Mechatronics 24 (2014) 975–985

Contents lists available at ScienceDirect

Mechatronics

journal homepage: www.elsevier.com/locate/mechatronics

Tuning hydrostatic two-output drive-train controllers using reinforcement learning

Kevin Van Vaerenbergh [a,*], Peter Vrancx [a], Yann-Michaël De Hauwere [a], Ann Nowé [a], Erik Hostens [b], Christophe Lauwerys [c]

[a] Computational Modeling Lab, Vrije Universiteit Brussel, Pleinlaan 2, 1050 Brussels, Belgium
[b] Flanders' Mechatronics Technology Centre vzw, Celestijnenlaan 300 D Bus 4027, 3001 Heverlee, Belgium
[c] XTOCON BVBA, Tervuursevest 23 Bus 1806, 3001 Heverlee, Belgium

ARTICLE INFO

Article history:
Received 30 September 2013
Revised 2 June 2014
Accepted 23 July 2014
Available online 20 August 2014

Keywords:
Interacting subsystems
Reinforcement learning
Global controller
Hydrostatic two-output drive-train

ABSTRACT

When controlling a complex system consisting of several subsystems, a simple divide and conquer approach is to design a controller for each system separately. However, this does not necessarily result in a good overall control behavior. Especially when there are strong interactions between the subsystems, the selfish behavior of one controller might deteriorate the performance of the other subsystems. An alternative approach is to design a global controller for the entire mechatronic system. Such a design procedure might result in more optimal behavior, however it requires a lot more effort, especially when the interactions between the different subsystems cannot be modeled exactly or if the number of parameters is large.

In this paper we present a hybrid approach to this problem that overcomes the problems encountered when using several independent subsystems. Starting from such a system with individual subsystem controllers, we add a global layer which uses reinforcement learning to simultaneously tune the lower level controllers. While each subsystem still has its own individual controller, the reinforcement learning layer is used to tune these controllers in order to optimize global system behavior. This mitigates both the problem of subsystems behaving selfishly without the added complexity of designing a global controller for the entire system. Our approach is validated on a hydrostatic drive train.

1. Introduction

When multiple independent subsystems are present in a global mechatronic system, the subsystem controllers are not aware of the interactions among them and will behave selfishly. This may not only deteriorate the performance of the other subsystem controllers but also put the control of the entire system at risk [1,2]. Alternatively, instead of designing individual controllers for the subsystems, one global controller can be designed. This approach requires more effort as a more complex problem needs to be solved as a whole, but this extra effort will result in the overall system to become more optimal when compared to multiple independent subsystems [3,4]. This paper presents the results of tuning a global controller, which controls the gains of 2 local PI controllers [5], using reinforcement learning (RL) [6] on a hydrostatic two-output drive-train. Tuning of PI(D) controllers with reinforcement learning was previously proposed in e.g. [7,8], however in this work we consider the novel scenario of jointly tuning multiple local controllers in order to optimize global performance.

In the following section, we give a brief explanation of the hydrostatic two-output drive-train followed by a description of the distributed control architecture for the hydrostat set-up, using PI controllers to track the reference speeds for individual hydromotors. We then explain how we use reinforcement learning to jointly tune the gains for these controllers in order to optimize the global system behavior. Reinforcement learning (RL) is a model free learning method for achieving optimal control in Markov decision processes. This approach does not require any prior knowledge of the dynamics of the system to be controlled. Instead, optimal controls are discovered through repeated interaction with the system. RL has been successfully applied in a wide range of control applications, both within and outside the field of mechatronics [9–12]. We explain RL in more detail in Section 4. The results of our method, both in simulation and on the real setup are described in Section 6.

* Corresponding author. Tel.: +32 2 629 37 07; fax: +32 2 629 37 08.
E-mail address: kevvaere@vub.ac.be (K. Van Vaerenbergh).

http://dx.doi.org/10.1016/j.mechatronics.2014.07.005
0957-4158/

2. Hydrostatic two-output drive-train

The hydrostatic drive-train (Fig. 2) [13] was developed as a demonstrator for advanced control techniques. Fig. 1 shows the working of the drive-train. An electric motor drives a hydraulic pump, which is hydraulically connected to two hydraulic motors. Each hydraulic motor drives a flywheel, that is in turn connected to another electric motor. The electric motors can generate any torque, and as such simulate any load. The electric motor driving the pump can also be used to simulate any engine, e.g. an internal combustion engine that cannot recuperate energy.

The hydraulic pump and hydraulic motors have a variable displacement volume, which is the amount of fluid that flows for one revolution of the pump/motor. Neglecting the losses, the ratio of the pump displacement volume over the motor displacement volume equals the transmission ratio between the pump shaft and the motor shaft.

To develop a simulation of this setup, an in-depth search was performed to mathematically define the system with system equations. These equations are then used to model the hydrostatic drive-train in SimScape.

