1

Introduction

This thesis focuses on the constrained optimal control of hybrid systems. Hybrid systems are heterogenous systems incorporating both continuous-valued components, which are governed by differential or difference equations, and discrete-valued components, such as finite state machines, if-then-else rules, on/off switches, digital circuits and software code. Hybrid systems switch between different operating modes, where each mode is governed by a dynamical law. Mode transitions are triggered by variables crossing specific thresholds, by the elapse of certain time periods, or by external inputs. With respect to hybrid systems, modelling, controller synthesis, monitoring schemes for estimation and fault detection, verification of safety properties, and analysis of stability, robustness and performance have recently become a very active area of research, both in control engineering and computer science. This is due to the fact that hybrid systems not only pose many theoretical challenges, but also offer novel ways and opportunities to improve on traditional control schemes. Most important, enabled by the rapid and steady progress in the computational power available, many intrinsically "hard" yet practically relevant and traditionally unsystematically solved problems related to the modelling, analysis and control of hybrid systems can be nowadays successfully tackled. We will consider linear hybrid systems in the discrete-time domain given in Mixed Logical Dynamical (MLD) [BM99a] or Piecewise Affine (PWA) [Son81] form.

For hybrid systems, control schemes are proposed that are based on discrete-time constrained finite-time optimal control with a receding horizon policy, often referred to as Model Predictive Control (MPC) [Mac02]. In MPC, the current control input is obtained by solving at each sampling instant an open-loop constrained optimal control problem over a finite horizon using the current state of the plant as the initial state. The underlying optimization procedure yields an optimal control sequence that minimizes a given objective function. The receding horizon policy refers to only applying the first control input of this sequence and to recomputing the control sequence at the next sampling instant over a shifted horizon, thus providing feedback and closing the control loop. The significant advantages of MPC, including its ability to systematically cope with hard constraints on manipulated variables, states and outputs, and to easily address systems with multiple inputs and outputs, have led to its success and widespread use, which initiated in the process industry more than two decades ago. Recent survey papers on MPC include [MRRS00] and [QB03], while the monographs [CB99] and [Mac02] provide introductions to the various classes of MPC.

As introduced in [BM99a], the MLD framework can be straightforwardly embedded in MPC allowing one to use hybrid models given in MLD form as prediction models for MPC. The underlying optimization problem is a Mixed-Integer Program, for which efficient off-the-shelf solvers exist [ILO02]. In practical applications, however, the computation time for solving the optimal control problem on-line often exceeds the sampling interval thus prohibiting the direct implementation of the controller. This obstacle is overcome by pre-computing off-line the solution to the optimal control problem for the whole state-space using Dynamic Programming and multi-parametric programming, where the state vector is treated as a parameter. For hybrid systems, such a method has been recently introduced, which is based on a PWA description of the controlled system [Bor03]. The result is a PWA state-feedback control law that can be easily implemented in form of a look-up table.

1.1 Mode Enumeration and Optimal Complexity Reduction

Introduction

When deriving models and designing controllers, in particular for hybrid systems, two issues commonly arise: The derivation of the set of (feasible) modes of the model, and the complexity reduction of PWA systems comprising both PWA models and PWA state-feedback control laws.

The modelling stage is often and most easily performed using the HYbrid Systems Description Language (HYSDEL). Hybrid models given in HYSDEL can be considered as compositions of Discrete Hybrid Automata (DHA) [TB04], which are a mathematical abstraction of the features provided by other computation oriented and domain specific hybrid system frameworks. In general, compositions of hybrid models are very complex, as the number of different operational modes depends exponentially on the number of component systems. The explosion of the number of possible modes leads to computational difficulties as the time and space complexity of most algorithms depends on it. Yet, most of these modes are infeasible due to the model dynamics, their interaction and additional constraints. To compute the set of (feasible) modes of a composition of models is beneficial for several reasons, among them being that these modes not only allow for reducing the computational burden of related algorithms, but also form a basis to efficiently translate HYSDEL code into PWA form.

On the other hand, the complexity of PWA systems, which is approximately given by the number of polyhedra, is often required to be as small as possible. As an example consider the case where the state-feedback control law is computed off-line based on a PWA model. Due to the combinatorial nature of the problem, both the computation time and the controller complexity are in the worst case exponential in the number of polyhedra of the PWA model [Bor03]. On the other hand, once the PWA state-feedback control law has been derived and is implemented as a look-up table in hardware, the memory requirement and the on-line computation time are linear in the number of polyhedra of the feedback law when using standard search techniques. Hence it is often of major importance to obtain a control law of minimal complexity.

