Proceedings of the
44th IEEE Conference on Decision and Control, and
the European Control Conference 2005
Seville, Spain, December 12-15, 2005

Implementation of Model Predictive Control for Glucose Regulation on a General Purpose Microprocessor

Leonidas G. Bleris and Mayuresh V. Kothare

Abstract— **In order to effectively control nonlinear and multivariable models, and to incorporate constraints on system states, inputs and outputs (bounds, rate of change), a suitable (sometimes necessary) controller is Model Predictive Control (MPC), also known as receding horizon control and moving horizon control. For processes with slow dynamics and low sampling rates, MPC is typically implemented on a dedicated workstation. In this work we present the implementation of MPC on a general-purpose processor, providing a low-cost, low-power consumption and small in size implementation. In order to test the performance of the embedded MPC we examine a glucose regulation problem. Additionally, profiling results of the performance of the processor are provided.**

I. Introduction

Model Predictive Control originated in the chemical process industries mainly due to the need for a controller capable of handling constrained and multivariable nonlinear processes. Because of the computational requirements of the optimizations associated with MPC, it has primarily been applied to processes with slow dynamics. Existing implementations of MPC typically perform numerical calculations using workstations (currently) in 64-bit Floating Point (FP) arithmetic, which is too expensive, power demanding and large in size. For systems with fast dynamics, where the size and the application precludes the use of a dedicated workstation, alternative solutions have to be investigated. Thus, there is an increasing need for reducing MPC on a chip. Some approaches to this direction have been reported in literature [1], [2], [3].

We have proposed in [4], [5] an efficient implementation of MPC by reducing the precision of the microprocessor to the minimum, while maintaining optimal control performance. By reducing the precision we increase the optimization speed, and at the same time we minimize the power consumption and the overall chip area (thereby the cost). This reduced precision application-specific processor can achieve sampling speeds as low as 0.032 seconds for relatively large problems [5]. To quantify the advantage of reducing the precision, estimations for both 64-bit FP and 16-bit LNS circuits show that for an arithmetic unit that computes addition, subtraction, multiplication and division, the size required is about 17 times smaller than a 64-bit floating point processor. The advantage of formulating the on-line MPC problem in this fashion is that it directly extends to other objective functions, as opposed to the MP approach which applies to QPs [2], [3].

In this paper we provide the results of the implementation of MPC on a general-purpose processor that provides a low-power, small in size and relatively inexpensive solution. We examine the performance of MPC using processor in the loop co-simulations; the optimizations associated with MPC are carried out on the general processor and control is applied to a model that runs on a host PC. A case study of glucose regulation for diabetic patients is examined. Additionally the performance of MPC running on this board is profiled and the results are presented.

The paper is organized as follows: Section II contains the main theoretical aspects of MPC and information on the optimization associated with the MPC operations. In Section III we provide the details of the hardware implementation of MPC, information about the profiling and the co-simulations. The glucose regulation case study is given in Section IV. We conclude the paper with remarks on our research results and an analysis of issues that are currently under investigation.

II. Theoretical background

Controllers belonging to the MPC family are generally characterized by the following steps [6] (The schematic of the overall system is given in Figure 1):

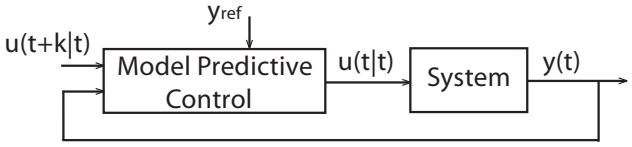

Fig. 1. Block diagram of MPC.

1. Initially the future outputs are calculated at each sample interval over a predetermined horizon N, the prediction horizon, using the process model. These outputs $y(t+k|t)$ for $k = 1...N$ depend up to the time t on the past inputs and on the future signals $u(t+k|t)$, $k = 0...N-1$ which are those to be sent to the system.

2. The next step is to calculate the set of future control moves by optimizing a determined criterion in order to keep the process as close as possible to a predefined reference trajectory. This criterion is usually a quadratic function of the difference between the predicted output signal and

Partial financial support for this research from the US National Science Foundation grant CTS-0134102 and the Pittsburgh Digital Greenhouse is gratefully acknowledged.

