

SIMULTANEOUS PERTURBATION STOCHASTIC APPROXIMATION FOR REAL-TIME OPTIMIZATION OF MODEL PREDICTIVE CONTROL

Irina Baltcheva
Felisa J. Vázquez-Abad ¹(member of GERAD)
Université de Montréal
DIRO, CP 6128 succ Centre Ville,
Montreal, QC, H3C 3J7 Canada
{baltchei,vazquez}@iro.umontreal.ca

Smaranda Cristea
César De Prada
Universidad de Valladolid
DISA, P Prado de la Magdalena s/n,
Valladolid, 47005 Spain
{smaranda,prada}@autom.uva.es

KEYWORDS

SPSA, Non-linear Model Predictive Control.

ABSTRACT

The aim of this paper is to suggest a global optimization method applied to a model predictive control problem. We are interested in the Van der Vusse reaction found in multiple chemical processes. The control variables (temperatures and input flow rates) are real time continuous processes and a target reference level must be reached within certain operational constraints. The canonical model discretises time using a sampling interval, thus translating the control problem into a non-linear optimization problem under constraints. Because of the non-linearity of the cost function, common methods for constrained optimization have been observed to fail in experiments using ECOSIM to simulate the production process. The controllers do not achieve their optimal values and the numerical optimization based on approximating gradients and Hessians cannot be performed in real time for the operation of the plant to be successful. In this research we implement a methodology for global optimization adding noise to the observations of the gradients, which will perform much better than deterministic methods.

INTRODUCTION

Model Predictive Control (MPC) is now recognized in the industrial world as a proven technology, capable of dealing with a wide range of multivariable constrained control problems. Nevertheless, most of the industrial controllers are based on linear internal models which limits its applicability. Because of it, non-linear model predictive control (NMPC) has received a lot of attention in the latest years, both from

the point of view of its properties [2] and implementation. Referring to this last aspect, the main drawback is the computational burden that NMPC implies. While linear MPC with constraints can solve the associated optimization problem each sampling time using QP or LP algorithms for which very efficient codes are available, NMPC relies on non-linear programming (NLP) methods such as SQP, that are known to be far more CPU demanding. Several schemes have been proposed to deal with this problem, among them the well known sequential and simultaneous approaches. ¹

For sequential solutions, the model is solved by integration at each iteration of the optimization routine. Only the control parameters remain as degrees of freedom in the NLP. Simulation and optimization calculations are performed sequentially, one after the other. The approach can easily be coupled with advanced simulation tools. In contrast, simultaneous model solution and optimization includes both the model states and controls as decision variables and the model equations are appended to the optimization problem as equality constraints. This can greatly increase the size of the optimization problem, leading to a trade-off between the two approaches. In both cases, computation time remains a difficulty in order to implement NMPC in real processes.

This paper shows a global optimization method oriented to reduce the difficulties associated with the computation of the gradients, in order to facilitate the implementation of NMPC algorithm, using the sequential approach, applied to a benchmark problem: Van der Vusse reactor.

MODEL DESCRIPTION

The Van der Vusse reaction is described in detail in [4] and the references therein. To summarize, there is a substance A in input, a chemical reaction and a substance B in output.

We'll denote by:

c_B the concentration of product B (controlled)

¹Also Principle Investigator, Department of Electrical and Electronic Engineering, The University of Melbourne.

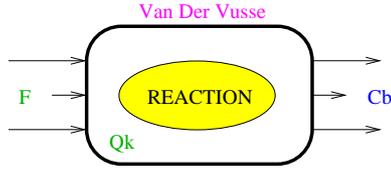


Figure 1: Van der Vusse

- c_A the concentration of product A (measured)
- T the temperature in the reactor (measured)
- T_K the temperature in the coolant (measured)
- F the input flow of product A (manipulated)
- Q_K the heat removal (manipulated)
- $c_{A0}, T_0, k_1, k_2, k_3, k_\omega, \rho, c_p, V_R, \Delta H_1, \Delta H_2, \Delta H_3, m_K$ constants.

To the input we associate the **control variables** F and Q_K . The **controlled variable** c_B is associated with the output. Thereby, we want to control the output c_B by manipulating the input F and Q_K .

