Engineering Applications of Artificial Intelligence 23 (2010) 249–261

Contents lists available at ScienceDirect

Engineering Applications of Artificial Intelligence

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

Energy dispatch controllers for a photovoltaic system

Ganesh Kumar Venayagamoorthy*, Richard L. Welch

Real-Time Power and Intelligent Systems Laboratory Missouri University of Science and Technology, Rolla, MO 65409, USA

ARTICLE INFO

Article history:
Received 16 August 2008
Received in revised form
24 July 2009
Accepted 19 November 2009
Available online 12 January 2010

Keywords:
Adaptive critic design
Energy dispatch controller
Neural network
Photovoltaic
Solar
PV-priority controller

ABSTRACT

In this paper two energy dispatch controllers for use in a grid-independent photovoltaic (PV) system are presented. The first, an optimal energy dispatch controller, is based on a class of Adaptive Critic Designs (ACDs) called Action Dependent Heuristic Dynamic Programming (ADHDP). This class of ACDs uses two neural networks to evolve an optimal control strategy over time. The first neural network or "Action" network dispenses the actual control signals while the second network or "Critic" network uses these control signals along with the system states to provide feedback to the action network, measuring performance using a utility function. This feedback loop allows the action network to improve behavior over time. The optimal energy dispatcher places emphasis on always meeting the critical load, followed by keeping the charge of the battery as high as possible so as to be able to power the critical load in cases of extended low output from the PV array, and lastly to power the non-critical load in so far as to not interfere with the first two objectives. The second energy dispatch controller is a smart energy dispatch controller and is built using knowledge from an expert, codified into a series of static rules. This smart energy dispatch controller is called the "PV-priority 2" controller. These energy dispatchers are compared with a static scheme called the "PV-priority 1". The PV-priority 1 controller represents the standard control strategy. Results show that the ADHDP-based optimal energy dispatcher (or controller) outperforms the standard PV-priority 1 energy dispatcher in meeting the stated objectives, but trails the PV-priority 2 energy dispatcher. However, the major advantage of the ADHDP controller is that no expert is required for designing the controller, whereas for a rule-based controller such as the PV-priority 2 controller, an expert is always required.

1. Introduction

As the costs of fossil fuels continue to rise, it is becoming economically important to investigate other sources of energy. Additionally, increased customer demands and loads are beginning to outstrip the grid's capacity to serve these loads. As such, distributed generation of alternative energy sources is becoming a very heated research and development area.

Currently, there are several alternative energy sources available: wind, solar, hydro-electric and geo-thermal, to name a few. Hydro-electric and geo-thermal plants often require large foot prints that are at odds with developed areas and environmental considerations. Wind energy is currently enjoying very energetic growth, at around 24% per year since the year 2000 in the US (US Department of Energy's Office of Energy Efficiency And Renewable Energy, 2007). But, wind power has its downsides as well since not all locations receive enough sustained winds to be productive, and even then production is sometimes sporadic. Of all of the mentioned sources, solar power seems to be the most promising in that all locations on Earth receive predictable sunlight to some degree, and the solar arrays used to convert sunlight into electricity (photovoltaic or PV arrays) scale very well from very small sizes for calculators to very large sizes used in centralized power plants.

Another benefit of solar energy is that the PV arrays contain no moving parts and can last several decades before needing to be replaced. During this time, the only maintenance that may need to be done to PV arrays is dust or snow clearing and checking for alignment problems. And while the PV arrays are within their rated lifetimes, they generally perform reliably while the Sun is shining.

Even with these advantages, there is one major drawback of PV systems that limits their adoption rate and that is the cost of these systems. The cost of energy derived from these systems makes them not currently competitive with other sources. However, due to technological, manufacturing and resource improvements, these costs have steadily fallen in previous years (Messenger and Ventre, 2004) and are expected to continue to do so. As the price falls, this source will be more competitive. Even so, the payback time (the amount of time required for the initial investment in a PV system to equal the costs accrued from purchasing power through the grid) can be lengthy, and in some cases as long as 30 years or more.

* Corresponding author. Tel.: +1 573 3416641; fax: +1 573 3414532.
E-mail address: gkumar@ieee.org (G.K. Venayagamoorthy).

0952-1976/$ - see front matter © 2009 Elsevier Ltd. All rights reserved.
doi:10.1016/j.engappai.2009.11.001

In order to make PV systems cheaper (and shorten the payback period), optimal control can be used to more effectively utilize energy generated by the PV array, resulting in smaller required batteries and arrays while still maintaining the critical load. This reduction in system component size leads to a direct reduction in cost for the entire system.

The traditional energy dispatcher for a PV system is called the "PV-priority" (called "PV-priority 1" in this paper) control scheme (Henze and Dodier, 2003) and will first attempt to power all loads using energy from the PV array; if there is not enough energy available from the PV array then energy from the battery is used to make up the shortfall (if available), and if there is more energy available from the PV array then the batteries are charged with the difference (if possible).

