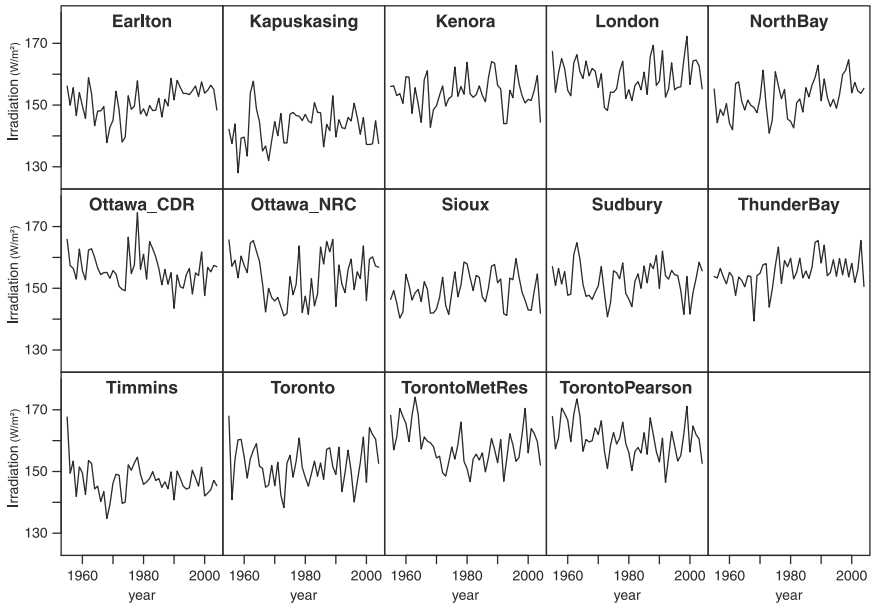

Fig. 1 Yearly radiation for 14 sites in Ontario from 1955 to 2004.

that some of the locations are within the same city, we expect to observe high correlations between some of the sites, which is indeed the case. Outside the multivariate normal paradigm, looking at correlations is not sufficient to determine the dependence structure within the data. A copula-based approach (see e.g. Genest & Favre 2007 for an introduction) is more flexible and can capture better the occurrence of simultaneous extreme events. Using copulas for our problem would lead to significant technical complications. Since there is no clear indication that a multivariate normal model is inappropriate, we adopt this simple model for modelling the data. Therefore our descriptive analysis of the data leads us to assume that there is no temporal correlation from one year to the next and that the annual radiation follows a multivariate normal distribution.

4 Portfolio optimization problem

In portfolio theory, a certain amount of money is allocated to different products in order to maximize yield or minimize risk. For our solar portfolio problem, we assume that \$ C are to be invested across the 14 potential sites in Ontario that we discussed previously. We let s_i denote the size (in square meters) of the power plant to be developed at site i and v_i the development cost per square meter at site i. By grouping costs and sizes in vectors, we may write our development budget constraint as $\mathbf{v}^{\mathsf{T}}\mathbf{s} \leq C$.

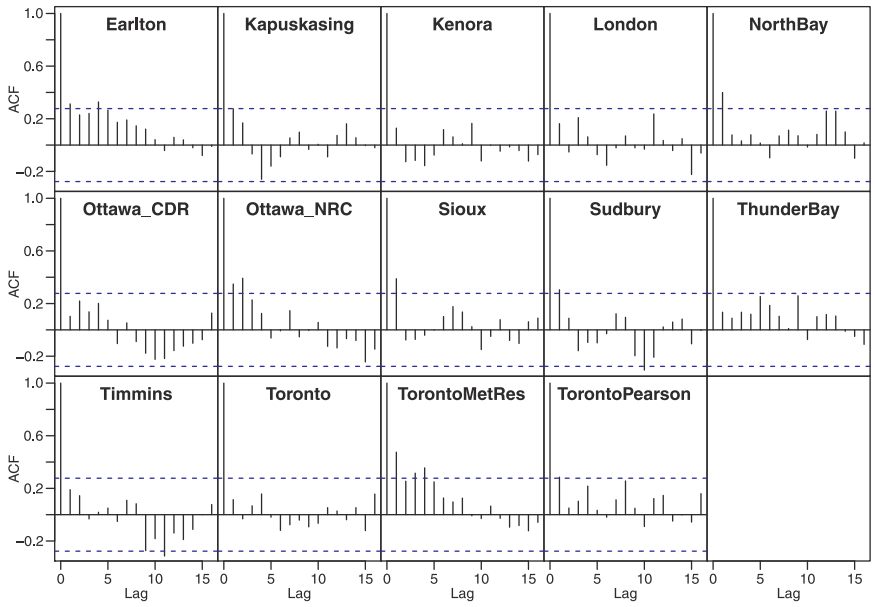

Fig. 2 Autocorrelation function of the yearly radiation for 14 sites in Ontario from 1955 to 2004.