Mayuresh V. Kothare is with the Faculty of the Department of Chemical Engineering, Lehigh University, Bethlehem, PA, 18015 mayuresh.kothare@lehigh.edu

Leonidas G. Bleris is with the Department of Electrical and Computer Engineering, Lehigh University, Bethlehem, PA, 18015 bleris@lehigh.edu

0-7803-9568-9/05/$20.00 ©2005 IEEE

the reference trajectory. In some cases the control moves $u(t+k|t)$ are included in the objective function in order to minimize the control effort:

$$J_P(k) = \sum_{k=0}^{P} \{[y(t+k|t) - y_{ref}]^2 + Ru(t+k|t)^2\} \quad (1)$$

$$|u(t+k|t)| \leq b \quad , \quad k \geq 0 \quad (2)$$

where $y(t+k|t)$ are the predicted outputs, y_{ref} is the desired set reference output, $u(t+k|t)$ the control sequence and R is the weighting on the control moves, a design parameter. This system is subject to input constraints given by the vector b.

3. Finally, the first control move $u(t|t)$ is sent to the system while the rest are rejected. This is because at the next sampling instant the output $y(t+1)$ is measured by the system and the procedure is repeated with the new values so that we get an updated control sequence.

A. Optimization

The optimization is a fundamental part of MPC since it results in optimal control inputs for the process. The computational effort required in MPC derives almost entirely from the optimization algorithm. Logically the choice of the optimization technique is decisive to the performance of the controller [7]. Typically, quadratic functions are used as cost functions in MPC. Equation 1 can be simplified and rewritten as:

$$\underset{x}{minimize}\, q(x) = \frac{1}{2}x^T Gx + g^T x \quad (3)$$

$$subject \quad to \quad a_i^T x = b_i, \quad i \in E \quad (4)$$

$$a_i^T x > b_i, \quad i \in I \quad (5)$$

where x is the vector of variables (corresponds to the actuation u), and the constraints have been generalized and divided into two sets, E or equality constraints and I or inequality constraints. Any constant term (not depending on x) appearing after the simplification of equation 1 can be ignored since it will not affect the solution of the minimization problem.

We use Newton's algorithm to solve problem by combining the constraints into the cost function, using barrier functions for the inequality constraints, and penalty functions for equality constraints, defined as:

$$d_i(x) = \mu_i (a_i^T x - b_i)^2 \,, \quad i \in E \quad (6)$$

$$d_i(x) = \mu_i \log(a_i^T x - b_i) \,, \quad i \in I \quad (7)$$

resulting in the *unconstrained* non-linear problem:

$$\underset{x}{minimize} \quad f(x) = \frac{1}{2}x^T Gx + g^T x + \sum_i d_i(x) \quad (8)$$

The problem of equation 8 can be solved numerically approximating $f(x)$ by a quadratic function around x, obtaining the gradient $\nabla f(x)$ and Hessian $H(x)$, and iterating:

$$x^{(k+1)} = x^{(k)} - H^{-1}(x^{(k)}) \cdot \nabla f(x^{(k)}) \quad (9)$$

The complexity of this algorithm is dominated by the computation of H^{-1}, which requires $O(n^3)$ operations, where n is the number of variables to optimize.

III. Hardware implementation of MPC

For the real-time implementation of model predictive control we use the high-performance single board computer phyCORE-MPC555, illustrated in Figure 2. This board packs the power of Motorola's embedded 32-bit MPC555 microcontroller within a miniature footprint. The MPC555 is a high-speed 32-bit Central Processing Unit that contains a 64-bit floating point unit designed to accelerate the advanced algorithms necessary to support complex applications. All signals and ports of the MPC555 extend to two Molex high density (0.635 mm pitch) 160 pin header connectors. These high density pins allow it to be plugged like a "big chip" into user target hardware.