Contributions

Both problems are addressed here. Using the notion of cell enumeration in hyperplane arrangements from computational geometry, we propose in Chapter 3 an algorithm that efficiently enumerates all feasible modes of a composition of DHAs. The impact of those techniques on applications is threefold. At the modelling stage, the enumeration of modes allows the designer to understand the real complexity of the compound model. After the modelling, the model can be efficiently translated into a PWA representation, which the model is generally required to be in when deriving the PWA state-feedback control law. Compared to a recently published related algorithm for deriving the PWA model [Bem02], the one presented here is of one to two orders of magnitudes faster. During the computational stage (i.e. analysis and control), the explicit computation of the set of feasible modes of the compound system can be used as structural information to prune infeasible modes from the resulting model and thus to reduce the computational burden of related algorithms, like optimal control schemes. Furthermore, the presented algorithm is able to deal with loops that may be present in compositions, and to determine if a composition is well-posed or not.

The information provided by the mode enumeration algorithm, namely the so-called markings, can be also used to determine *a priori* – i.e. without solving any Linear Program (LP) – if a given combination of polyhedra is convex. Exploiting this fact, we propose in Chapter 4 two algorithms that solve the problem of deriving a PWA model that is both equivalent to the former and minimal in the number of regions. The first algorithm executes a branch and bound on the markings yielding a new set of disjoint polyhedra, where additional heuristics on the branching strategy are employed to reduce the computation time. The second approach relies on the fact that the optimal complexity reduction problem can be reformulated as a logic minimization problem by replacing the markings by Boolean variables and minterms [Kat94]. Logic minimization is a fundamental problem in digital circuit, and efficient tools have been developed to successfully tackle these problems, which often encounter hundreds or thousands of variables [BHMS84]. The resulting polyhedra are in general not disjoint and thus overlapping. As both algorithms refrain from solving additional LPs, they are not only optimal but also computationally feasible. The applicability of the algorithms can be extended to general PWA systems lacking the hyperplane arrangement (like PWA state-feedback control laws) by first computing the hyperplane arrangement. In many cases, the optimal complexity reduction enables the implementation of the optimal controllers as look-up tables in hardware.

Publications

Chapter 3 is almost entirely based on

[GTM03a] Geyer, T., F.D. Torrisi and M. Morari: *Efficient Mode Enumeration of Compositional Hybrid Models*. Technical Report AUT03-01, Automatic Control Laboratory ETH Zurich, http://control.ee.ethz.ch/, 2003.

Preliminary results have appeared in

[GTM03b] Geyer, T., F.D. Torrisi and M. Morari: *Efficient Mode Enumeration of Compositional Hybrid Systems*. In Pnueli, A. and O. Maler (editors): *Hybrid Systems: Computation and Control*, volume 2623 of *Lecture Notes in Computer Science*, pages 216–232. Springer-Verlag, 2003.

Apart from the logic minimization scheme, which is an extension, Chapter 4 is based on

[GTM04] Geyer, T., F.D. Torrisi and M. Morari: *Optimal Complexity Reduction of Piecewise Affine Models Based on Hyperplane Arrangements*. In *Proceedings of the American Control Conference*, pages 1190–1195, Boston, MA, June 2004.

Software Codes

The mode enumeration algorithm is implemented in Matlab and assumes that the composition of DHAs is given as Hysdel code. The latest version of the algorithm can be downloaded from http://control.ethz.ch/~hybrid/hysdel. Also the complexity reduction algorithms are written in Matlab. They are included in the multi-parametric toolbox (MPT) [KGBM04], which is freely available from http://control.ee.ethz.ch/~mpt/.

1.2 Direct Torque Control

Introduction

The rapid development of power semiconductor devices led to the increased use of adjustable speed induction motor drives in a variety of applications. In these systems, DC-AC inverters are used to drive induction motors as variable frequency three-phase voltage or current sources. One methodology for controlling the torque and speed of induction motor drives is Direct Torque Control (DTC) [TN86], which features very favorable control performance and implementation properties.

The basic principle of DTC is to exploit the fast dynamics of the motor's stator flux and to directly manipulate the stator flux vector such that the desired torque is produced. This is achieved by choosing an inverter switch combination that drives the stator flux vector to the desired position by directly applying the appropriate voltages to the motor windings. This choice is made usually with a sampling time $T_s = 25\,\mu$s using a pre-designed switching table that is derived in a heuristic way and, depending on the particularities of the application, addresses a number of different control objectives. These primarily concern the induction motor – more specifically, the stator flux and the electromagnetic torque need to be kept within pre-specified bounds around their references. In high power applications, where three-level inverters with Gate Turn-Off (GTO) thyristors are used, the control objectives are extended to the inverter and also include the minimization of the average switching frequency and the balancing of the inverter's neutral point potential around zero. Due to the discrete switch positions of the inverter, the DTC problem is a hybrid control problem, which is complicated by the nonlinear behavior of the torque, length of stator flux and the neutral point potential.

Contributions

We aim at deriving MPC schemes that are conceptually and computationally simple yet yield a significant performance improvement with respect to the state of the art. More specifically, the term *conceptually simple* refers to controllers allowing for straightforward tuning of the controller parameters or even a lack of such parameters, and easy adaptation to different physical setups and drives, whereas *computationally simple* implies that the control scheme does not require excessive computational power to allow the implementation on DTC hardware that is currently available or at least will be so within a few years.