Now given a site i and a year t, let I_{it} denote the average hourly amount of radiation (in watts per square meter) received at site i for year t. The data in Section 3 are observations from these random variables. Based on our previous analysis, we assume that the vectors $\mathbf{I}_t = [I_{1t}, \ldots, I_{14t}]^\mathsf{T}$ are independent realizations of a multivariate normal distribution with mean μ_I and covariance Σ_I.

We assume that the performance of the plant can be summarized by a single capacity factor that takes into consideration the efficiency of the panels (e_1) and that of the entire plant (e_2). Multiplying I_{it} by the capacities, the number of hours in a year, and dividing by 10^6 to convert the watts in MW, we can then write a simple conversion factor $\kappa = e_1 e_2 (365.25 \cdot 24)/10^6$ and define $X_{it} = \kappa I_{it}$, the amount of electricity (in MWh) produced per square meter at site i in year t. Accommodating different efficiencies for the considered sites is straightforward, but it makes the notation more cumbersome. By the properties of the multivariate normal distribution, the $\mathbf{X}_t$ are independent vectors from a multivariate distribution with mean $\mu_X = \kappa \mu_I$ and covariance $\Sigma_X = \kappa^2 \Sigma_I$. As a consequence, the total quantity of energy produced in year t, given by the expression $\mathbf{s}^\mathsf{T}\mathbf{X}_t$, follows a normal distribution with mean $\mathbf{s}^\mathsf{T}\mu_X$ and covariance $\mathbf{s}^\mathsf{T}\Sigma_X\mathbf{s}$.

For a typical solar energy investment project under the FIT regime, 70% of the money invested is financed through a loan. Let us assume an annual interest rate of u is paid on the loan and reimbursed in n equal annual payments. Then in year t, the amount

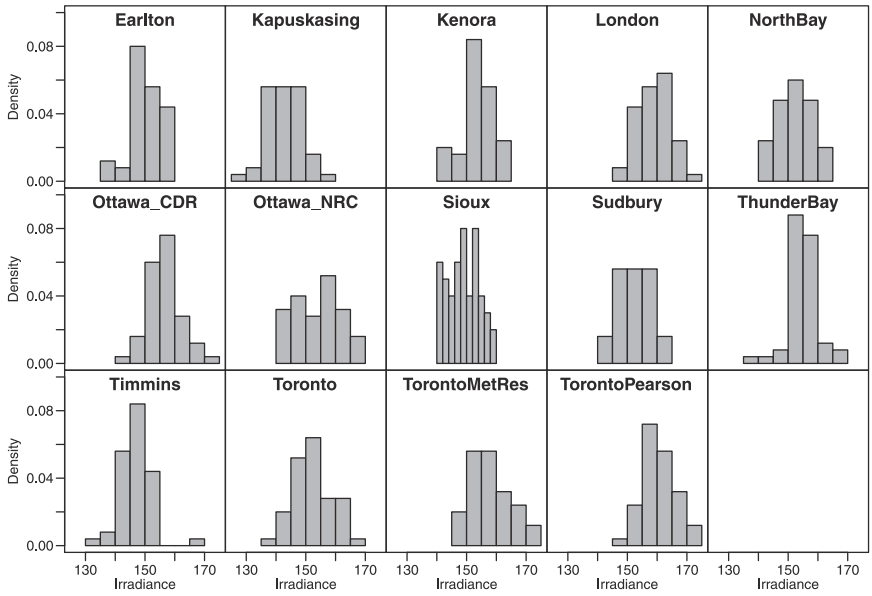

Fig. 3 Histograms of the yearly radiation for 14 sites in Ontario from 1955 to 2004.

$$M_t = \begin{cases} 0.7(\mathbf{s}^\mathsf{T}\mathbf{v})\left\{\frac{u}{1-(1+u)^{-n}}\right\} & \text{if } t \leq n \\ 0 & \text{if } t > n \end{cases}$$

will be paid. This information will be used to verify that the sales of electricity are sufficient to avoid defaulting on the loan.