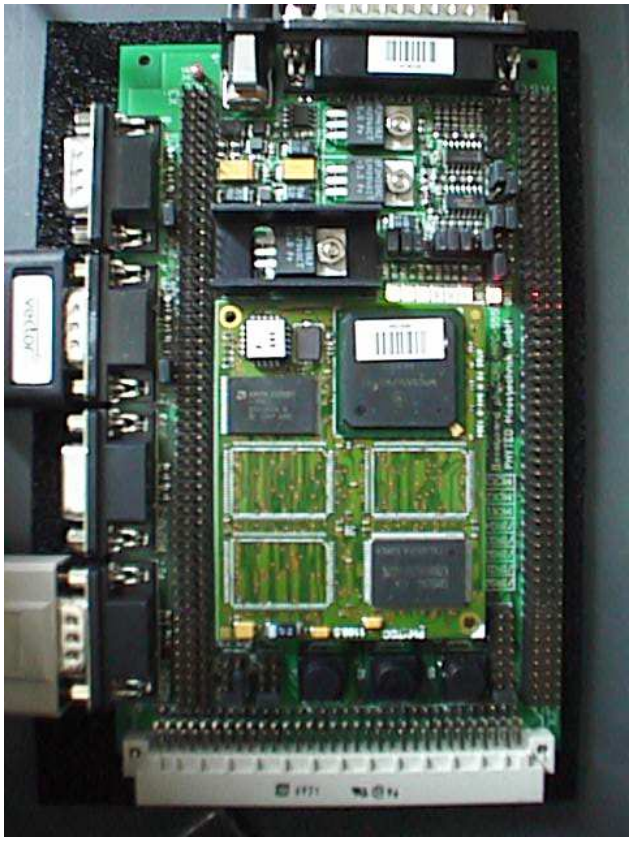

Fig. 2. PhyCORE-MPC555 board and MPC555 processor (2cm×2cm).

In order to implement the optimization required by MPC on this target we use a combination of software tools: CodeWarrior Integrated Development Environment (IDE) [8], MATLAB, Real-Time Workshop and SIMULINK. By combining state-of-the-art MATLAB/SIMULINK environment with the robust CodeWarrior Development Studio, we were able to compile a C++ implementation of optimization and download the operations on the MPC555 processor. The size of the code downloaded was on average 30KBytes. The phyCORE-MPC555 can be populated with a maximum of 4 MB flow-through synchronous external BURST-SRAM for DATA storage and a maximum of 4 MB Flash-ROM memory at 0 Wait-States.

A. Code profiling

A profiler can analyze the amount of time a program spends performing various tasks of the optimization. A profiler can help detect bottlenecks (time consuming routines that data passes through) and routines that are inordinately slow. Clearly a good profile of the runtime performance of code requires more information than a raw count. Most of the profilers perform statistical sampling of the runtime environment; these profilers are called passive or sampling profilers. A passive profiler divides the program being profiled into evenly-sized "buckets" in memory. It then samples the processor's program counter at regular intervals to determine which bucket the counter is in. The main advantage of a passive profiler is that it requires no modification to the program under observation. Also, passive profilers distribute the overhead that they incur evenly over time, allowing the post-processing steps to ignore it. On the other hand, they cannot sample too frequently or the sampling interrupt will overwhelm the program being sampled.

B. Co-simulation

In order to test the performance of MPC on the embedded target we use Processor-In-the-Loop (PIL) co-simulations. The board is connected to a host workstation and in PIL mode, a plant model runs in non-real-time on the host workstation in Simulink. Meanwhile, generated code running on the MPC555 exchanges signals via RS232 serial communication with the Simulink simulation running on the workstation. At each sample interval, Simulink performs model updates and sends output signal data via RS232 to the MPC555. During PIL testing, one can include monitoring probes in the Simulink plant model. These probes may assist the code validation by displaying signal data. The signal data retrieved from the microcontroller is used for the next simulation step of the plant model.

IV. Active Drug Delivery Devices

The development of novel medical sensors and drug delivery methods have introduced multiple benefits to medical devices. These include safety, efficacy, robustness, and patient compliance. For example, advances over the past few years in microfabrication technologies have allowed researchers to create microneedles [9], [10] that painlessly cross the uppermost layer of the skin to effectively, efficiently, and painlessly deliver drugs to depths as shallow as $1mm$. Their sizes range from one millimeter to one micron in diameter, being too minute to stimulate nerve endings and cause pain. In parallel several other drug delivery approaches are under research and development [11], including non-invasive needle-free injection devices. The key to needle-free injection is to release the drug in liquid form at the proper (extremely high) velocity so that it diffuses through the skin. These drug delivery devices have the potential to maximize patient compliance, especially for diseases that require frequent drug doses.