To describe the evolution of this dynamical system, we introduce the non-linear differential equations related to mass and energy conservation:

$$\begin{aligned}
 c_A' &= \frac{F}{V_R}(c_{A0} - c_A) - k_1 c_A - k_3 c_A^2 \\
 c_B' &= -\frac{F}{V_R} c_B + k_1 c_A - k_2 c_B \\
 T' &= \frac{F}{V_R}(T_0 - T) - \frac{1}{\rho C_p}(k_1 c_A \Delta H_1 + k_2 c_B \Delta H_2 + \\
 &\quad + k_3 c_A^2 \Delta H_3) + \frac{k_\omega A_R}{\rho C_p V_R}(T_K - T) \\
 T_K' &= \frac{1}{m_K C_{pK}}(Q_K + k_\omega A_R(T - T_K))
 \end{aligned}$$

The concentration of product B must not exceed some upper and lower limits. This gives the following constraints, which will be included in the objective function later on:

$$l \leq c_B \leq L.$$

MODEL BASED PREDICTIVE CONTROL

Let us denote by:

- $x_t = (c_A, c_B, T, T_K)$ the state at time t
- $y_t = (c_B)$ the controlled variable at time t
- $u_t = (F, Q_K)$ the control at time t (manipulated variable)

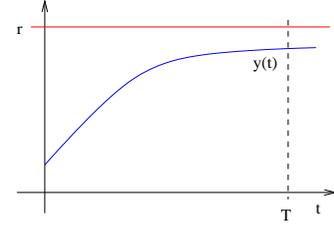


Figure 2: Control

r : reference level.

The objective of the predictive control is to find the future optimal control sequence $u_t^* = (u^*(i), i = 0, 1, \dots, N(T))$ over a finite horizon time T , which minimizes the *quadratic error* between the controlled variable y_t and its reference level r . In addition, u_t^* must assure that the trajectory of y_t is smooth, i.e. the range of values taken by the manipulated variables F and Q_K varies gradually and does not go from very high values to very low ones. The objective function is then:

$$J(x_t, y_t, u_t) = \int_0^T \|y_t - r\|^2 dt + \beta \int_0^T \|\dot{u}_t\|^2 dt,$$

where $\dot{u} = \frac{du}{dt}$, T is the prediction horizon and $\beta \geq 0$ is constant. The optimization is done over the set of feasible controls \mathcal{U} :

$$\min_{u \in \mathcal{U}} J(x_t, y_t, u_t)$$

The *receding horizon* mode of operation is used here: once the optimal control sequence $u_t^* = (u^*(i), i = 0, 1, \dots, N(T))$ is found, only $u_t^*(0)$ is applied to the system at time t . At time $t + 1$, another optimal sequence is found and again, only the first component is applied. The control u_t is kept constant during a so called *sampling time* h , when all the simulation parameters are kept at a fixed value. In the experiments already done, the sampling time was 20 seconds. At the next sampling time $t + h$, another optimal sequence is found and again, only the first component $u_t^*(0)$ is applied. Thus, the control is better adapted to the system's actual state and the noise is reduced. Notice that the time horizon is moved ahead (receding-horizon): $T(t + h) = T + h$.

In order to solve the problem it is necessary to parameterize the manipulated variable u_t , otherwise an infinite number of decision variables would appear in the problem. An usual approach is to discretize u along the control horizon N_u when the input remains constant over the sampling period h :

$$u_t = u(k), \quad kt \leq t < (k + 1)t,$$

$$u(k) = u(N_u - 1), \quad \forall k > N_u - 1.$$

Thus, the constrained optimization problem, subject to the continuous model equations and to the typical restrictions

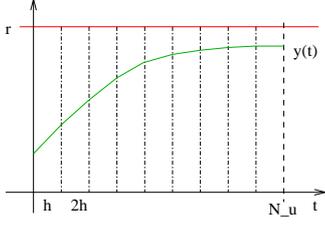


Figure 3: Discretized Control

applied to the manipulated and controlled variables, can be written as:

$$\min_{u(\frac{k}{h}), \dots, u(\frac{k+N_u-1}{h})} J = \int_{t_k}^{t_k+N\tau} \gamma [y(t) - r(t)]^2 dt + \sum_{i=0}^{N_u-1} \beta [\Delta u(k+i|t)]^2$$