As we have mentioned before, the price of electricity is guaranteed by the Ontario government for 20 years. Let this price be p\$ per MWh. Then taking into consideration the cost of repaying the loan, the amount of profit for the operations of year t is $P_t = p\mathbf{s}^\mathsf{T}\mathbf{X}_t - M_t$. Considering a fixed risk free interest rate of j, the accumulated value at year t is given by $L_t = \sum_{k=1}^{t}(1+j)^{t-k}P_k$, which means that the project would default from its loan obligations if $L_t < 0$ held at some point.

The future value of the electricity sold in 20 years with all debts reimbursed is represented by the random variable L_{20}, which by our assumptions follows a normal distribution with mean $\mu_{L_{20}} = pk_{20}\mathbf{s}^\mathsf{T}\mu_X - k_n(1+j)^{20-n}M_1$ and covariance $\Sigma_{L_{20}} = (pk_{20})^2\mathbf{s}^\mathsf{T}\Sigma_X\mathbf{s}$, where $k_t = \{(1+j)^t - 1\}/j$. The goal of the portfolio problem is to optimize properties of that variable. Different measures of risk could be considered, but we decided on

$$cVaR_\alpha = \mathrm{E}\{L_{20}|L_{20} < q_\alpha(L_{20})\},$$

which is the expected value of the lower quantiles of L_{20}. The idea is to minimize the risk ($cVaR_\alpha$) while achieving a certain return. Let V denote the value of the power plants in 20 years and r the equivalent annual return. To get a return of r, the initial investment of $0.3C$ should have a value of $0.3C(1+r)^{20}$ after 20 years. This

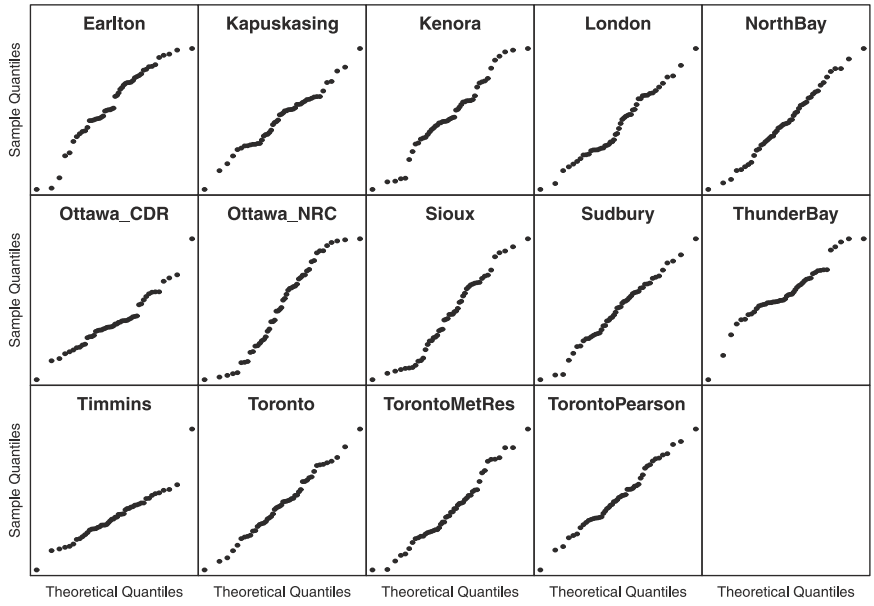

Fig. 4 QQ-plots of the yearly radiation for 14 sites in Ontario from 1955 to 2004.

expression must thus be a lower bound for the expected value of the investment in 20 years, i.e., the expected amount of money accumulated from selling electricity $\mathrm{E}(L_{20})$ and the value of the plant V.