Looking at the future, the end target is to develop small and compact drug delivery devices, that require minimal power and can maximize the therapeutic results of a drug. Noteworthy [12], for the former, integrated batteries can be used in order to provide enough power to an implant for five to eight years of life, without causing side effects. The latter can be accomplished by minimizing selected cost functions in a control procedure (the actual cost function of the optimization will depend upon the particular application and control objectives). Classic control algorithms (i.e. ON/OFF, PID, gain scheduling etc) can be used for drug-delivering medical devices. These algorithms can be implemented easily on chip, are robust, and have relatively low energy consumption. Nevertheless, the human organism is one of the most complex dynamical systems. Because of the inherent existence of nonlinearities, constraints, patient variability and delays the need for advanced control schemes becomes essential. The increased understanding of physiological, pharmacokinetic and pharmacodynamic models gives us the ability to apply advanced model based control schemes. In order to effectively control nonlinear and multivariable models, and to incorporate constraints in the manipulated variables (bounds, rate of change), a suitable controller is model predictive control.

A. Blood Glucose Control in Diabetic Patients

Diabetes mellitus is a chronic health condition where the human body is unable to produce insulin and properly breakdown sugar (glucose) in the blood. The insufficient insulin production or lack of responsiveness to insulin, results to hyperglycemia (high blood glucose levels, 120mg/dL). There are two primary types of diabetes mellitus, type I (insulin-dependent or juvenile-onset), which may be caused by an autoimmune response, and type II (non-insulin-dependent or adult-onset). Symptoms include hunger, thirst, excessive urination, dehydration and weight loss. Complications can include heart disease, stroke, neuropathy, poor circulation leading to loss of limbs, hearing impairment, vision problems and death.

The treatment of diabetes requires regular insulin injections, proper nutrition and exercise in order to maintain normoglycemia, defined as blood glucose $70-100$mg/dL. Insulin and glucagon are the hormones responsible for glucose regulation. Both insulin and glucagon are secreted from the pancreas, and thus are referred to as pancreatic endocrine hormones. Most of the long-term complications associated with diabetes result from sustained hyperglycemia, but hypoglycemia can result in very acute symptoms such as coma and death.

A significant effort has been devoted to the development of closed-loop controllers for blood glucose control. In the following sections we provide the results of co-simulations, where we apply MPC to maintain normoglycemia. For approaches to applying control to diabetic subjects the reader is referred to [13], [14], [15].

B. Minimal Glucose Model

Minimal models of glucose and insulin plasma levels have been developed [16], [17], [18] for humans using frequently-

sampled intravenous glucose tolerance (FSIGT) tests. During a FSIGT test, a single intravenous injection of glucose is given to a fasting subject and blood samples are collected at regular timed intervals. The blood samples are then analyzed for glucose and insulin concentration. Figure 3 shows a typical response from a normal subject.

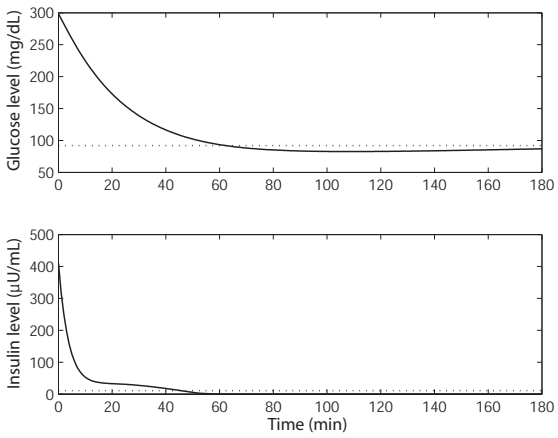

Fig. 3. Typical glucose and insulin response from a normal subject.

As illustrated in Figure 3, the glucose level in plasma starts at a peak due to the injection, drops to a minimum which is below the basal glucose level, and then gradually returns to the basal level (dashed line). The insulin level in plasma rapidly rises to a peak immediately after the injection, drops to a lower level which is still above the basal insulin level, rises again to a lesser peak, and then gradually drops to it's basal level. Depending on the state of the subject, there can be wide variations from this response that can determine the condition of the patient [17].