$$y_{min} \leq y(k) \leq y_{max}$$

$$u_{min} \leq u(k) \leq u_{max}$$

$$\Delta u_{min} \leq \Delta u(k) \leq \Delta u_{max}$$

After penalizing the constraints on the controlled variable y , the objective function becomes:

$$\min_u J = \int_{t_k}^{t_k+N\tau} \gamma [y(t) - r(t)]^2 dt + \sum_{i=0}^{N_u-1} \beta [\Delta u(k+i|t)]^2 + \sum_{j=0}^{N_u-1} \alpha_1 [y_{max} - y(k+j|t)]^2 + \sum_{l=0}^{N_u-1} \alpha_2 [y_{min} - y(k+l|t)]^2$$

where β and γ are non-negative constants, and α_1, α_2 represent a penalty function. Remark that this optimization problem cannot be solved analytically. In fact, it is hard to calculate the gradient $\nabla_u J$, because y is non-linear in u . This is why we need to estimate it.

In order to treat the control constraints, we'll make a projection of the Lagrangian function J over the set of feasible controls. We'll truncate it in the sense that if $u(k) \geq u_{max}$, we'll set $u(k) = u_{max}$, and if $u(k) \leq u_{min}$, we'll set $u(k) = u_{min}$.

Previous approaches to the optimization problem were based on the SQP algorithm implemented in the NAG library, which uses finite differences to estimate the gradient ∇J . This approach has the disadvantage of losing precision and is quite slow because of the large number of func-

tion evaluations. Also, it may get trapped in local minima.

SYSTEM REQUIREMENTS

To see how fast the optimization must be, let us summarize the simulation by the algorithm below:

WHILE (simulation time) DO

WHEN (sampling) DO

1. Calculate the optimal control ($\min\{u \in \mathcal{U} : J(x_t, y_t, u_t)\}$);
2. Apply the control u_t^* ;
3. Measure C_a, T, T_K ; $x_t \leftarrow (C_a, T, T_K)$;
4. Take new value of C_b ; $y_{t+1} \leftarrow C_b$;
5. $t = t + h$; go to 1.

END WHEN

END WHILE

It is clear that the optimal control must be found in less than $h = 20$ seconds (the sampling time) if we want the control to be done online. The optimization must be fast and precise.

The Van der Vusse model presents important nonlinearities which makes the problem quite difficult. Another source of problems is the possible existence of local minima. Among others, this is a reason why we'll apply a method of global optimization. In the next section, we present this method, which is known for its efficiency in high dimensional problems. It could replace successfully the finite differences approximations in the Van der Vusse model and will perform as well in the presence of a larger number of variables.

SIMULTANEOUS PERTURBATION STOCHASTIC APPROXIMATION (SPSA)

SPSA is a descent method capable of finding global minima. Its main feature is the gradient approximation that requires only two measurements of the objective function, regardless of the dimension of the optimization problem. Recall that we want to find the optimal control u^* , with loss function $J(u)$:

$$u^* = \arg \min\{J(u) : u \in \mathcal{U}\}$$

Both Finite Differences Stochastic Approximation (FDSA) and SPSA use the same iterative process:

$$u_{n+1} = u_n - a_n \hat{g}_n(u_n),$$

where $u_n = ((u_n)_1, (u_n)_2, \dots, (u_n)_p)^T$ represents the n^{th} iterate, $\hat{g}_n(u_n)$ is the estimate of the gradient of the objective function $g(u) = \frac{\partial}{\partial u} J(u)$ evaluated at u_n , and $\{a_n\}$ is a positive number sequence converging to 0. If u_n is a p -dimension

vector, the i^{th} component of the symmetric finite difference gradient estimator is:

$$\text{FD: } (\hat{g}_n(u_n))_i = \frac{J(u_n + c_n e_i) - J(u_n - c_n e_i)}{2c_n},$$

$1 \leq i \leq p$, where e_i is the unit vector with a 1 in the i^{th} place, and c_n is a small positive number that decreases with n . With this method, $2p$ evaluations of J for each g_n are needed. Clearly, when p is large, this estimator loses efficiency.

