NONLINEAR PREDICTIVE CONTROL IN THE LHC ACCELERATOR

C. de Prada*, E. Blanco, S. Cristea*, J. Casas****

***Dpt. of Systems Eng. and Automatic Control, University of Valladolid, Spain*
fax: +34 983 42 3161, prada@autom.uva.es, smaranda@autom.uva.es
***LHC Division, CERN, CH-1211, Geneva (Switzerland)*
fax: +41 22 767 4620, Enrique.Blanco@cern.ch, Juan.Casas@cern.ch

Abstract: This paper describes the application of a nonlinear model-based control strategy in a real challenging process. A predictive controller based on a nonlinear model derived from physical relationships, mainly heat and mass balances, has been developed and commissioned in the Inner Triplet Heat Exchanger Unit (IT-HXTU) prototype of the LHC particle accelerator being built at CERN, operating at a temperature of about 1.9 K. The development includes a state estimator with a receding horizon estimation procedure to improve the regulator predictions. *Copyright © 2003 IFAC*

Keywords: Nonlinear Predictive control, Nonlinear estimators, Cryogenic

1. INTRODUCTION

MPC strategies have become the preferred control technique for many process control problems, most of the industrial MPC controllers using an internal linear dynamic model. Nevertheless, many common processes exhibit nonlinear behavior and they may be required to operate over a wide range of conditions, therefore, these controllers are often tuned in a conservative way, which can result in serious degradation of controller performance. An alternative is to use a non-linear internal model. Many controllers, using different types of internal models, have been proposed in the literature, but the number of non-linear MPC is still low in industry.

In this paper we present an implementation of a non-linear MPC to a challenging process, which requires to operate under strong requirements. It corresponds to the new particle accelerator, the LHC, which is under construction at CERN, Geneva. In its final form it will expand in a circumference of about 27 Km in France and Switzerland. In order to test the proposed design some prototypes such as the String1 (Casas et al., 1998) and the IT-HXTU (Byrns et al., 1998) were built (Fig.1).

The aim of the LHC is to accelerate particles at speeds close to the one of the light in order to study the results of its collisions. For this purpose, the particles are driven within the LHC accelerator using very strong magnetic fields, which requires high electrical currents of about 12 kA for its magnets. A practical operation of the magnets requires operating with no electrical resistance in the coils, superconductivity condition, that can be maintained only at extremely low temperatures of around 1.9K. The main aim of the control system presented in this paper is to maintain this temperature in a narrow range, 50 mK, in spite of the unknown disturbances acting on the process. This is required in order to avoid a "quench" that will stop the operation.

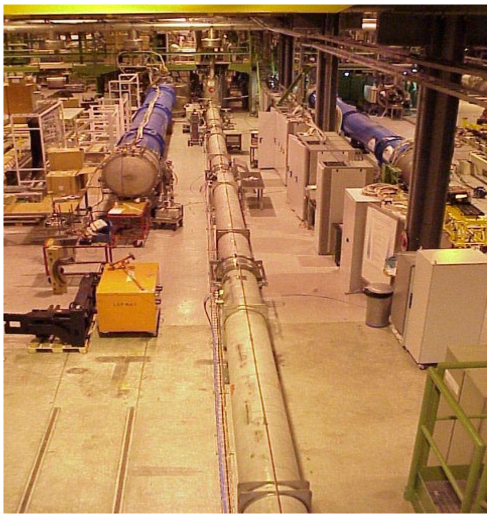

Fig. 1 Inner Triplet Prototype (length: 30 meters)

Several linear control strategies has been tested at String1, including PID and linear MPC (Blanco, 1999; Cristea, 1998) but all of them suffer from the above mentioned problem: their performance is degraded when, due to different heat load charges, the unit must operate in different working conditions. This was the main reason to implement NMPC. This paper presents the non-linear approach, as well as state and disturbance estimations that were not taken into account in previous versions, and it is organized as follows: After the introduction, section 2 describes the String1 and its cryogenic system. Section 3 is devoted to the process model and section 4 to the non-linear controller including the state estimator. Experimental results are given in section 5. The paper ends with some brief conclusions.