The glucose minimal model involves two physiologic compartments: an interstitial tissue compartment and a plasma compartment. Insulin leaves or enters the interstitial tissue compartment at a rate proportional to the difference between the plasma insulin level, I(t), and the basal level I_b (equation 11). If the plasma insulin level falls below the basal level, insulin leaves the interstitial tissue compartment, and if the plasma insulin level rises above the basal level, insulin enters the interstitial tissue compartment. Insulin also disappears from the interstitial tissue compartment via a second pathway at a rate proportional to the amount of insulin in the interstitial tissue compartment. Glucose leaves or enters the plasma compartment at a rate proportional to the difference between the plasma glucose level, G(t), and the basal level G_b; if the plasma glucose level falls below the basal level, glucose enters the plasma compartment, and if the glucose level rises above the basal level, glucose leaves the plasma compartment. Glucose also disappears from the plasma compartment via a second pathway at a rate proportional to the amount of insulin in the interstitial tissue. The differential equations corresponding to the above analysis are:

$$\frac{dG(t)}{dt} = k_1(G_b - G(t)) - X(t)G(t) \tag{10}$$

$$\frac{dX(t)}{dt} = k_2(I(t) - I_b) - k_3X(t) \tag{11}$$

where t is time, G(t) is the plasma glucose concentration at time t, I(t) is the plasma insulin concentration at time t, and X(t) is the interstitial insulin at time t, with G(t_0)=G_0 and X(t_0)=0. G_b is the basal plasma glucose concentration and I_b is the basal plasma insulin concentration. The insulin sensitivity is defined as $S_I = k_2/k_3$ and the glucose effectiveness is defined as $S_G = k_1$. Basal plasma concentrations of glucose and insulin are typically measured either before, or 180 minutes after, administration of glucose. There are four unknown parameters in this model: k_1, k_2, k_3, and G_0 that depend on the particular subject and can be estimated experimentally. We use the following parameters adopted from [18], [19]:

G_b	92 (mg/dL)
I_b	11 (μU/mL)
G_0	92 (mg/dL)
k_1	0.038291 (1/min)
k_2	$3.5364 * 10^{-6}$ (1/min)
k_3	0.0015002 (1/min)

TABLE I

PARAMETERS USED FOR THE MINIMAL MODEL

C. Multi-model MPC

In the minimal glucose model, glucose is utilized at the constant rate k_1, having the negative feedback effects of the interstitial insulin, represented by the term -X(t)G(t). An additional amount of plasma insulin will cause the amount of interstitial insulin X(t) to change, which will cause the rate of glucose utilization to change. In order to capture efficiently this effect we use a multi-model MPC approach. That is, we capture the glucose levels responses to pulse inputs of insulin for different initial conditions for the interstitial insulin. Under the assumption that the interstitial insulin is estimatable [20], [21] at each sampling interval, the predictive controller utilizes the proper model that best describes the dynamics of glucose-insulin in the human body.

As illustrated in Figure 4 we obtain the glucose level responses to a pulse input of 100(μU/ml) of insulin, for different initial conditions of interstitial insulin (from top starting at 0 with intervals of 0.005). For all these cases the glucose level returns to the basal level, slowly after reaching the minimum at 30min.

D. Co-simulation results

In order to test the performance of this multi-model MPC in real time we proceed with the processor-in-the-loop co-simulations. The setup is illustrated in Figure 5, where we have the MPC chip in closed loop with the monitored patient (the minimal model on the host workstation). The sampling time for the system is set at one second.

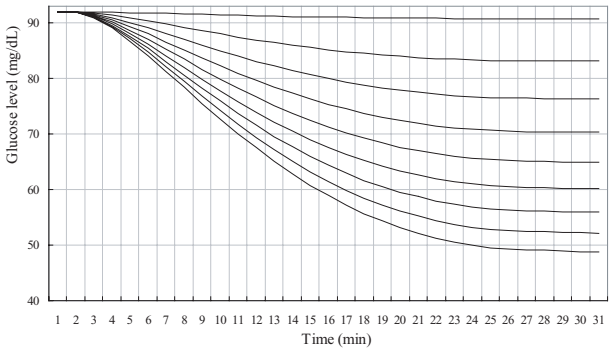

Fig. 4. Glucose level responses to pulse input of insulin, for different initial conditions of interstitial insulin.