Bringing all those equations together, we obtain the following mathematical program.

$$\begin{aligned} \min \quad & cVaR_\alpha \\ w.r.t. \quad & 0.3C(1+r)^{20} \leq \mathrm{E}(L_{20}) + V \\ & \mathbf{v}^\mathsf{T}\mathbf{s} \leq C \\ & 0 \leq s_i \leq S_i,\, i = 1,\ldots,14 \\ & P(L_t < 0) \leq \beta,\, t = 1,\ldots,20 \end{aligned}$$

The risk is thus minimized under constraints ensuring that

- a minimum annual return is obtained,
- the development budget is respected,
- development occurs within the physical limits of the potential sites,
- the probability of default is bounded from above.

As stated, this program would be difficult to solve, but some simplifications can help greatly. First, note that the probability of default is decreasing with time because of the lack of temporal dependence in the production and the fact that not having defaulted yet means we have some money set aside. For the probabilities of default, it is thus sufficient to keep only the constraint $P(L_1 < 0) \leq \beta$.

Rockafellar & Uryasev (2000) show that the minimization of $cVaR_\alpha$ is equivalent to the minimization of the variance (here a function of $\mathbf{s}^\mathsf{T}\Sigma_X\mathbf{s}$) when the target random variable is normal, which is the case under our assumptions. Our program can thus be solved by solving an equivalent quadratic program, a type of problem that can be addressed easily by a number of computer packages.

		A	B	C	D	E	F	G	H	I	J	K	L	M	N
Earlton	A	1	0.41	0.21	0.41	0.65	0.29	0.55	0.20	0.55	0.50	0.60	0.41	0.33	0.22
Kapuskasing	B	0.41	1	0.20	0.02	0.37	0.38	0.25	0.48	0.56	0.36	0.56	0.05	0.05	0.08
Kenora	C	0.21	0.20	1	0.08	0.09	-0.01	0.06	0.78	0.09	0.66	0.20	0.03	0.04	0.07
London	D	0.41	0.02	0.08	1	0.56	0.32	0.60	-0.10	0.45	0.15	0.20	0.76	0.76	0.72
NorthBay	E	0.65	0.37	0.09	0.56	1	0.23	0.49	0.06	0.52	0.40	0.35	0.42	0.29	0.33
Ottawa_CDR	F	0.29	0.38	-0.01	0.32	0.23	1	0.47	0.05	0.44	-0.02	0.59	0.50	0.48	0.50
Ottawa_NRC	G	0.55	0.25	0.06	0.60	0.49	0.47	1	-0.02	0.68	0.10	0.43	0.63	0.64	0.55
Sioux	H	0.20	0.48	0.78	-0.10	0.06	0.05	-0.02	1	0.13	0.63	0.20	-0.11	-0.14	-0.15
Sudbury	I	0.55	0.56	0.09	0.45	0.52	0.44	0.68	0.13	1	0.34	0.55	0.50	0.42	0.43
ThunderBay	J	0.50	0.36	0.66	0.15	0.40	-0.02	0.10	0.63	0.34	1	0.28	0.19	-0.05	-0.01
Timmins	K	0.60	0.56	0.20	0.20	0.35	0.59	0.43	0.20	0.55	0.28	1	0.36	0.29	0.30
Toronto	L	0.41	0.05	0.03	0.76	0.42	0.50	0.63	-0.11	0.50	0.19	0.36	1	0.73	0.73
TorontoMetRes	M	0.33	0.05	0.04	0.76	0.29	0.48	0.64	-0.14	0.42	-0.05	0.29	0.73	1	0.87
TorontoPearson	N	0.22	0.08	0.07	0.72	0.33	0.50	0.55	-0.15	0.43	-0.01	0.30	0.73	0.87	1

Table 1 Correlations of yearly radiation for 14 sites in Ontario based on data from 1955 to 2004.

In the end, the simplifying assumptions for this problem lead to a quadratic program with linear constraints, except for the constraint $P(L_1 < 0) \leq \beta$. In the examples that we solved, this constraint is seldom active. Therefore we solve the mathematical program without taking it into account and verify that the solution we have found satisfies it. The problem to be solved is therefore as the following.

$$\begin{aligned} \min \quad & \mathbf{s}^\mathsf{T} \Sigma_X \mathbf{s} \\ w.r.t. \quad & 0.3C(1+r)^{20} \leq \mathrm{E}(L_{20}) + V \\ & \mathbf{v}^\mathsf{T} \mathbf{s} \leq C \\ & 0 \leq s_i \leq S_i,\ i = 1, \ldots, 14 \end{aligned}$$