2. PROCESS DESCRIPTION

The LHC 1.8 K Cooling Loop represents a structure of four magnets (four quadropoles in the IT-HXTU and one quadrupole and three dipoles in the String1 prototypes) mounted at a slope of 1.4% to match the steepest inclination in the real accelerator tunnel. The superconducting magnets operate below 1.9 K in a bath of pressurized helium.

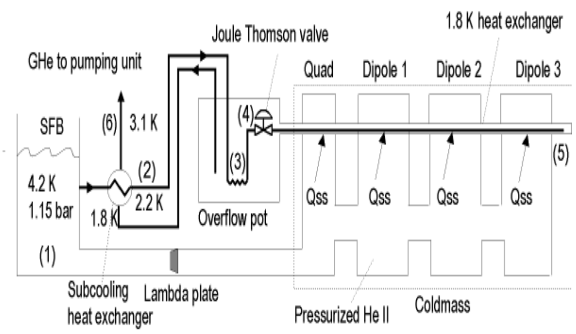

Fig. 2. Process and Schematic Diagram (String1)

Referring to Fig.2, the heat deposited on the bath is extracted by gradual vaporization of saturated superfluid helium flowing along the wetted length of a heat exchanger (HX) tube threading the string of magnets. The tube is only partially wetted, being the wetted length a main control variable. The liquid helium used for cooling at the 1.8K level is taken from the main reservoir (SFB) at 4.2 K and 1.15 bar (1). The helium is subcooled in the subcooling-heat exchanger (2) to 2.2 K, and then it is sent through the heat exchanger in the overflow pot (3). The subcooled liquid is then expanded to saturation at 17 mbar and 1.8 K in the Joule-Thomson valve (4), where a vapor fraction is created as well. The helium is led to the end of String (5). Here, it is let out in the HX, and flows back towards the overflow pot that, in normal operation, is empty. The helium vapor at 17mbar is taken out from the overflow pot (3) and through the subcooling-heat exchanger (6), thus providing the subcooling for the incoming pressurized liquid at 4.2 K.

The regulation goal is keeping the temperature of the superconducting magnets as constant as possible within strict operating constraints imposed by the maximum temperature at which the magnets can operate, the cooling capacity of the cryogenic system, the heat loads, and at last, the accuracy of the instrumentation. A small margin of a few mK is allowed before the superconductivity of the magnet coils is lost. If this happens, a potentially dangerous situation (quench) is created because of the heat released in the new conditions and the sudden helium vaporization it implies.

The Joule-Thomson valve opening is the manipulated variable, and the temperature sensors located in the cold mass (two per magnet) provide the controlled variables, the warmest temperature is taken as controlled variable at every time step. Disturbances are of two different types: general heat loads and variations in the flow through the Joule-Thomson valve. Heat loads are produced mainly by heat inleaks from the higher temperature levels, current magnet ramping and particle beam losses (simulated by electrical heaters). The set point is the saturation temperature plus a certain ΔT, typically 0.03 K.

This process has shown difficult dynamic behavior, being a non-self regulating process (integrating response), with variable dead time (transportation lag between 6-12 minutes) and exhibiting inverse response.

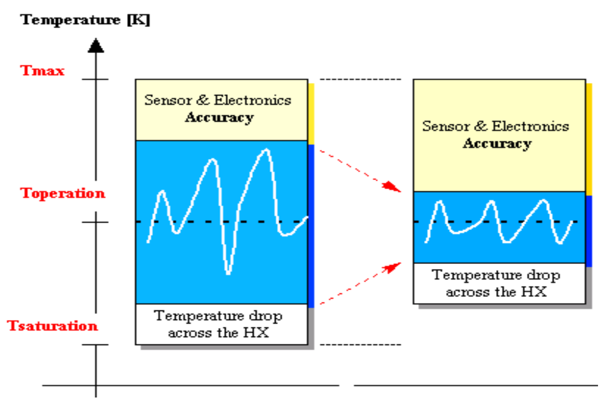

Fig. 3 - Advanced control motivation

An additional aim in implementing MPC on the plant was optimizing its operation. As can be seen in Fig. 3 reducing the variance in the magnets temperature will allow, either operating at a higher temperature setpoint without violating the upper constraints (which implies money savings because of the reduced demand on the cryogenic system) or admitting less instrumentation accuracy (which also implies savings in design and construction).

