

Real-Time Optimization for Kites

Sean Costello* Grégory François* Dominique Bonvin*

* *École Polytechnique Fédérale de Lausanne, Switzerland (e-mails: sean.costello@epfl.ch, gregory.francois@epfl.ch, dominique.bonvin@epfl.ch).*

Abstract: Over the past decade, a large number of academics and start-ups have devoted themselves to developing kites as a renewable energy source. Determining the trajectories the kite should follow is a modeling and optimization challenge. We present a dynamic model and analyse how uncertainty affects the resulting optimization problem. We show how measurements can be used to rapidly correct the model-based optimal trajectories in real time. This novel real-time optimization approach does not rely on intensive online computation. Rather, it uses knowledge of the structure of the optimal solution, which can be studied offline.

Keywords: Kite Power, Optimal Control, Real-Time Optimization, Kite Dynamics.

1. INTRODUCTION

The concept of using kites, or airplanes on tethers, as a renewable energy source has received increasing attention in recent years, both from industry and academia. A kite is a wing (just like the blades of a turbine, the wings of an airplane or the sails of a yacht) that is connected to the ground by a cable. Kites can fly at speeds many times that of the the wind. At such high speeds a large aerodynamic force acts on the kite, and the force is transmitted to the ground via the cable. A number of different ways of using kites as an energy source are being investigated by ‘kite-power’ start-ups (Skysails GmbH, 2012; Makani Power, 2012; Ampyx Power, 2012). For example, the cable can be wound around a reel whose shaft is connected to a generator. When the tension in the cable forcibly unwinds the reel, the generator produces electricity. In a second phase, the generator works as a motor to reel the cable back in. The kite flies a trajectory designed to maximize the cable tension during the reel-out phase, and to minimize it during the reel-in phase, resulting in a net production of power. Kites are thought to have two main advantages over conventional wind turbines:

- The cable transmits the forces directly to the ground, so there is no need for a costly tower.
- Kites can fly at altitudes of several hundred meters, where the wind is stronger and more stable (Thuillier and Lappe, 1964). The highest wind turbine towers stand 135 metres tall.

The trajectory the kite follows directly determines the aerodynamic force on the kite, and thus the power the kite generates. Determining the optimal trajectory is a periodic dynamic optimization problem, which can be solved using a mathematical model of the system. However, constructing an accurate dynamic model of a kite is still an open problem. Standard airplane dynamic models cannot be used to model a kite, as the cable inhibits many of the dynamics found in un-tethered flight, and introduces new ones (Breukels, 2010). Various models have been proposed



Fig. 1. A superimposed sequence of frames shows a kite following a closed path during testing of EPFL’s lab-scale system.

(Diehl, 2001; Erhard and Strauch, 2012; Argatov et al., 2009; Breukels, 2010) and some comparisons have been made with experimental data (Erhard and Strauch, 2012; Dadd et al., 2010; Breukels, 2010). While those authors were able to glean valuable insight about certain aspects of a kite’s behavior from these experiments, no one model has been shown to very accurately reproduce a kite’s behavior. Even were an accurate model available, the kite’s behavior will depend on external time-varying disturbances, mainly the wind speed and direction. These disturbances cannot be directly measured as, due to wind shear, the wind at the kite’s altitude will generally not be the same as the wind measured at ground level.

Real-Time Optimization (RTO) is a family of techniques for achieving optimal operation despite uncertainty. Measurements are used to counteract modeling error and disturbances. The measurements contain information about the real process and, if they are correctly interpreted, can be used to correct the model-based trajectories during

operation. RTO for dynamic optimization problems has often been applied to batch chemical reactors (Srinivasan et al., 2003; François et al., 2004), as these are repeated (periodic) optimization problems. RTO is particularly applicable to these problems because measurements from past iterations can be used to improve the performance of the current one. Luckily, the optimization problems in kite power are also periodic, as the kite will typically trace out the same path in the sky many times over. This means that many of the ideas from RTO for chemical problems can be carried over to the problem of determining optimal trajectories for kites. Indeed, some characteristics of the kite problem make it even easier to do RTO:

- The period of the kite’s trajectory is measured in seconds, rather than hours. This is an advantage because measurements quickly become available, however on-line computations must be fast.
- A large number of continuous measurements are available, such as position, velocity, orientation, wind-speed at the ground, and line tension.