To achieve this, we exploit in Chapter 6 a number of physical properties of DTC drives to derive discrete-time models of DTC drives with two- or three-level inverters tailored to our needs, more specifically, models that are of low complexity yet of sufficient accuracy to serve as prediction models for our model-based control schemes. These properties are the (compared with the stator flux) slow rotor flux and speed dynamics, the symmetry of the voltage vectors, and the invariance of the motor outputs under flux rotation. The low-complexity models are derived by assuming constant speed within the prediction horizon, mapping the states (the fluxes) into a 60 degree sector, and aligning the rotor flux vector with the d-axis of a reference frame rotating with the rotational speed of the rotor. The benefits of doing this are a reduction of the number of states from five to three, and a highly reduced domain on which the nonlinear functions need to be approximated by PWA functions.

Based on the hybrid models of the DTC drive, we propose in Chapter 7 three novel control approaches to tackle the DTC problem, which are inspired by the principles of MPC and tailored to the peculiarities of DTC. The first scheme uses soft constraints to model the hysteresis bounds on the torque, stator flux and neutral point potential, and approximates the average switching frequency (over an infinite horizon) by the number of switch transitions over a short horizon. To make this approximation meaningful and to avoid excessive switching, the *Late Switching Strategy* has to be added, which favors the postponement of switch transitions. Three penalty levels with corresponding penalties of different orders of magnitude provide clear controller priorities and make the fine-tuning of the objective function obsolete, and the *Multiple Prediction Model Approach* allows us to extend the prediction interval without increasing the computational burden. This control scheme not only leads to short commissioning times for DTC drives, but it also leads to a performance improvement in terms of a reduction of the switching frequency in the range of 20 % with respect to the industrial state of the art, while simultaneously reducing the torque and flux ripples. Yet the complexity of the control law is rather excessive.

The second scheme exploits the fact that the control objectives only weakly relate to optimality but rather to feasibility, in the sense that the main objective is to find a control input that keeps the controlled variables within their bounds, i.e. a control input that is feasible. The second, weaker objective is to select among the set of feasible control inputs the one that minimizes the average switching frequency, which is again approximated by the number of switch transitions over the (short) horizon. We therefore propose an MPC scheme based on feasibility in combination with a move blocking strategy, where we allow for switching only at the current time-step. For each input sequence, we determine the number of steps the controlled variables are kept within their bounds, i.e. remain feasible. The switching frequency is emulated by the cost function, which is defined as the number of switch transitions divided by the number of predicted time-steps an input remains feasible, and the control input is chosen so that it minimizes this cost function. The simplicity of the control methodology translates into a state-feedback control law with a complexity that is of an order of magnitude lower than the one of the first scheme, while the performance is improved.

The third scheme can be considered as a combination of the two preceding concepts. Specifically, we use a rather short horizon and compute for the input sequences over the horizon the evolution of the controlled variables using the prediction model. To emulate a long horizon, the "promising" trajectories are extrapolated and the number of steps is determined when the first controlled variable hits a bound. The cost of each input sequence is then determined by dividing the total number of switch transitions in the sequence by the length of the extrapolated trajectory. Minimizing this cost yields the optimal input sequence and the next control input to be applied. The major benefits of this scheme are its superior performance in terms of switching frequency, which is reduced over the whole range of operating points by up to 50 %, with an average reduction of 25 %. Furthermore, the controller needs no tuning parameters. As the computation of an explicit solution is avoided, all quantities may be time-varying including model parameters, set points and bounds. Those can be adapted on-line, making the concept applicable to the whole range of operating points. As all computations are performed on-line, the prediction model is not restricted to be PWA, allowing us to use the nonlinear (and more accurate) discrete-time model.

Summing up, all control schemes are based on minimizing an approximate of the average switching frequency, they use an internal model of the DTC drive to predict the output response to input sequences, and they are tailored to the specific DTC problem set-up. Starting from the first scheme, the complexity of the controllers in terms of computation times and the memory requirement is steadily reduced by several orders of magnitude, while the performance is steadily improved. In particular the last control scheme is expected to be implementable on the currently available DTC hardware.

Publications

Chapters 5, 6 and 7 are based on a re-arrangement and slight extension of

[PGM04c] Papafotiou, G., T. Geyer and M. Morari: *Optimal Direct Torque Control of Three-Phase Symmetric Induction Motors*, 2004. submitted to journal.

[GP05] Geyer, T. and G. Papafotiou: *Direct Torque Control for Induction Motor Drives: A Model Predictive Control Approach based on Feasibility.* In Morari, M. and L. Thiele (editors): *Hybrid Systems: Computation and Control*, volume 3414 of *Lecture Notes in Computer Science*, pages 274–290. Springer-Verlag, 2005.