The fact that trying to squeeze as much as possible the control band is a strong constraint and a full justification for the choice of a MPC technology. The violation of this constraint would imply an eventual high-cost shutdown during normal operation.

3. PLANT MODELLING

Nonlinear predictive control (NMPC) is a natural extension of the linear MPC technique. The algorithm is again based in the use of an internal plant model, this time a nonlinear one which captures the main process characteristics. A key element in NMPC is the nature of the internal model. Several alternatives are possible including first-principle models, neural nets, Volterra series, etc. In our case a physical model was used trying to balance the capture of the process dynamics under several operating conditions and the simplicity of the representation.

A non-linear model based on physical laws and balances has been developed and validated using real experimental data obtained in the IT-HXTU installation. The implementation of this first principles model provides precise predictions over a very different operational conditions having into account changes in the saturation pressure and existing dynamic heat loads. Nevertheless, for control purposes, a simplified model is considered, based on some assumptions and equations (Blanco, 2001).

All magnets are assumed to operate at equal temperature T. Considering a single cold mass temperature simplifies the model and the heat transfer calculation through the interconnections. An energy balance leads to:

$$\frac{d}{dt}(m_{cm}C_p(T)T) = Q_{ss} + Q_{tr} - q_{cool} \tag{1}$$

where Q are heat loads and the cooling provided by the heat exchanger, q_{cool} , is calculated through

$$q_{cool} = HA_{ws}(T - T_s) \tag{2}$$

where the heat transfer coefficient, H, is estimated experimentally, the saturation temperature, T_s, is obtained by direct measurement of the saturation pressure, and the wetted area A_{ws} is estimated from the helium II mass accumulated in the heat exchanger tube by the following calculation

$$A_{ws} = f_1 L_{ws} \tag{3}$$

where L_{ws} is the wetted length on the HX and it is calculated through.

$$L_{ws} = f_2 \frac{m_{HX}}{\rho(T_s)} \tag{4}$$

f_1 and f_2 being strong non-linear functions of mass depending of the geometry of the pipe.
The accumulated helium II mass in the heat exchanger is calculated by

$$\frac{dm_{HX}}{dt} = lfjt - Fr - Fv \tag{5}$$

where the l_{fjt} is the liquid flow passing through the Joule-Thomson valve, Fr, the helium II liquid overflow and Fv the helium which evaporates in the HX, having also into account the vapour fraction, flash, produced by the Joule-Thomson valve.

$$l_{fjt} = m_{fjt}(1 - vflash) \tag{6}$$

where m_{fjt} represents the total mass flow passing through the JT valve which depends on the valve characteristic, and *vflash* the vapour fraction produced. This is computed from an enthalpy balance between the incoming high-pressure stream and the two coexisting phases at saturation pressure at the output of the JT valve.

$$vflash = \frac{H_{liq}(before) - H_{liq}(after)}{H_{gas}(after) - H_{liq}(after)} \tag{7}$$

Finally, the vapour on the heat exchanger is composed by two components, vapour fraction produced by the JT valve, and evaporation of the He II liquid.

$$Fv = g_{fjt} + \frac{q_{cool}}{h_{fg}} \tag{8}$$

g_{fjt} represents the gas flow produced by the JT valve and it is calculated by (9), being h_{fg} the latent heat of vaporization of liquid helium.

$$g_{fjt} = m_{fjt} \cdot vflash \tag{9}$$

The JT valve is characterized by a calibration curve and its opening represents the input of the model, v_{open}, and the constants, *cti*, give the valve characteristic.