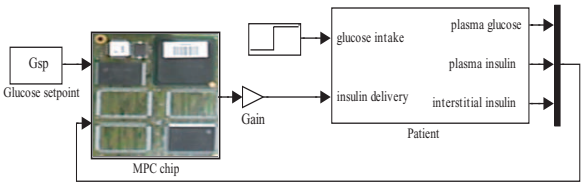

Fig. 5. Block diagram of the MPC chip in closed loop with the monitored patient.

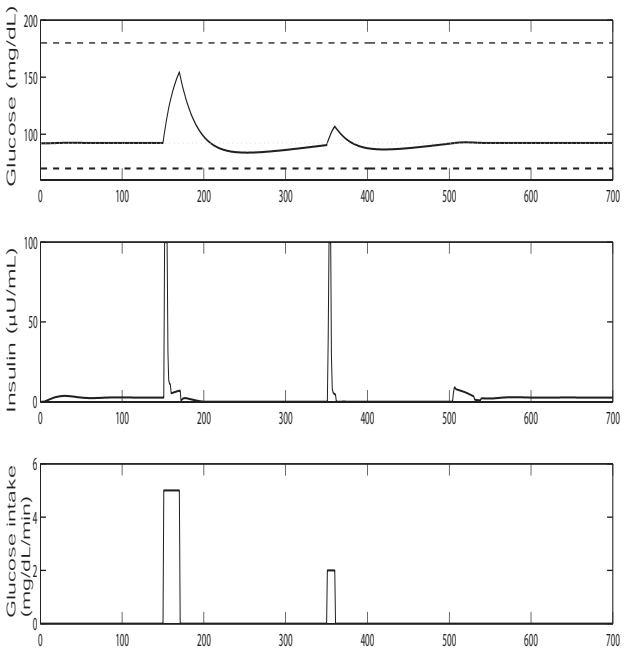

Fig. 6. Simulation results for control horizon of 5 and prediction horizon of 31.

The simulation results for control horizon of 5, a prediction horizon of 31 and a fixed number of optimizations at 40, are given in Figure 6. The first from the top subplot represents the concentration of glucose in the blood, the second is the manipulated variable (insulin injection), and the third subplot is the glucose intake (due to two small meals). Due to the insulin injection we manage to keep the glucose concentration between the desired limits, avoiding hyperglycemia that would result in the absence of control and hypoglycemia that would result for overdosing insulin.

The profiling results, given in Figure 7, illustrate that the time required for the computation of the optimal insulin dosage from the predictive controller remains under one second for all examined control horizon cases, and prediction horizons of 21 and 31. Therefore, there is sufficient time for increasing the number of optimizations and adding fail-safe approaches that will guarantee an optimal, safe and consistent behavior by this embedded model predictive control implementation. In Figure 8 we provide the profiling results for variable number of iterations in the optimization, control horizon 7, and prediction horizons of 21 and 31. As expected the computational time grows linearly.

V. Concluding Remarks

A real-time implementation of model predictive control using a general purpose processor has been provided in this paper. Profiling results illustrate that the optimizations require less than one second for most cases. An implementation of real-time MPC on-a-chip for diabetes control, would additionally require the incorporation of a glucose estimator, some changes in the optimization code (so as to successfully deal with disturbances and uncertainty), the addition of fail-safe and scheduling routines, and finally the integration of the sensors and actuators. Therefore we should expect an increase in the computational time. Nevertheless, for the particular case of glucose regulation the sampling time is not a restriction, since most of the glucose sensors today operate with a sampling rate of 5 minutes.