In Fig. 3 we have a schematic representation of the whole setup. For our experiments, both T_{des1} and T_{des2} are zero. Therefore we can define the setup as being an interaction between the three motors (PUMP, HM1 and HM2). All three motors are similar and are simulated using the following system equations. Eq. (1) is the first-order model for the stroke volume and Eq. (2) defines the velocity of the hydraulic component. While Eq. (3) is the speed-flow hydraulic component associated with volumetric efficiency. The last two Eqs. (4) and (5) are two equations responsible for the torque-pressure hydraulic component associated with hydrostatic yield.

$$S_{x\ act}(s) = \frac{1}{\tau_p \cdot s + 1} \cdot S_x(s) \tag{1}$$

$$V_x = \frac{V_{x\ 210} - V_{x\ 35}}{210bar - 35bar} \cdot (\Delta p_x - 35bar) + V_{x\ 35} \tag{2}$$

$$Q_x = -\eta_{x\ vol}(\Delta p_x, \omega_y, S_{x\ act}) \cdot V_x \cdot S_{x\ act} \cdot \omega_y \tag{3}$$

$$T_{fric\ x} = T_x \cdot (1 - \eta_{x\ mec}(\Delta p_x, \omega_y, S_{x\ act}) \cdot \tanh\left(\frac{\omega_y}{0.01\,\frac{rad}{s}}\right) \tag{4}$$

$$T_x = \Delta p_x \cdot V_x \cdot S_{x\ act} + T_{fric\ x} \tag{5}$$

where the subscript x stands for either one of the hydraulic components (pump and hydromotors) and y is the *DrAx* for the PUMP and the *LdAx* for HM1 and HM2. For the hydromotors 1 and 2 the T_x is negative. Some of the values are constants and are defined here:

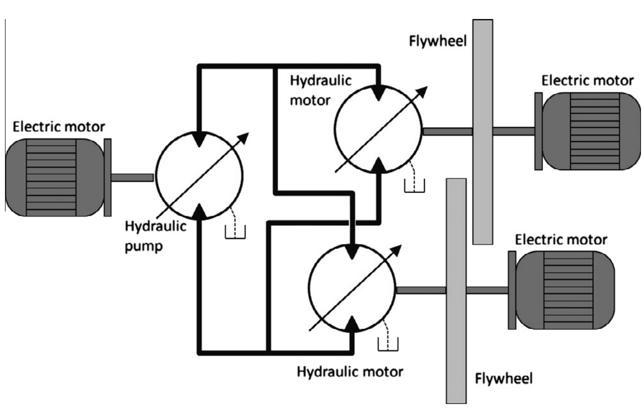

Fig. 1. Schematic drawing of the hydrostatic system.

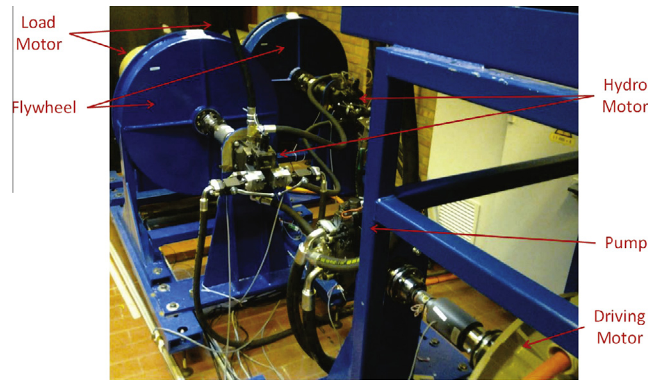

Fig. 2. Photograph of the real hydrostatic system.

$$\begin{aligned} V_{p\ 210} &= 113.31\,\frac{cc}{rev} \quad \text{and} \quad V_{p\ 35} = 113.31\,\frac{cc}{rev} \\ V_{HM1\ 210} &= 59.9\,\frac{cc}{rev} \quad \text{and} \quad V_{HM1\ 35} = 59.3\,\frac{cc}{rev} \\ V_{HM2\ 210} &= 59.7\,\frac{cc}{rev} \quad \text{and} \quad V_{HM2\ 35} = 59.1\,\frac{cc}{rev} \end{aligned} \tag{6}$$

In this paper we adopt a simplified control strategy using standard industrial PI-controllers to control the hydrostatic drive-train. This control strategy is explained in the following section (used to be pressure dependent).

3. Hydrostat controller

The distributed controller for the hydrostat designed by XTO-CON consists of three individual modules with a limited set of available measurements, setpoints and controlled manipulated variables. The goal of this controller is to track dynamic reference trajectories for both hydraulic motors. The first controller, *the pump controller*, has a measurement of the oil pressure and the reference trajectories of the two motors available. This controller steers the swash plate position and the velocity of the pump engine. The second and third controllers, *the motor controllers*, each have a measurement available of their corresponding individual hydraulic motor velocity and reference trajectory. Additionally, they are also allowed to use a measurement of the pressure.

The control strategy is based on the two physical insights: given a constant pressure, increasing a motor swash position plate leads to an increased motor torque which accelerates the motor and leads to a higher motor velocity, and given fixed motor swash plate positions, increasing the pump flow (by either increasing the engine speed or the pump swash plate positions) leads to a higher pressure. This control strategy leads to the basics of the controllers' design strategy, let the flow controller (pump controller) steer the pressure and let the motor controllers track the reference trajectories. The design is based on the following observations of the physics of the hydrostat. The basic ideal state equations are:

$$\begin{aligned} &S_1 * w_1 + S_2 * w_2 = S_p * w_p && \text{(flow balance)} \\ &\frac{T_p}{S_p} = \frac{T_1}{S_1} = \frac{T_2}{S_2} && \text{(torque balance)} \\ &w_1 = f_1(T_1) \text{ with } \tfrac{df_1}{dT_1} < 0 && \text{(stable load on motor 1)} \\ &w_2 = f_2(T_2) \text{ with } \tfrac{df_2}{dT_2} < 0 && \text{(stable load on motor 2)} \end{aligned}$$

S_1 and S_2 are the swash plate positions of the motors. T_1 and T_2 are the torques of the motors. S_p is the swash plate position and T_p the torque of the pump. f_1 and f_2 are two load functions, expressing the relation between the speed on the load and the motors. Then the following assumptions are made which were experimentally validated. Increasing the pump speed (7) and increasing the pump