$$mfjt = ct1 \cdot v_{open}{}^2 + ct2 \cdot v_{open} + ct3 \tag{10}$$

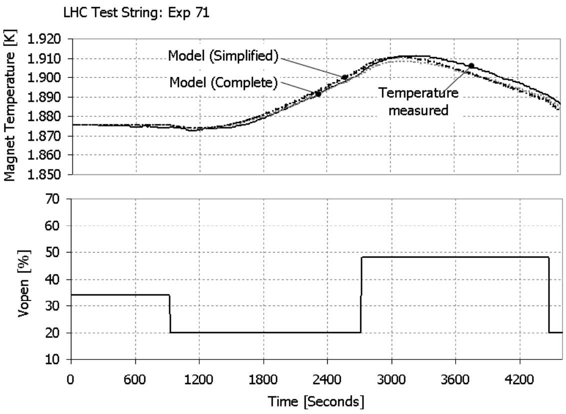

Fig. 4 Simplified vs. Complete first principles model

A comparison between the simplified model versus the complete model shows the magnet temperature when the JT valve is moved (Fig.4). A good trade off between complexity and quality of the model is obtained.

The value of some parameters, i.e. the cold mass and the heat transfer, was estimated by off-line optimization of the model errors.

4. PREDICTIVE CONTROLLER

The objective of the non-linear model predictive control (NMPC) is finding the future optimal manipulated variable sequence in order to minimize a function based on a desired output trajectory over a prediction horizon. The cost function is the integral of the sum of squares of the residuals between the model predicted outputs and the setpoint values over the prediction horizon N_2, plus a penalty term. A typical formulation is

$$\min_{u(k/k),\ldots,u(k+Nu-1/k)} \left(\int_{t_k}^{t_k+N_2} [\gamma(y(t) - r(t))]^2 \, dt + \sum_{j=0}^{Nu-1} \beta [\Delta u(k+j)]^2 \right) \tag{11}$$

where y and u are the process output and input. The minimization (11) is done subject to the continuous model equations and the typical restrictions applied on the manipulated and controlled variables:

$$\begin{aligned} u_{\min} &\leq u(k) \leq u_{\max} \\ \Delta u_{\min} &\leq \Delta u(k) \leq \Delta u_{\max} \\ y_{\min} &\leq y(k) \leq y_{\max} \end{aligned} \qquad (12)$$

Of the Nu moves optimal control sequence, only the first component is implemented. The optimization is solved using the scheme of Fig.5. Within this schema, the model equations are not explicit restrictions to the optimisation problem, being the manipulated variables the only decision variables. The simulation package will integrate the model equations along the prediction horizon taking as initial conditions the current process state and evaluating the formulated objective at the end of the integration. Path constraints are implemented as penalty functions. A simultaneous solution approach was also tested, but convergence and computation time did not improve the sequential one.

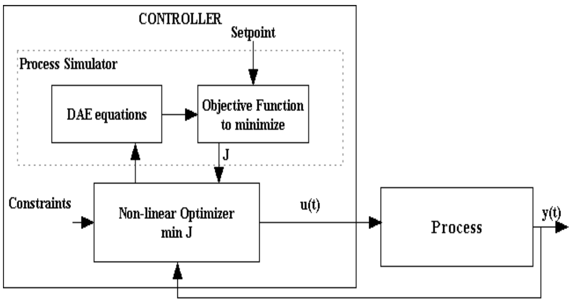

Fig. 5 Nonlinear controller – Continuous implementation framework

4.1 Nonlinear State Estimator

In our plant, the liquid helium II accumulated in the HX tube is not measurable. This is a critical factor of the model predictions, not only because it is a state but because it provides the wetted area in the heat exchanger from where the heat in the pressurized helium is removed. So, as we have an incomplete state vector, in order to apply the NMPC, a method of reconstructing the current state of the system from the measured outputs must be included. There are different approaches to the state estimation problem. We have chosen a receding horizon one (Muske and Edgar, 1997) because it match very well within the predictive control framework, it allows easy extensions to the non-linear case maintaining the same model as in the controller and with explicit inclusion of constraints in the variables and, finally, because state disturbances can be computed as a sub-product of the estimation. In our case it is really important estimating the non measurable overall *heat load* because it highly influences the model predictions. LHC prototypes operation has shown the variance on this disturbance in short periods of time.