Novel microfabrication technologies have allowed researchers to create minute, hollow needles that painlessly cross the uppermost layer of the skin to deliver drugs to the dermal layer. Some products under development can effectively, efficiently, and painlessly deliver drugs by bolus injection or infusion to depths as shallow as 1 mm. Coupled with additional advances in microelectronics some companies are developing disposable pumps for use with chronically administered drugs. The end target is to integrate a dedicated microprocessor for the application of real-time feedback control. There are numerous potential applications in this area: control of physiological processes, glucose control, muscle control, respiration control, drug infusion control, cardiac pacemakers and defibrillators, heart rate control, blood flow and pressure control, HIV control, and neurological implants. Drug-delivering closed-loop controllers can improve the effectiveness and safety of drugs, but also the quality of life of patients. A sensor can monitor the effect of the drug and through a control algorithm adjust the dosage to the optimum levels for effectiveness, reducing thereby side effects and potentially minimizing the costs. On the other hand, several small-scale industrial and consumer

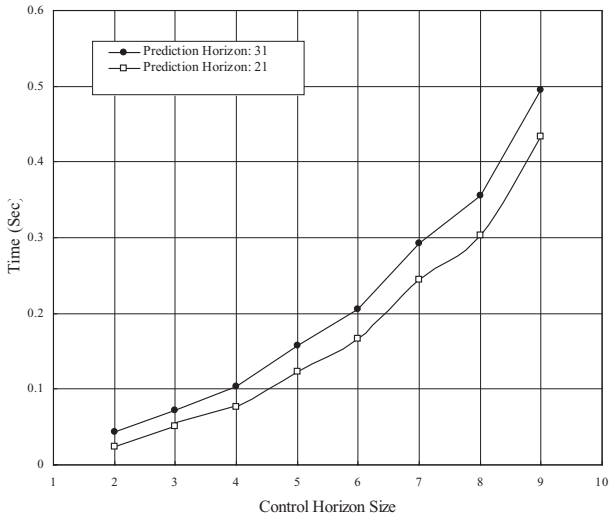

Fig. 7. Profiling results for variable control horizons, and prediction horizons of 21 and 31.

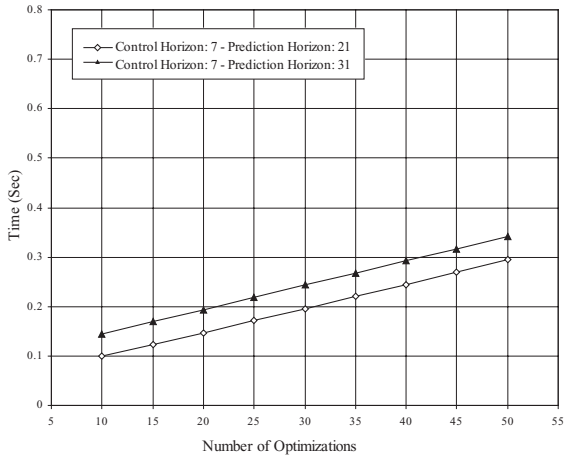

Fig. 8. Profiling results for variable number of iterations in the optimization, control horizon 7, and prediction horizons of 21 and 31.

electronics application areas that are characterized by fast sampling or low-cost properties will benefit from a real-time implementation of MPC. These areas include: MEMS, automotive/aerospace control, power electronics, microchemical systems. The advantages of MPC such as the ability to handle constraints, the applicability to nonlinear processes and to multivariable problems, constitute this control method a necessary choice for many of these control problems.

We are currently investigating [22], [23] the implementation of MPC on an 16-bit fixed-point microprocessor, assisted by a specially customized Field Programmable Gate Array (FPGA) that will accelerate the matrix operations associated with MPC. Thus, this is a mixed software-hardware embedded controller for real-time control applications. We intend to provide a practical co-design framework for the implementation of real-time MPC for large scale systems.

REFERENCES

[1] George Hassapis. Implementation of model predictive control using real-time multiprocessing computing. *Microprocessors and Microsystems*, 27:327–340, 2003.

[2] A. Bemporad, M. Morari, V. Dua, and E. N. Pistikopoulos. The explicit linear quadratic regulator for constrained systems. *Automatica*, 38(1):3–20, January 2002.

[3] E. N. Pistikopoulos. On-line optimization via off-line optimization! - a guided tour to parametric programming and control. Invited plenary lecture at the 7^{th} IFAC Symposium on Dynamics and Control of Process Systems (DYCOPS-7), Boston, MA, July 2004.