The simultaneous perturbation estimator uses a p -dimensional random perturbation vector $\Delta_n = ((\Delta_n)_1, (\Delta_n)_2, \dots, (\Delta_n)_p)^T$ and the i^{th} component of the gradient estimator is, $1 \leq i \leq p$:

$$\text{SP: } (\hat{g}_n(u_n))_i = \frac{J(u_n + c_n \Delta_n) - J(u_n - c_n \Delta_n)}{2c_n (\Delta_n)_i}.$$

Remark that FD perturbs only one direction at the time, while the SP estimator disturbs all directions in the same time (the numerator is identical in all p components). The number of loss function measurements needed in the SPSA method for each g_n is always 2, independent of the dimension p . Thus, SPSA uses p times fewer function evaluations than FDSA, which makes it a lot more efficient.

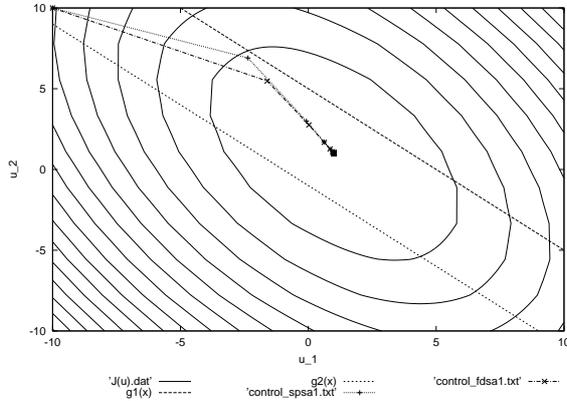


Figure 4: SPSA vs FDSA

Simple experiments with $p = 2$ showed that SPSA converges in the same number of iterations as FDSA. The latter follows approximately the steepest descent direction, behaving like the gradient method (see Figure 4). On the other hand, SPSA, with the random search direction, does not follow exactly the gradient path. In average though, it tracks it nearly because the gradient approximation is an almost unbiased estimator of the gradient, as shown in the following lemma found in [3].

Lemma 1 Denote by

$$b_n = \mathbb{E}[\hat{g}_n | \theta_n] - \nabla J(\theta_n)$$

the bias in the estimator \hat{g}_n . Assume that $\{(\Delta_n)_i\}$ are all mutually independent with zero-mean, bounded second moments, and $\mathbb{E}|(\Delta_n)_i|^{-1}$ uniformly bounded on \mathcal{U} . Then $b_n \rightarrow 0$ w.p. 1.

The detailed proof is in [3]. The main idea is to use conditioning on Δ_n to express $\mathbb{E}[(\hat{g}_n)_i]$ and then to use a second order Taylor expansion of $J(u_n + c_n \Delta_n)$ and $J(u_n - c_n \Delta_n)$. After algebraic manipulations implying the zero mean and the independence of $\{(\Delta_n)_i\}$, we get

$$\mathbb{E}[(\hat{g}_n)_i] = (g_n)_i + \mathcal{O}(c_n)$$

The result follows from the hypothesis that $c_n \rightarrow 0$.

Next we resume some of the the hypotheses under which u_t converges in probability to the set of global minima of $J(u)$. For details, see [5], [3] and [6]. The efficiency of the method depends on the shape of $J(u)$, the values of the parameters a_k and c_k and the distribution of the perturbation terms Δ_{ki} . First, the algorithm parameters must satisfy the following conditions:

- $a_t > 0$, $a_t \rightarrow 0$ when $t \rightarrow \infty$ and $\sum_{t=1}^{\infty} a_t = \infty$; a good choice would be $a_k = \frac{a}{k}$, $a > 0$;
- $c_t = c/t^\gamma$, where $c > 0$, $\gamma \in [\frac{1}{6}, \frac{1}{2}]$;
- $\sum_{t=1}^{\infty} (a_t/c_t)^2 < \infty$.
- Δ_{ki} must be mutually independent zero-mean random variables, symmetrically distributed about zero, with $|\Delta_{ki}| \leq \alpha_1 < \infty$ and $\mathbb{E}|\Delta_{ki}^{-2}| \leq \alpha_2 < \infty$ a.s., $\forall i, k$. A good choice for Δ_{ki} is Bernoulli ± 1 with probability 0.5. The uniform and normal distributions do not satisfy the finite moment conditions, so can not be used.

The loss function $J(u)$ must be thrice continuously differentiable and the individual elements of the third derivative must be bounded: $|J^{(3)}(u)| \leq \alpha_3 < \infty$. Also, $|J(u)| \rightarrow \infty$ as $u \rightarrow \infty$.