The first contribution of this article is to review the kite-modeling literature and present a modified dynamic model of a kite that is suitable for dynamic optimization (Section 2). We focus on the trajectory optimization problem that arises when the kite pulls a boat, a system that has been implemented on several large German vessels (Skysails GmbH, 2012). As other kite-power trajectory optimization problems are usually a variation of this basic problem, the technique we propose could be adapted to other kite-power setups. An analysis of how *uncertainty* affects the optimal solution is presented in Section 3. The second contribution is a novel RTO scheme for repeated dynamic optimization (Section 4). In Section 5 the RTO scheme is applied to the kite and boat system in simulation, achieving optimal performance in a matter of minutes. The difficult path following control problem is not addressed in this paper. In order to focus on RTO, an ideal (simulation-based) controller is used.

2. PROBLEM FORMULATION

2.1 Dynamic Model

A number of kite models have been proposed in the literature. The simplest model assumes the kite has no mass (Erhard and Strauch, 2012; Argatov et al., 2009). A slightly more complex model considers the mass of the kite to be lumped (Diehl, 2001). Embellishments of the lumped-mass model take into account tether drag (Canale et al., 2010; Houska and Diehl, 2006) and rotational inertia (Breukels, 2010). Far more complex models have also been developed, such as the multi-body kite simulator developed by Breukels (2010). For trajectory optimization we choose to work with the zero-mass model. Dadd et al. (2010) showed experimentally that this model gives a reasonable prediction of line forces during dynamic flight, while Erhard and Strauch (2012) experimentally identified a turning law for this model that agrees well with measured data. The generated power in our optimization problem is based on line forces, which depend chiefly on the position of the kite. It is mainly the turning behavior that determines how the kite’s steering input (the manipulated

variable) influences the kite’s position. Thus, although we have chosen a simple model, it should reasonably accurately represent the dependence of power generation on how the kite is steered.

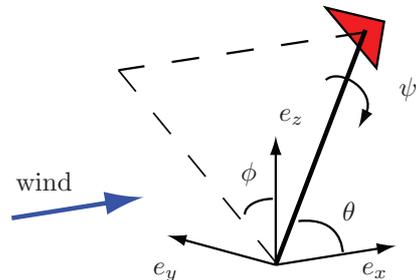


Fig. 2. The spherical co-ordinate system used.

We first give the dynamics of a kite attached to a fixed point by a tether of fixed length, as shown in Figure 2. The model is copied almost exactly (our co-ordinate system differs slightly) from Erhard and Strauch (2012), where a full derivation can be found. Here we merely describe the underlying physical assumptions. We use a spherical co-ordinate system with the tether point as the origin (Figure 2). The zenith angle, θ , is the angle between the wind vector and the tether. The azimuth angle, ϕ , is the angle between the vertical (z axis) and the tether’s projection onto the y - z plane. As the kite will not fly upwind of its tether point, it is free to move on a quarter sphere: $\theta \in [0, 90^\circ]$, $\phi \in [-90^\circ, 90^\circ]$. A third angle, ψ , defines the kite’s orientation. This is the angle by which the kite’s longitudinal axis is rotated around the tether (see Erhard and Strauch (2012) for an exact definition). The reference angle, $\psi = 0$, occurs when the kite is aligned with the longitudinal lines running from the downwind position to the windward edges of the quarter-sphere. The velocity of the kite on the sphere is given by $\dot{\theta}$ and $\dot{\phi}$. These two degrees of freedom are determined such that the following two assumptions are satisfied:

- (1) The velocity is such that the angle between the kite’s longitudinal axis and the apparent wind (the air velocity relative to the kite) is minimized. This is comparable to the way an airplane’s tail keeps it pointing towards the oncoming air.
- (2) The kite’s inertia is negligible compared with the aerodynamic force acting upon it. Thus, when it is acted upon by a net external force, it accelerates infinitely fast, instantaneously reaching a velocity at which there is no longer a net force acting upon it. This means the aerodynamic force must always be directly opposed by the line tension (as these are the only forces acting on the kite). By applying basic aerodynamic theory (Anderson and Eberhardt, 2000) (which necessitates several more assumptions), we can conclude that the angle between the apparent wind and the kite’s tether is constant. This angle depends on an aerodynamic coefficient: the kite’s glide ratio E .