[4] L. G. Bleris, M. V. Kothare, J. G. Garcia, and M. G. Arnold. Embedded model predictive control for system-on-a-chip applications. In *Proceedings of the 7^{th} IFAC Symposium on Dynamics and Control of Process Systems (DYCOPS-7)*, Boston, MA, July 2004.

[5] L. G. Bleris, M. V. Kothare, J. G. Garcia, and M. G. Arnold. Towards embedded model predictive control for system-on-a-chip applications. *To appear: Journal of Process Control*, 2005.

[6] E. F. Camacho and C. Bordons. *Model Predictive Control in the Process Industry*. Springer, New York, 1995.

[7] R. Fletcher. *Practical Methods of Optimization*. John Wiley & Sons, 1987.

[8] Metrowerks Corporation, Austin, TX. *CodeWarrior Development Studio, MPC5xx Edition*, 2004.

[9] R. Langer. Transdermal drug delivery: past progress, current status, and future prospects. *Advanced Drug Delivery Reviews*, 56:557–558, March 2004.

[10] M. R. Prausnitz. Microneedles for transdermal drug delivery. *Advanced Drug Delivery Reviews*, 56:581–587, March 2004.

[11] S. L. Taoa and T. A. Desai. Microfabricated drug delivery systems: from particles to pores. *Advanced Drug Delivery Reviews*, 55:315–328, March 2003.

[12] Orhan Soykan. Power sources for implantable medical devices. *Device Technology & Applications*, pages 76–79, 2002.

[13] R. S. Parker, F. J. Doyle, and N. A. Peppas. A model-based algorithm for blood glucose control in type I diabetic patients. *IEEE Transactions on Biomedical Engineering*, 46(2):148157, 1999.

[14] E. Ruiz-Velazquez, R. Femat, and D. U. Campos-Delgado. Blood Glucose Control for type I diabetes mellitus: A robust tracking Hinf problem. *Control Engineering Practice*, 12:1179–1195, 2004.

[15] R. S. Parker and F.J. Doyle III. Control-relevant modeling in drug delivery. *Advances in Drug Delivery Reviews*, 48(2):211–248, 2001.

[16] R. N. Bergman, Y. Z. Ider, C. R. Bowden, and C. Cobelli. Quantitative estimation of insulin sensitivity. *American Journal of Physiology*, 236:E667–E677, 1979.

[17] S. E. Kahn G. M. Steil, A. Volund and R. N. Bergman. Reduced sample number for calculation of insulin sensitivity and glucose effectiveness from the minimal model. *Diabetes*, 42:250–256, 1993.

[18] G. Pacini and R. N. Bergman. A computer program to calculate insulin sensitivity and pancreatic responsivity from the frequently sampled intravenous glucose tolerance test. *Computer Methods and Programs in Biomedicine*, 23:113–122, 1986.

[19] N. V. Riel. Minimal models for glucose and insulin kinetics. In *Technique Report*, Eindhoven University of Technology, February 2004.

[20] S. M. Lynch and B. W. Bequette. Model predictive control of blood glucose in type 1 diabetics using subcutaneous glucose measurements. In *Proceedings of the 2002 American Control Conference*, pages 4039–4043, Anchorage, AK, June 2002.

[21] R. Hovorka, B. Canonico, L. J. Chassin, U. Haueter, M. M. Benedetti, M. O. Federici, T. R. Pieber, H. C. Schaller, L. Schaupp, T. Vering, and M. E. Wilinska. Nonlinear model predictive control of glucose concentration in subjects with type 1 diabetes. *Physiological Measurement*, pages 905–920, 2004.

[22] L. G. Bleris, P. Vouzis, M. G. Arnold, and M. V. Kothare. Submitted: A Co-Processor FPGA Platform for the Implementation of Real-Time Model Predictive Control. In *2006 American Control Conference*, Minneapolis, MI, July 2006.

[23] P. Vouzis, L. G. Bleris, M. V. Kothare, and M. G. Arnold. Towards a Co-design Implementation of a System for Model Predictive Control. In *2005 AIChE Annual Meeting*, Cincinnati, OH, November 2005.