In addition, ∇J must be Lipschitz continuous, bounded and the ODE $\dot{u} = g(u)$ must have a unique solution for each initial condition.

Under these conditions and a few others (see [5]), u_k converges in probability to the set of global minima of $J(u)$.

RESULTS

The experiments done with the ECOSIM simulator showed that effectively, SPSA outperforms the SQP method using the finite differences gradient approximation. We realized that the quality of the control was very much depending on SPSA parameters (in particular of the initial values of a , c and the value of γ), so we performed a few pilot simulations in order to find the right values. The fastest simulation did twice better than FDSA, i.e. the same experiment was

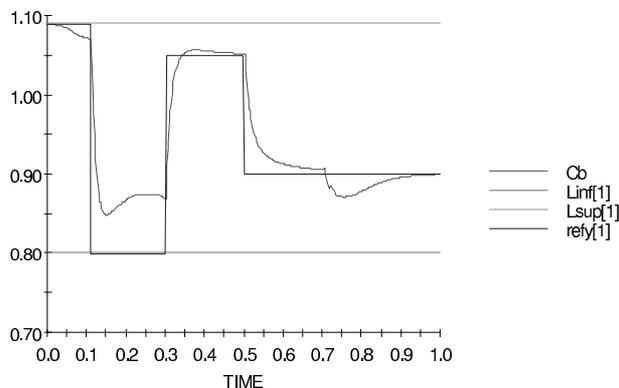


Figure 5: Concentration of Product B

two times faster when using SPSA. Like expected, the latter needed $p = 2$ times less cost function evaluations (in fact, the dimension of the problem here is 2: $u_t = (F, Q_k)$). The resulting control can be seen in Figure 5, where are plotted the reference level changing over time, the upper and lower bounds of the controlled variable c_B and the concentration of product B (y_t) itself. The small perturbation at time 0.7 is due to the change of temperature provoked on purpose. We see that the control is handling it well. Though, the controlled variable seems having some difficulties when near its upper and lower bounds, indicating that our way of treating the limit constraints may be incorrect. We used a dynamic stopping criteria of the form $\|\Delta u\| \leq \epsilon$, where $\epsilon = 0.1$ in most of the experiments, implying a small number of SPSA iterations. A higher precision was very costly, indicating that what we gained in speed was unfortunately lost in precision. However, we believe that a better choice of parameters can make the method more robust and this makes part of our ongoing work.

CONCLUSION

In this research we present a methodology for global optimization adding noise to the observations of the gradients in order to achieve better performance of the model. Theory showed that the addition of random noise can make the control variables attain near optimality much faster than deterministic methods as confirmed by some of our experiments. In addition, this method can provably overcome the curse of dimensionality and thus be used in larger problems. However, finding the suited parameters proved to be challenging, showing that a particular attention should be given to this point.

References

- [1] D. P. Bertsekas. *Nonlinear Programming*. Mass. USA, athena scientific edition, 1999.
- [2] H. Chen and F. Allgower. A quasi-infinite horizon nonlinear model predictive control scheme with guaranteed stability. In *Automatica*, 1998.
- [3] M. C. Fu and S. D. Hill. Optimization of discrete event systems via simultaneous perturbation stochastic approximation. *IEEE Transactions*, 29:233–243, 1997.
- [4] A. K. H. Chen and F. Allgower. Nonlinear predictive control of a benchmark cstr. In *Proceedings of 3rd European Control Conference*, pages 3247–3252, 1995.
- [5] J. L. Maryak and D. C. Chin. Global random optimization by simultaneous perturbation stochastic approximation. In *Proceedings of the American Control Conference*, pages 756–762, 2001.
- [6] J. C. Spall. An overview of the simultaneous perturbation method for efficient optimization. *John Hopkins APL Technical Digest*, 19(4):482–492, 1998.
- [7] C. D. P. W. Colmenares, S. Cristea and al. Mld systems: Modeling and control experience with a pilot process. In *Proceedings of IMECE 2001*, New York, USA, 2001.

ACKNOWLEDGEMENTS

The work of the first two authors was sponsored in part by NSERC and FCAR Grants of the Government of Canada and Quebec.