Under these two assumptions the velocity of the kite is given by:

$$\dot{\theta} = \frac{v_a}{L} \left(\cos \psi - \frac{\tan \theta}{E} \right), \quad (1)$$

$$\dot{\phi} = -\frac{v_a}{L \sin \theta} \sin \psi, \quad (2)$$

where L is the tether length and v_a is the magnitude of the apparent wind vector projected onto the quarter sphere's tangent plane. The later is given by:

$$v_a = v_0 E \cos \theta, \quad (3)$$

where v_0 is the wind speed. Note that ϕ does not affect the kite's dynamics; having neglected mass, and thus gravity, the kite's behavior is unchanged by rotation around the wind vector.

While the differential equations for $\dot{\theta}$ and $\dot{\phi}$ (the kite's velocity) follow directly from physical assumptions, the equation for $\dot{\psi}$ (the turning behavior) does not. Erhard and Strauch (2012) experimentally validated the law:

$$\dot{\psi} = g v_a u + \dot{\phi} \cos \theta, \quad (4)$$

where g is a constant and $u(t)$ is the steering deflection applied to the kite. Using the expression for v_a (3) and Assumption 2, we derive an expression for the line tension¹:

$$T = P_D A \cos^2 \theta (E + 1) \sqrt{E^2 + 1}, \quad (5)$$

where $P_D = \rho v_0^2 / 2$ is the wind's dynamic pressure and A is the area of the kite. Note that the line tension depends on the angle between the tether and the wind, θ , and the glide ratio, E . In Erhard and Strauch's model E is constant. This is a reasonable approximation if the model is used for controller design, but not if it is used for optimization. A constant E would imply that the tension in the tether is maximized if the kite is kept as directly downwind as possible, regardless of the large steering deflections which are required to keep it there. In our experiments, it was observed that steering deflections cause a reduction in tether tension. This is particularly noticeable for very large steering deflections, which will almost cause the kite to stall, drastically reducing the apparent windspeed, and hence the tether tension. We propose the following law to model this behaviour:

$$E = E_0 - c u^2, \quad (6)$$

where c is a constant that determines how much the kite's glide ratio is penalized for a steering deflection. The basic idea is that, unlike the way an airplane steers (by rolling), a kite steers by increasing the angle of attack on one side of the kite. This will decrease the kite's overall glide ratio. A detailed justification goes beyond the scope of this paper.

Consider that the tether is attached to a boat, which is moving at constant velocity. This is a reasonable assumption, because (i) the boat's huge inertia will dampen any inconsistency in the kite's pull, and (ii) the current industrial system is designed to supply only a fraction of the boat's thrust (Skysails GmbH, 2012). Figure 3 illustrates the relationship:

$$\mathbf{V}_{wb} = \mathbf{V}_{wind} - \mathbf{V}_{boat}, \quad (7)$$

where \mathbf{V}_{wb} is the wind relative to the boat. In our kite model, $v_0 = |\mathbf{V}_{wb}|$, and the x-axis in Figure 2 is aligned with \mathbf{V}_{wb} . We are interested in maximizing the component

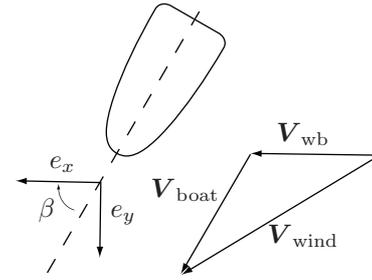


Fig. 3. Vector diagram illustrating how the wind experienced on a moving boat depends on the wind (relative to the water) and the boat's velocity.