5 Solved example of a solar portfolio

To illustrate the portfolio problem, we solve the program developed in Section 4 for hypothetical parameter values that could be similar to those found in a real-life situation (where a total investment valued at $\$\, C = 20\,000\,000$ is planned). We use the data to estimate the mean and covariance of I_{it} and obtain the following values, where the 14 potential sites appear in alphabetical order (the same order as the one used in the tables and figures).

$$\hat{\mu}_I = \begin{bmatrix} 150.8 \\ 143.0 \\ 154.0 \\ 159.1 \\ 152.3 \\ 156.2 \\ 153.9 \\ 149.1 \\ 152.5 \\ 155.3 \\ 147.1 \\ 152.1 \\ 158.3 \\ 160.4 \end{bmatrix}, \quad \hat{\Sigma}_I = \begin{bmatrix} 26.6 & 12.4 & 5.6 & 11.6 & 18.9 & 8.1 & 21.4 & 5.3 & 15.5 & 12.7 & 16.1 & 13.8 & 10.8 & 6.5 \\ 12.4 & 33.6 & 6.1 & 0.6 & 12.0 & 12.2 & 10.8 & 14.5 & 17.7 & 10.2 & 16.7 & 1.8 & 1.8 & 2.6 \\ 5.6 & 6.1 & 27.5 & 2.2 & 2.8 & -0.2 & 2.3 & 21.5 & 2.7 & 17.1 & 5.3 & 1.0 & 1.3 & 2.2 \\ 11.6 & 0.6 & 2.2 & 29.8 & 17.2 & 9.6 & 24.9 & -3.0 & 13.4 & 4.0 & 5.6 & 27.0 & 26.5 & 22.8 \\ 18.9 & 12.0 & 2.8 & 17.2 & 32.0 & 7.3 & 21.0 & 1.6 & 16.1 & 11.1 & 10.3 & 15.4 & 10.7 & 10.9 \\ 8.1 & 12.2 & -0.2 & 9.6 & 7.3 & 30.2 & 19.7 & 1.4 & 13.2 & -0.6 & 16.9 & 17.8 & 17.1 & 15.8 \\ 21.4 & 10.8 & 2.3 & 24.9 & 21.0 & 19.7 & 56.7 & -0.6 & 27.8 & 3.7 & 16.8 & 30.9 & 30.9 & 24.1 \\ 5.3 & 14.5 & 21.5 & -3.0 & 1.6 & 1.4 & -0.6 & 27.6 & 3.8 & 16.2 & 5.5 & -3.7 & -4.8 & -4.6 \\ 15.5 & 17.7 & 2.7 & 13.4 & 16.1 & 13.2 & 27.8 & 3.8 & 29.5 & 9.0 & 15.6 & 17.8 & 14.6 & 13.6 \\ 12.7 & 10.2 & 17.1 & 4.0 & 11.1 & -0.6 & 3.7 & 16.2 & 9.0 & 24.3 & 7.1 & 6.1 & -1.6 & -0.3 \\ 16.1 & 16.7 & 5.3 & 5.6 & 10.3 & 16.9 & 16.8 & 5.5 & 15.6 & 7.1 & 27.0 & 12.3 & 9.8 & 9.1 \\ 13.8 & 1.8 & 1.0 & 27.0 & 15.4 & 17.8 & 30.9 & -3.7 & 17.8 & 6.1 & 12.3 & 42.2 & 30.4 & 27.5 \\ 10.8 & 1.8 & 1.3 & 26.5 & 10.7 & 17.1 & 30.9 & -4.8 & 14.6 & -1.6 & 9.8 & 30.4 & 41.3 & 32.3 \\ 6.5 & 2.6 & 2.2 & 22.8 & 10.9 & 15.8 & 24.1 & -4.6 & 13.6 & -0.3 & 9.1 & 27.5 & 32.3 & 33.4 \end{bmatrix}$$

We assume that the panels have an efficiency of $e_1 = 0.13$ and that the sites have an efficiency of $e_2 = 0.9$. These numbers can be used to calculate μ_X and Σ_X with the formulas provided in Section 4.