In this paper, two additional controllers are developed: the "PV-priority 2" and an optimal controller based on a class of adaptive critic designs (ACDs) called action-dependent heuristic dynamic programming, or ADHDP (Werbos, 1992; Venayagamoorthy et al., 2002; Prokhorov and Wunsch, 1997). The "PV-priority 2" control scheme is similar to the PV-priority 1 scheme in that its performance is rule based, but is different in that it attempts to always power the critical load first, then charge the battery to 70%, and finally use whatever energy is available from the PV array and anything over 70% state of charge in the battery to power the non-critical load.

The ADHDP-based optimal energy dispatcher on the other hand is not rule based, and develops its control action strategy by adapting its performance in response to a measured metric value. Adaptive critic designs use a combination of dynamic programming and reinforcement learning, and the ADHDP method is the simplest of the ACD family (it uses only 2 neural networks). One of the neural networks (called the "action" network or "actor") is responsible for providing the control signals while the second (called the "critic" network) critiques these control signals over time. The objectives of this energy dispatcher are the following:

i). Completely power the critical load at all times.
ii). Maintain the battery state of charge as high as possible so as to be able to meet the critical load during times of reduced (or non-existent) energy from the PV array.
iii). Power the non-critical load such that the controller is still able to meet the first two objectives.

Another advantage of the ADHDP-based controller is that since it is not rule based, its behavior can be modified over timntetro6(objectives.)]TJ-mleri5-1.]TJ2gidingcdiactives-1.]0(Ong)-261(such5627(simoodvaer device. Likewise, new energy storage systems (Jiang and Dougal, 2006; Lemofouet and Rufer, 2006) may eclipse the performance of the standard lead-acid-type battery technology primarily used today. Because the focus of this study is primarily to evaluate the performance of the presented control strategies, the assumption of 100% efficiency of the photovoltaic system is made. If other efficiencies are desired, they can be set by modifying the appropriate values within the models, as depicted in Fig. 1.

Also, the PV array is simulated to be tilted south at an angle equal to the latitude of each test city and the efficiency of the PV array model is taken as 11% to account for various non-optimal conditions (such as array misalignment, dust on the arrays, etc.). This value is representative of the current commercially available range of efficiencies for PV arrays. Generally, PV panels vary in efficiency from 6% to up to 30%; although the high efficiency panels are generally reserved for spacecraft usage because of their high radiation tolerances and higher power-to-weight ratio. A rough equivalent to the PV arrays being simulated in this paper would be an array of eight Kyocera KC200GT panels. These panels are over 16% efficient and will output 200 W during optimal conditions (Kyocera, 2007). The minimum charge for the battery of 30% is required to supply energy to the loads (this is consistent with standard deep cycle lead-acid batteries).

Due to insufficient PV energy during winter months and no PV energy at night, a control system is required to decide the amount of energy to be dispatched to the different loads, including the charging of the battery. The complete system in schematic diagram form is shown below in Fig. 1 (energy flow depicted by arrows).

3. PV-priority controllers

3.1. PV-priority 1 controller

The standard controller called the “PV-priority 1” controller is
a very simple controller which always tries to meet 2(is)-2resenta26(accou

This controller works well when there is sufficient PV energy.

non-critical loads.

$$\begin{aligned}U(t) = {} & A_1^* abs(1-(E_{CL}/(L_{CL}+M_{non-zero}^* L_{CL\,max}))) \\ & + A_2^* abs(1-(E_{CB}/((E_{B\,max}-E_B)+M_{non-zero}^* E_{B\,max}))) \\ & + A_3^* abs(1-(E_{NCL}/(L_{NCL}+M_{non-zero}^* L_{NCL\,max})))\end{aligned} \quad (2)$$

In the $U(t)$ function given in (2), a higher priority is given to meeting the critical load at all times over the batteries being charged or the non-critical load being supplied by assigning different weightings – 30/23 to the CL term, 15/23 to the BC term and 13/23 to the NCL term. This $U(t)$ meets the threefold objective for the optimal PV controller design.

In the training of the Critic network, the objective is to minimize (3) given below.

$$\sum_{t=0}^{\infty} E^2(t) \quad (3)$$

where

$$E(t) = U(t) + \gamma J(t) - J(t-1) \quad (4)$$

The weight change and update equations for the Critic network using the BP algorithm is given by (5) and (6), respectively.

$$\Delta W_c(t) = \eta_c\, E(t) \frac{\partial J(t)}{\partial W_c} \quad (5)$$

$$W_c(t+1) = W_c(t) + \Delta W_c(t) \quad (6)$$

where η_c and W_c are the learning rate and the weights of the Critic neural network, respectively.