of the tether tension in the boat's direction of motion. The unit vector in this direction is:

$$e_{wb} = (\cos \beta) e_x + (\sin \beta) e_y, \quad (8)$$

while the unit vector aligned with the tether tension is:

$$e_T = (\cos \theta) e_x + (\sin \theta \sin \phi) e_y + (\sin \theta \cos \phi) e_z. \quad (9)$$

The component of the tether force in the boat's direction of movement (referred to as thrust) is the the dot product of these vectors:

$$T_F = T (\cos \theta \cos \beta + \sin \theta \sin \beta \sin \phi). \quad (10)$$

2.2 Optimization Problem

How should the kite be controlled to maximize the average thrust? The input and state *trajectories* (functions of time), $u(t)$ and $x(t) = [\theta(t) \ \phi(t) \ \psi(t)]^T$, must be periodic, i.e. the kite must return to its original position at the end of each cycle. In addition, the kite must respect a minimum altitude constraint to ensure it will not crash in the case of disturbances. The dynamic optimization problem can be formulated mathematically as

$$\underset{\tilde{u}(t), t_f, x_0}{\text{maximize}} \quad \bar{T}_F = \frac{1}{t_f} \int_0^{t_f} T_F dt \quad (11)$$

$$\text{subject to} \quad \dot{x} = \begin{bmatrix} \dot{\theta} \\ \dot{\phi} \\ \dot{\psi} \end{bmatrix} = \begin{bmatrix} \frac{v_a}{L} \left(\cos \psi - \frac{\tan \theta}{E} \right) \\ -\frac{v_a}{L \sin \theta} \sin \psi \\ \frac{v_a}{L} \tilde{u} + \dot{\phi} \cos \theta \end{bmatrix}, \quad (12)$$

$$v_a = v_0 E \cos \theta, \quad (13)$$

$$E = E_0 - \tilde{c} \tilde{u}^2, \quad (14)$$

$$T_F = \left(P_D A \cos^2 \theta (E + 1) \sqrt{E^2 + 1} \right) \times (\cos \theta \cos \beta + \sin \theta \sin \beta \sin \phi), \quad (15)$$

$$x(0) = x_0, \quad (16)$$

$$x(t_f) = x_0, \quad (17)$$

$$L \sin \theta \cos \phi \geq h_{min}, \quad (18)$$

where t_f is the final time, \bar{T}_F is the average thrust, x_0 is the initial state vector and h_{min} is the kite's minimum permissible altitude. We used the following substitutions to simplify the problem formulation:

$$\tilde{u} = g L u, \quad (19)$$

$$\tilde{c} = \frac{c}{(g L)^2}. \quad (20)$$

Equations (11) to (18) consist of a) the integral cost, b) the dynamic system equations, and c) the periodicity and height constraints (the dashed lines indicate this separation). The solution to this problem depends on

¹ To do so we assume the kite operates with a lift coefficient of 1. While this may not be exactly true, it only introduces a proportional error in T , which does not affect the kite's optimal path.

numerous model parameters, the nominal values of which are given in Table 1. These parameters have been chosen to be consistent with the current commercial system (g and c were estimated). Among these we distinguish the uncertain parameters, $\pi = [\tilde{c}, v_0, E_0, \beta]$, that are prone to variations or inaccuracy.

Table 1. Nominal and perturbed parameters.

Parameter	Units	Nominal value		
L	m	400		
A	m ²	300		
h_{min}	m	100		
c	-	0.25	Perturbed scenario	
g	m ⁻¹	7.5×10^{-3}		
π	\tilde{c}	-	0.028	.005
	v_0	m s ⁻¹	10	15
	E_0	-	5	6.5
	β	degrees	0	15

The nominal optimal path (the kite's position on the sphere) is shown in Figure 4. The kite follows a roughly circular path. This allows it to stay close to the downwind position in which it generates maximum thrust. Note that it takes only 11 seconds for the kite to complete one period, as it flies at approximately 130 km/h.