In analogy to the model predictive control concept, the estimation problem is formulated as an optimal control problem on a finite horizon into the past. In the framework of the receding horizon estimation a quadratic cost function penalizing, among other things, model and measurement errors, is minimized. The optimisation problem is subject to model equations. Physical limits on the process variables are incorporated through inequality constraints.

More precisely, the problem is to estimate the initial conditions at time instant *t-N*, and the state disturbances, which have driven the process to its present state applying the past control sequence, by minimizing the difference between the outputs given by the evolution of the system from its initial conditions and the actual measured outputs (Fig. 6). In our case this can be translated into estimating the liquid helium mass in the HX tube at time instant t-N and the heat loads in the range [t-N, t-1], so that, if the JT valve were operating as in the real plant in this time interval, the computed temperature would approach the measured one and the heat loads are as small as possible.

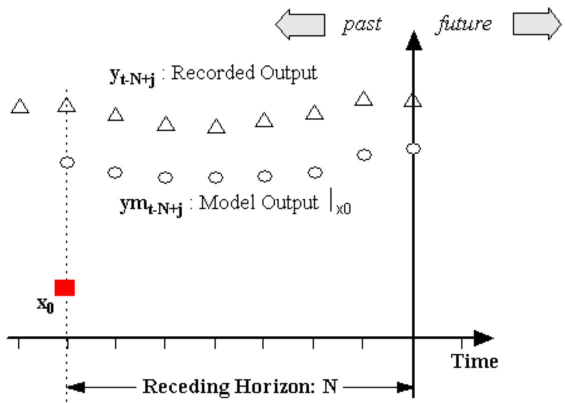

Fig. 6 Receding Horizon State Estimation

The standard approach can be synthesized as the optimization problem:

$$\min_{m_0, Q_{k-N},\ldots \hat{w}_k} J = \gamma_1 \sum_{j=0}^{N} \left(ym_{k-N+j} - y_{k-N+j} \right)^2 + \gamma_2 \sum_{j=0}^{N} \left(Q_{k-N+j} \right)^2 \qquad (13)$$

where N represents the receding horizon, one of the parameters to tune. y_{k-N+j} is the real process output and ym_{k-N+j} is the output of the model starting with $m_0=mHX$ at the time (k-N) when the past controls and the estimated process disturbances Q_{k-N+j} are applied.

Once the initial state and disturbances are estimated, the unknown state at time t can be computed integrating the model with the optimal values obtained.

Nevertheless, when applying (13) we realize that, due to the particular estructure of our process model, they were many solutions able to provide a perfect fit between the measured and computed temperatures. So, in our case, instead of (13) an alternative criteria was formulated, where, on one hand, the model temperature was equated to the measured one in the interval, and a new cost function was defined based on: (A) minimize the initial difference between the mass and its estimated value in the previous iteration mHX_{k-N} (B) penalization of the mass changes only if the JT valve was smooth during the horizon, in other case, where the JT was active, this contribution is cancelled, and finally (C) minimize the heat load change with respect to the value estimated in the previous iteration which provides a smooth JT valve moves.

With this structure the objective becomes

$$\min_{m_0} J = \gamma_0 [A]^2 + \gamma_1 \sum_{j=0}^{N-1} [B]^2 + \gamma_2 \sum_{j=1}^{N-1} [C]^2 \quad (14)$$

where γ_i weights the contribution of each factor and A, B, C represent

$$\begin{aligned} A &= m_{t-N} - m_{HX\,k-N}{}^{rec} \\ B &= \frac{m_{HX\,j+1}{}^{cal} - m_{HX\,j}{}^{cal}}{1 + \left[vopen_{j+1}{}^{rec} - vopen_{j}{}^{rec} \right]} \\ C &= Q_{tr_j}{}^{cal} - Q_{tr_j}{}^{rec} \end{aligned} \quad (15)$$

where $m^{cal}{}_{HXk-N}$ represents the liquid helium II in the heat exchanger already calculated N periods before, $m_{HXj}{}^{cal}$ the mass calculated along the receding horizon, $vopen_j{}^{rec}$ the manipulated variable (JT valve) moves recorded during the receding horizon, $Qtr_j{}^{cal}$ the calculated dynamic heat load and $Qtr_j{}^{rec}$ the previously calculated and recorded dynamic heat load during the horizon.