On the financial side, we assume that 70% of the project (i.e. $\$\,14$ million) is financed by a $n = 7$ year loan at a $u = 0.06$ interest rate. Profits will bear the risk-free interest rate of $j = 0.02$. The electricity has a price of $p = 820$\$/MWh, guaranteed through the FIT program. We assume that the development cost equals $v_i = 650$\$ per square meter for all sites. We do not impose upper bounds on the sizes of the different sites, meaning that we assume that every site considered can accommodate the panels of the whole project, requiring an area greater than three hectares.

TorontoPearson is the site that receives the most radiation with a yearly expected production of 0.164 MWh per square meter. Kapuskasing is the location receiving the least radiation with a yearly expected production of 0.147 MWh per square meter. With the parameters of our model, a 20-million dollar project in Kapuskasing would yield 85.76 million in 20 years from now, equivalent to an annual return of 14.22%. Similar calculations for TorontoPearson lead to 96.72 million dollars, equivalent to an annual return of 14.91%. In terms of annualized percentage, the returns in the best-case and worst-case scenarios are similar. Relatively speaking, however, there is a greater variation in the risk across the sites.

Figure 5 shows the expectation and risk associated with investments at each of the 14 locations considered for development. The efficient frontier corresponds to the characteristics of the optimal portfolio when the return is fixed to different values between 14.22% and 14.91%. The solution corresponding to the minimum risk is marked with a square. This solution proposes to develop 10 of the 14 sites in the proportions appearing on Figure 6.

Under the optimal solution, the probability of defaulting during the first year, $P(L_1 < 0)$, is about 10^{-73}. The constraint on the probability of default is thus inactive for any reasonable choice of β.

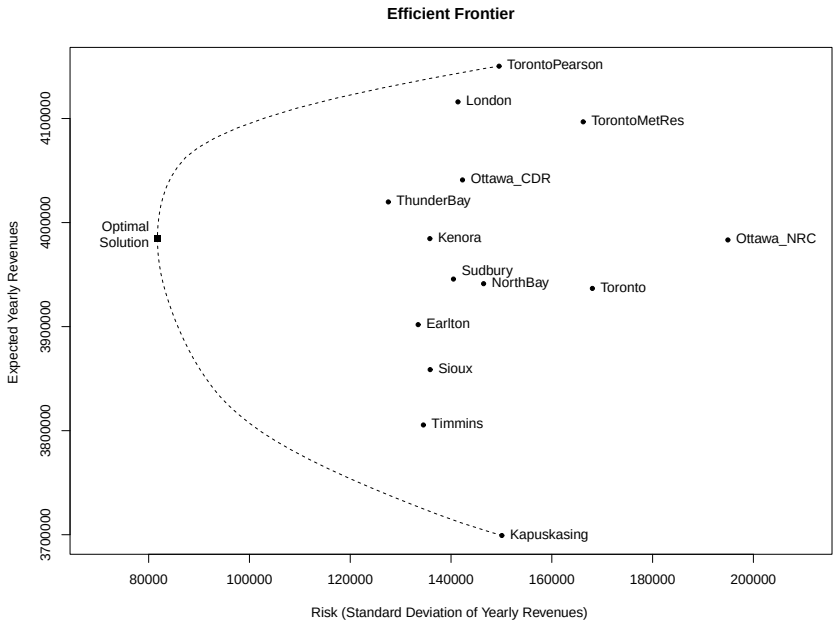

Fig. 5 Efficient frontier of the portfolio problem from a solar energy investment at 14 potential sites. The optimal solution (minimum risk) is marked with a square.

In any given year, the production of the optimal portfolio has a mean of 4859 MWh (with a standard deviation of 99.7 MWh), yielding a gross revenue of 3.98 million (with a standard deviation of 82 thousand dollars). At the end of the 20-year period, the accumulated profits, the interest they bore, and the salvage value of the plant add up to 92.7 million dollars (with a standard deviation of 1.99 million), which is equivalent to a 14.67% annual rate of return on the initial amount (6 million dollars) that was invested in the project.

Note that during the industrial workshop, Matlab was used to solve the portfolio problem. For the redaction of this report, the package quadprog of R was used to recreate the initial calculations and produce additional figures. Quadratic program with linear constraints arise in many different contexts and can be solved relatively easily by using several other software.