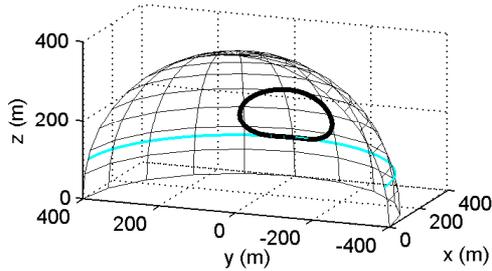


Fig. 4. Nominal optimal path on the sphere the kite is constrained to (the direction of flight is clockwise). The blue line is the height constraint.

3. UNCERTAINTY ANALYSIS

We now investigate how the optimal solution varies along with π , the vector of uncertain parameters. The windspeed v_0 is the greatest source of uncertainty. Interestingly, the windspeed does not affect the optimal *path* the kite should follow on the sphere. This is because the kite's dynamic equation is proportional to v_0 . Thus, an increase in windspeed simply accelerates the optimal input and state trajectories in time (and scales the cost function).

Figure 5 shows the optimal paths that are obtained when the parameter \tilde{c} is varied. There is a loss in optimality if the system follows the nominal optimal path (in red), as depicted in Figure 6. To illustrate this we use an ideal controller, which is capable of following the nominal path very precisely (of course, such a controller can only be achieved in simulation). The optimality loss can reach 10% of the thrust for unexpectedly small values of \tilde{c} . If \tilde{c} is much larger than the nominal value it becomes impossible to track the nominal optimal path, as the kite is not maneuverable enough.

The same analysis can be carried out for the two remaining uncertain parameters, β and E_0 . Varying β changes the optimal trajectory significantly (Figure 7), yet the optimality loss from tracking the nominal path is tiny (less than 1%). The effect of E_0 on the dynamics is similar to that of v_0 , thus it causes very little change in the optimal path (not shown), and tracking the nominal path leads to negligible optimality loss. However, if E_0 is much smaller than the nominal value, the trajectory-tracking problem becomes infeasible. The results of the uncertainty analysis are summarized in Table 2.

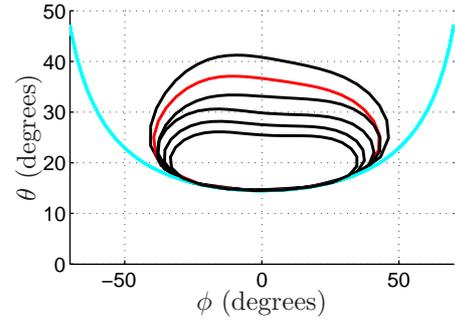


Fig. 5. The optimal path as \tilde{c} varies from 0.0041 to 0.0460 (the optimal path dilates as \tilde{c} increases). The blue line is the height constraint.

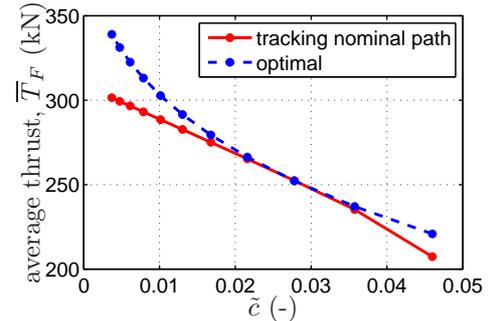


Fig. 6. The optimal thrust and the thrust obtained by following the nominal path, for different values of \tilde{c} .

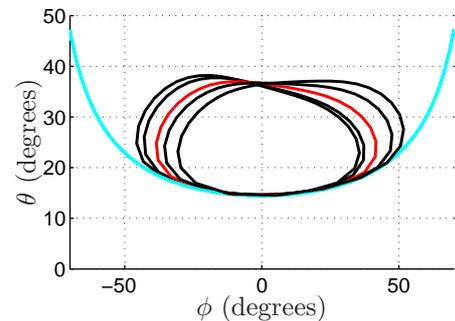


Fig. 7. The optimal path as β , the angle between the wind relative to the boat and the boat's velocity, varies from -15° to 15° .