Estimating only the initial state and heat load disturbances makes the problem tractable. The advantage of this solution is the smaller number of decision variables for any given horizon length, resulting in less computational time to solve the optimization problem (Gelb, 1974). In this case the state disturbance, the heat load present in the process, is included in the estimation procedure by means of the model equations. A good starting point helps in the estimation, which can be found if the process starts operating in a known steady state like helium overflowing. Besides (14), the optimization includes constraints on the values of the decision variables.

The control structure designed for the nonlinear controller incorporates a nonlinear predictive algorithm and a state estimator. The solution proposed yields a new approach based of an initial state estimate and of a moving horizon algebraic estimator in a combined structure. The state estimator provides the optimal mass accumulated in the heat exchanger tube and the dynamic heat load valuation. A block diagram of the structure can be seen in Fig. 7.

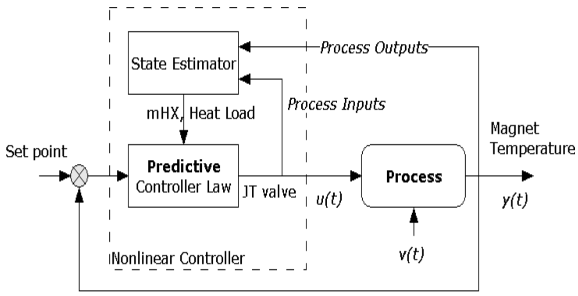

Fig. 7 NMPC proposed control structure

4.2 Simulation results

Simulation studies have been performed in order to verify the performance of the implemented estimator and the improvement on the control. An example is shown in the Fig. 8 where several steps of the heat load (12 Watts) were applied and the values of the estimated state disturbance are compared against the values given by the model showing good agreement. The only tuning parameter, apart from the weights, in the proposed objective structure, is the receding horizon value N (here, N=4). The process is represented by the full simulation model.

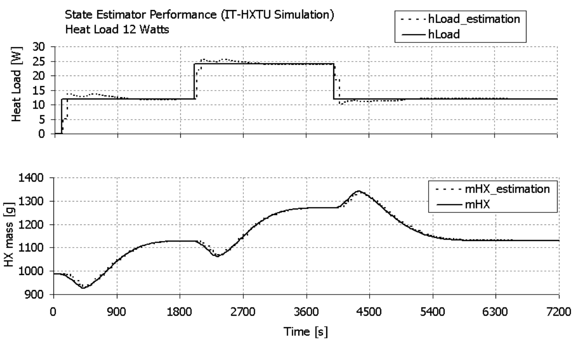

Fig. 8 Performance of the Nonlinear State Estimator

The complexity of the optimization problem is not growing proportionally to the length of the horizon due to the fact that only one initial value is considered as a decision variable and the remainders are algebraic calculations. N has not much influence on neither the controller nor the state estimator performance. In the test, only the parameter γ_0 conditions the way the heat load is estimated, the greater the number, the faster the heat load disturbance estimation.

4.3 Optimization: numerical solutions

Both, the controller law solution and the state estimation problems, presented in the nonlinear predictive controller framework, lead to the same nonlinear programming problem, which could be formulated generically as a real time minimization of a nonlinear function subject to constraints. These constraints could be simple bounds on the variables and both, linear and nonlinear constraints. In the case of the LHC 1.8 K Cooling Loop the method used is a SQP one, due to its ability to solve problems with nonlinear constraints.

5. EXPERIMENTAL RESULTS

The validation of the state estimator module based on the receding horizon was done experimentally by powering the electrical heaters located in the cold mass. These simulate a change in the overall heat load due to a unknown disturbance, then data was stored corresponding to the electrical watts applied and the heat load estimation carried out by the state estimator.