6 Discussion

Investments in solar energy are often made within the context of the simultaneous development of multiple sites. We have described the problem of selecting a port-

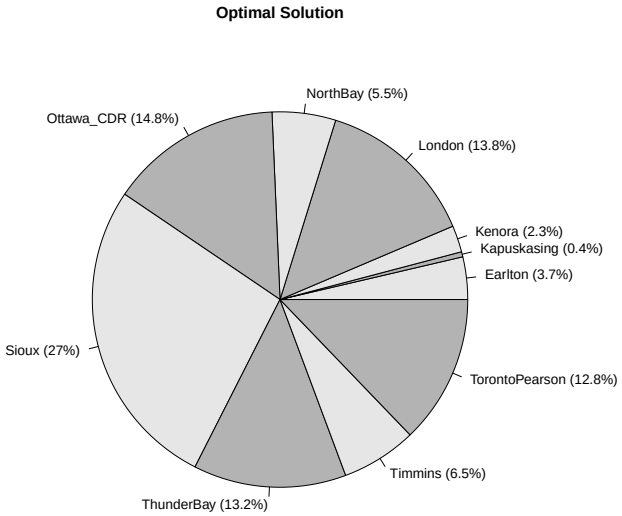

Fig. 6 Distribution of the investment for the optimal portfolio solution (minimum risk).

folio of investments and modelled it through a mathematical program, which we solved for a hypothetical case.

The variations in solar radiation were modelled by a multivariate normal distribution. We used data from Environment Canada to calibrate that distribution. The examination of the same data provided empirical evidence that the normality assumption is reasonable.

Although the normality assumption is based on empirical evidence, it plays an important role in the simplification of the problem. For instance the distribution of the total annual production and the final value of the projects are also normally distributed, and the problem of minimizing $cVaR_\alpha$ can be replaced by a simpler equivalent quadratic problem.

Another interesting feature of the data is that it shows no temporal dependence. At least for this data set, we did not observe an impact of the solar cycle or other natural variations in the Earth's climate.

The portfolio is designed with Ontario's FIT program in mind and its guarantee of a fixed rate for the electricity produced. This has major simplifying consequences. If the portfolio problem had to be solved for an open market environment without a fixed price schedule, a reliable pricing model would be required and the fluctuations in irradiation would have to be modelled at a more refined scale, creating serious technical complications.

Maintenance costs were neglected in the formulation of the problem but it is straightforward to include them into the model, and the latter is not more difficult to solve than before. On the other hand, including into the model a fixed cost for the development of a new site has major consequences for the model complexity. Indeed this modification implies that not all of the new constraints are linear; the resulting model is a mixed integer quadratic program, which is much more difficult to solve than a quadratic program with linear constraints.

In this article we have solved the portfolio problem for a group of 14 sites with data provided with Environment Canada. Turquoise Technology Solutions Inc. processes satellite images that can provide useful data as well, and they will be able to solve similar problems for many other locations in Canada and the United States.

References

1. M. Bazilian & F. Roques (2008). *Analytical Methods for Energy Diversity and Security: Portfolio Optimization in the Energy Sector : a Tribute to the Work of Dr. Shimon Awerbuch*, Elsevier, Amsterdam, The Netherlands, 321 p.
2. J. Boland (2008). Time series modelling of solar radiation, *Modeling Solar Radiation at the Earth's Surface: Recent Advances, Viorel Badescu ed.*, Springer-Verlag, Berlin, 517 p.

3. J. Davies & J. Binyamin (2008). Modelling UV-B irradiance in Canada, *Modeling Solar Radiation at the Earth's Surface: Recent Advances, Viorel Badescu ed.*, Springer-Verlag, Berlin, 517 p.
4. B. Drake & K. Hubacek (2007). What to expect from a greater geographic dispersion of wind farms?—A risk portfolio approach, *Energy Policy*, **35**, 3999–4008.
5. C. Genest & A.-C. Favre (2007). Everything you always wanted to know about copula modeling but were afraid to ask. *Journal of Hydrologic Engineering*, **12**, 347–368.
6. H. Markowitz (1952). Portfolio selection, *Journal of Finance*, **7**, 1, 77–91.
7. R. T. Rockafellar & S. Uryasev (2000). Optimization of Conditional Value-at-Risk, *Journal of Risk*, **2**,21–41.
8. A. Zaharim, A. M. Razali, T. P. Gim & K. Sopian (2009). Time series analysis of solar radiation data in the Tropics, *European Journal of Scientific Research*, **25**, 4, 672–678.