4. REAL-TIME OPTIMIZATION METHODOLOGY

Skysails, the company that developed kite power for boats, uses three manipulated variables to define the kite's trajec-

Table 2. Summary of the uncertainty analysis.

Uncertain parameter	v_0	\tilde{c}	β	E_0
effect on optimal path	none	large	large	small
effect on optimal thrust	large	small	tiny	large
potential optimality loss	none	10%	1%	1%
tracking can become infeasible	no	yes	no	yes

tory (Erhard and Strauch, 2012). These could be adjusted online to maximize the thrust (perhaps Skysails does so), but the resulting trajectory would probably be far from ‘optimal’ with respect to the optimization problem formulated in Section 2.2. We now propose a RTO technique that ensures this optimality.

Let us begin with the mathematical reasoning behind the RTO scheme. The solution to the dynamic optimization problem is the optimal path $z^*(l, \pi) = [\theta^*(l, \pi) \ \phi^*(l, \pi)]^T$, where $l \in [0, l_f]$ is the path length. The optimal path can be decomposed into two intervals:

$$z^*(l, \pi) = \begin{cases} z_{UC}^*(l, \pi) & \text{for } l \in [0, l_s] \\ h = h_{min} & \text{for } l \in]l_s, l_f] \end{cases} \quad (21)$$

where $z_{UC}^*(l, \pi)$ is the path during the *unconstrained* interval (when the height constraint is inactive), l_s is the point at which the constraint becomes active and $h := L \sin \theta \cos \phi$ is the kite’s altitude. It can be shown that in general, for parametric optimal control problems, small variations of the uncertain parameters π do not alter the sequence and type of intervals in the solution (Maurer and Büskens, 2001). In our case, simulation studies confirmed that this also holds for larger variations of π . This allows us to develop a Taylor series for the optimal path around the nominal parameter values, π_0 :

$$z^*(l, \pi) \simeq \begin{cases} z_{UC}^*(l, \pi_0) + \left. \frac{\partial z_{UC}^*(l, \pi)}{\partial \pi} \right|_{\pi_0} \Delta \pi & \text{for } l \in [0, l_s] \\ h = h_{min} & \text{for } l \in]l_s, l_f] \end{cases} \quad (22)$$

The Taylor-series analysis yields the directions, $\left. \frac{\partial z_{UC}^*(l, \pi)}{\partial \pi} \right|_{\pi_0}$, in which the path should be adapted in response to a perturbation in the parameters, $\Delta \pi = \pi - \pi_0$. The problem is that the magnitude of the adaptation in each direction is unknown. We propose to find this using real-time optimization.

A block diagram of our iterative RTO scheme is shown in figure 8. There are three components: (i) the Path Generator, which maps a vector of decision variables, α_k , to the reference path, $z_k(l, \alpha_k)$, for each period, k , (ii) the Controller, which tracks this path, and (iii) the Periodic Optimizer, which adjusts the decision variables α_k once per period, in order to eventually maximize the thrust. At the end of each period the resulting average thrust, $\bar{T}_{F,k}$, is fed back to the Periodic Optimizer.

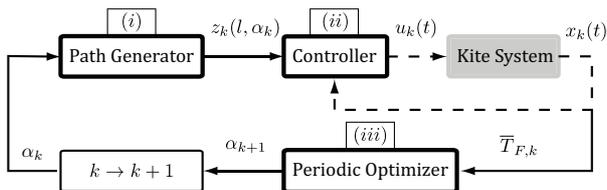


Fig. 8. The iterative RTO scheme.