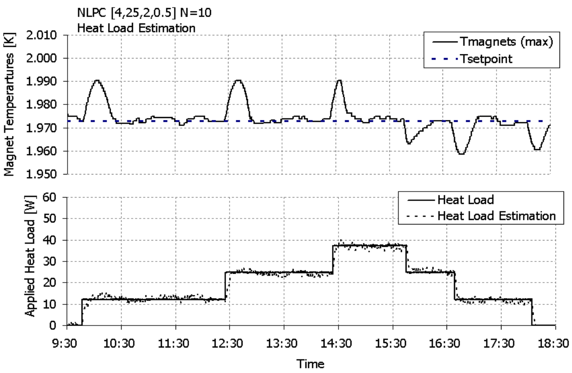

Fig. 9 . RHE performance. Heat load steps

In Fig. 9 several step changes on the heat load were applied to the process. Performance of the state estimator is fast and precise, and the heat load is estimated immediately after its change despite the abrupt jump. This situation could be produced by several factors in the real systems, for example, a degradation of the insulation vacuum which leads to a higher existing overall heat load. The other state estimated, the accumulated helium mass in the HX tube is not shown because is a non-measured variable and no comparison are possible. Performance of the controller is also displayed in the same Fig. 9. The temperature excursions, due to the heat load applied, are cancelled around 1.99 K in all the different operational zones showing a robust behavior of the regulator.

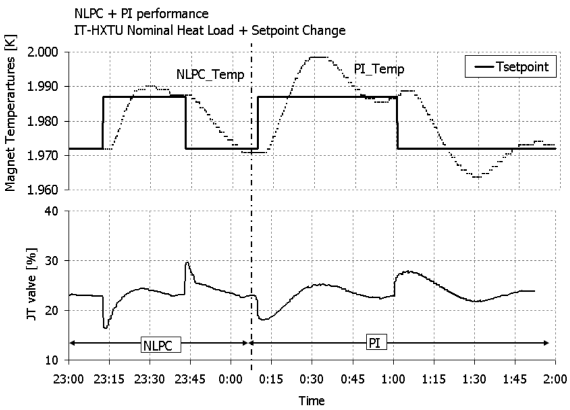

Fig. 10 NMPC vs. PI control: Tracking characteristics

Once the state estimator was tested and validated, more experiments were performed in order to validate the nonlinear predictive controller. Changes in the set point were considered to check also for tracking features (Fig. 10). A comparison with a classical PI controller is done in order to show the improvements in terms of control performance and robustness.

6. CONCLUSIONS

The LHC full-scale prototypes were employed as a test-bed of what advanced nonlinear control can do for improving for cryogenic processes regulation. The nonlinear process model construction gave a better understanding of the process, provided the ideas to overcome the usual changes in process dynamics and helped to improve the regulation strategies by means of the simulation. The response has been improved and optimized by the use of the nonlinear predictive controller with a receding horizon state estimator. The regulation structure proposed is based in a nonlinear predictive controller algorithm combined with a state estimator with an initial state estimate approach and a moving horizon algebraic calculation for the disturbance.

ACKNOWLEDGEMENTS

The authors would like to thank the financial support of the project PPQ2000-1075-C02-01 by the Spanish CICYT and are very grateful to the LHC Cryogenics group (ACR) and the String and IT-HXTU teams.

REFERENCES

Blanco, E., C. Prada, J. Casas and S. Cristea (1999). Predictive Temperature Control of the LHC Superconducting Magnets. *European Control Conference ECC'99*, Karlsruhe, Germany.

Blanco, E. (2001). Nonlinear Model-Based Predictive Control Applied to Large Scale Cryogenic Processes. In *Ph.D. Thesis, Universidad de Valladolid*, Valladolid, Spain.

Cristea, S. (1998). Control predictivo en el dominio delta. In *Ph.D. Thesis, Universidad de Valladolid*, Valladolid, Spain.

Casas, J., et al., (1998) Operation Testing and Long Term Behavior of the LHC Test String Cryogenic System. In *ICEC17 Proceedings*, edited by Hughes. D., et al., UK, pp. 747-750.

Gelb, A. (1974). Applied Optimal estimation. MIT Press, Massachusetts

Muske, K.R. and T.F. Edgar (1997). Nonlinear State Estimation. In *Nonlinear Process Control* (M.A. Henson and D.E. Seborg, (Eds)), pp. 311-370. Prentice-Hall, Inc., New Jersey