The path generator is based on the Taylor-series expansion (22). The vector of decision variables α_k determines the magnitude of the adaptation in the directions given by the expansion ²:

$$z_k(l, \alpha_k) := \begin{cases} z_{UC}^*(l, \pi_0) + \left. \frac{\partial z_{UC}^*(l, \pi)}{\partial \pi} \right|_{\pi_0} \alpha_k, & l \in [0, l_s] \\ h = h_{min}, & l \in]l_s, l_f] \end{cases} \quad (23)$$

The periodic optimizer solves an unconstrained optimization problem at each iteration. It chooses the decision variables α_{k+1} for the next period. From its point of view, each period is a function evaluation, and it can be any unconstrained optimization algorithm. In this paper we use a standard technique from numerical optimization, the Gradient Descent method. This method is attractive for RTO because it generally requires far fewer function evaluations to converge than more sophisticated algorithms. An unconstrained optimization algorithm suffices because the path generator only produces closed paths satisfying the height constraint.

The optimization method requires the experimental evaluation of the gradient of \bar{T}_F with respect to α . So, in general, the number of periods until the entire scheme converges will increase (at least) linearly with the dimension of α . As the dimension of α equals that of π , we must focus on a subset of the uncertain parameters if we wish to speed up the RTO scheme, while rejecting the effect of the most significant disturbances. The path generator then uses a Taylor-series expansion with respect to these ‘important’ parameters alone. Indeed, we do not assume the other uncertain parameters to be fixed, rather we deem it unnecessary to react to their variations.

5. RTO RESULTS

The RTO scheme was applied to the kite and boat system in simulation. In order to focus on the RTO algorithm without the results being biased by control error we again used the ‘ideal’ path-following controller. The Taylor-series expansion (22) was performed with respect to \tilde{c} , as (i) it is the only uncertain parameter whose variation incurs large optimality loss, and (ii) the resulting Taylor expansion allows the RTO scheme to be initialized at a path the controller can easily track. The result was less than 1% optimality loss for a wide range of uncertainties. Here we illustrate one scenario (labelled ‘Perturbed scenario’ in Table 1) in which all the uncertain parameters are significantly different from their nominal values.

The evolution of the average thrust per period is shown in Figure 9. Initially the kite generates about 25% less thrust than is possible. It is initialized with a large α_0 , which gives a dilated trajectory that is easy to track. After 15 periods the RTO scheme has reduced the optimality loss to about 1%. Interestingly, the path the RTO algorithm converges to (Figure 10) is quite different from the optimal path, yet this difference only results in a negligible optimality loss.

² In order for the path to be continuous, a continuity condition must be enforced at the junction between the two intervals. Methods for performing this slight correction, which is due to neglected higher-order terms in the Taylor expansion, are discussed by Pesch (1989).

This is due to the variation in β , which we (correctly) decided can be neglected by the RTO algorithm based on the uncertainty analysis in Section 3.

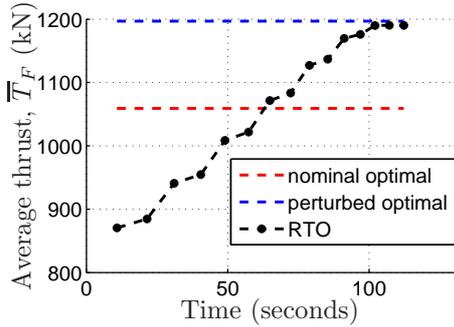


Fig. 9. Convergence of the RTO scheme from a conservative initial path. The red line is the thrust that would be obtained by following the nominal optimal path.

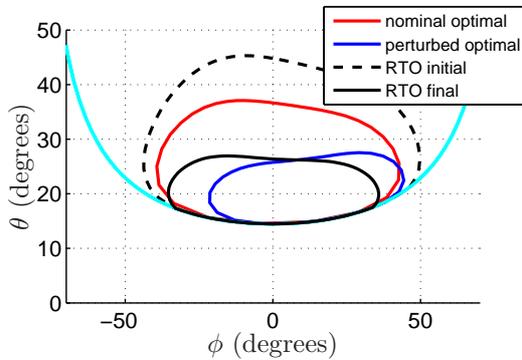


Fig. 10. The initial and final paths for the RTO scheme.

6. CONCLUSION

The aim of this article is to illustrate the usefulness of RTO for maximizing the efficiency of kite-power systems. Although we analyse a kite towing a boat, the technique we propose can readily be adapted to electricity-generating kites. We use the only partially validated, low-dimensional dynamic kite model available in the literature. An equation is proposed to model the negative influence of steering deflections on the tether tension, a phenomenon we have observed in practical tests. Unfortunately, the model differs substantially from reality, due to a number of uncertain parameters. The main source of uncertainty is the wind-speed, but luckily the optimal path the kite should follow is invariant with respect to changes in the windspeed. Hence, we concentrate on calculating this optimal path, with the assumption that a path-following controller tracks it. Our analysis indicates that the variation of one particular parameter is likely to cause the largest optimality loss. This parameter \tilde{c} models the kite's turning behavior. We propose a simple RTO scheme to reject the effect that any variation in this parameter has on the thrust.

An alternative approach would be to systematically estimate the uncertain parameters. The new estimates could then be used to recalculate the optimal path. This leads to a difficult estimation problem, as all the parameters must be estimated simultaneously, even though the only

one of real interest for optimization is \tilde{c} . Our RTO scheme essentially allows us to avoid this estimation problem, as we indirectly estimate \tilde{c} (which ideally should correspond to α), through its influence on the thrust.

Finally, we note that the successful implementation of this approach on a practical system (work in progress) depends very much on a functioning path controller. Although many simulation studies exist, at the point of writing, no practical implementation of such a controller has been reported in the literature. This would be an important step forward for kite power. For the same reason, we cannot claim our approach outperforms Skysail's. While in theory it achieves more thrust, it may be more difficult to implement in practice.

REFERENCES

- Ampyx Power (2012). <http://www.ampyxpower.com>.
- Anderson, D. and Eberhardt, S. (2000). *Understanding Flight*. McGraw-Hill Professional.
- Argatov, I., Rautakorpi, P., and Silvenoinen, R. (2009). Estimation of the mechanical energy output of the kite wind generator. *Renewable Energy*, 34(6), 1525 – 1532.
- Breukels, J. (2010). *An Engineering Methodology for Kite Design*. Ph.D. thesis, Delft Universitij of Technology.
- Canale, M., Fagiano, L., and Milanese, M. (2010). High altitude wind energy generation using controlled power kites. *IEEE Transactions on Control Systems Technology*, 18(2), 279 –293.
- Dadd, G.M., Hudson, D.A., and Sheno, R.A. (2010). Comparison of two kite force models with experiment. *Journal of Aircraft*, 47, 212–224.
- Diehl, M. (2001). *Real Time Optimization for Large Scale Nonlinear Processes*. Ph.D. thesis, Ruprecht-Karls-Universität Heidelberg.
- Erhard, M. and Strauch, H. (2012). Control of towing kites for seagoing vessels. *arXiv:1202.3641v1*.
- François, G., Srinivasan, B., Bonvin, D., Hernandez Barajas, J., and Hunkeler, D. (2004). Run-to-run adaptation of a semiadiabatic policy for the optimization of an industrial batch polymerization process. *Industrial & Engineering Chemistry Research*, 43(23), 7238–7242.
- Houska, B. and Diehl, M. (2006). Optimal control of towing kites. In *45th IEEE Conference on Decision and Control*, 2693 –2697.
- Makani Power (2012). <http://www.makanipower.com>.
- Maurer, H. and Büskens, C. (2001). Sensitivity analysis and real-time control of parametric optimal control problems using boundary value methods. *Online Optimization of Large Scale Systems*, 17–56.
- Pesch, H.J. (1989). Real-time computation of feedback controls for constrained optimal control problems, part 1: Neighbouring extremals. *Optimal Control Applications and Methods*, 10, 129–145.
- Skysails GmbH (2012). <http://www.skysails.info>.
- Srinivasan, B., Bonvin, D., Visser, E., and Palanki, S. (2003). Dynamic optimization of batch processes II. Role of measurements in handling uncertainty. *Computers & Chemical Engineering*, 27(1), 27–44.
- Thuillier, R.H. and Lappe, U.O. (1964). Wind and temperature profile characteristics from observations on a 1400ft tower. *Journal of Applied Meteorology*, 3, 299 –306.