One important note about the Critic network design is that when determining how many previous time steps to use, it is necessary to look at how the Critic network training is proceeding. If the critic network is not training properly, then it may be necessary to increase the number of time steps used as input to the critic network. The Critic network requires multiple time steps because it is implemented using a feedforward network and using multiple delayed time steps gives the network more information

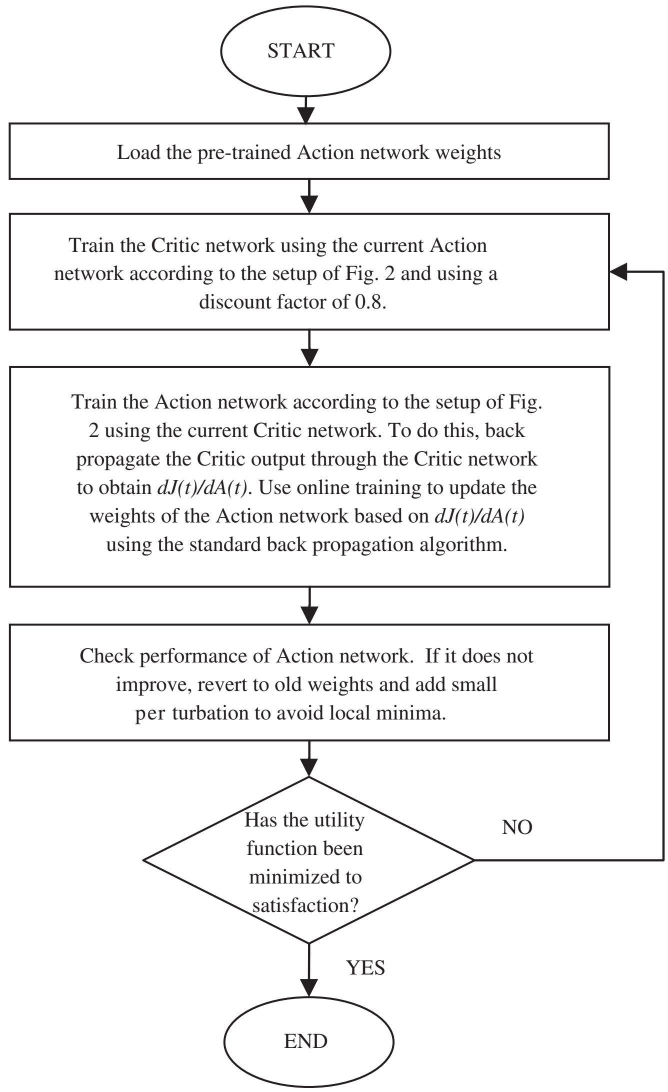

Fig. 8. Flowchart for Critic/Action network training.

Eqs. (8) and (9), respectively.

$$\Delta W_A(t) = \eta_A \, E_A(t) \frac{\partial A(t)}{\partial W_A} \tag{8}$$

$$W_A(t+1) = W_A(t) + \Delta W_A(t) \tag{9}$$

Here η_A and W_A are the learning rate and the weights of the Action neural network, respectively.

As with the Critic network, the size of the hidden layer in the Action network is determined by trial and error, while the size of the input and output layers are determined by the number of inputs and outputs, respectively.

4.3. *Actor/Critic training*

The flowchart in Fig. 7 outlines the pre-training steps for the Action network, while Fig. 8 details the iterative training technique used to develop the optimal controller over time. During the iterative training phase, several metrics can be used to determine if the Actor's performance is increasing. For this study, the simple sum of the utility function for each cycle of training the Action network is used. This means that when the sum of the utility function is decreasing, the performance of the Action network is improving. The simulation is run for a fixed number of iterations, but if the sum of the utility function increases during training then the new Action network weights are discarded and replaced with the previous best weights. When this happens, a very small perturbation (a random number between -0.01 and 0.01) is added to the Action network weights such that the network avoids getting stuck in a local minimum.

After the best Action network weights are found, these weights are then used to optimally dispatch energy to the critical loads, the non-critical loads and the battery.

5. Results

A one year simulation of the PV system is carried out for the following areas under varying conditions with multiple controllers: Phoenix, AZ, Miami, Fl, Boulder, CO, Springfield, MO, Caribou, ME and Fairbanks, AK. These simulations use data from the TMY2 database (National Renewable Energy Laboratory, 1995). Phoenix receives more solar radiation than Miami, Miami more than Boulder, Boulder more than Springfield, etc. The actual relationship between cities and total annual solar insolation is shown in Table 1. A map showing the solar insolation received on a flat plate collector facing south and tilted at the latitude angle for the United States (with investigated cities marked) is shown in Fig. 9 (National Renewable Energy Laboratory, 2004).

For three of the six cities (Phoenix, Springfield and Caribou), an action network is separately trained and then these networks are simulated against the data for each city. Selected results of these simulations (along with results from using the PV-priority

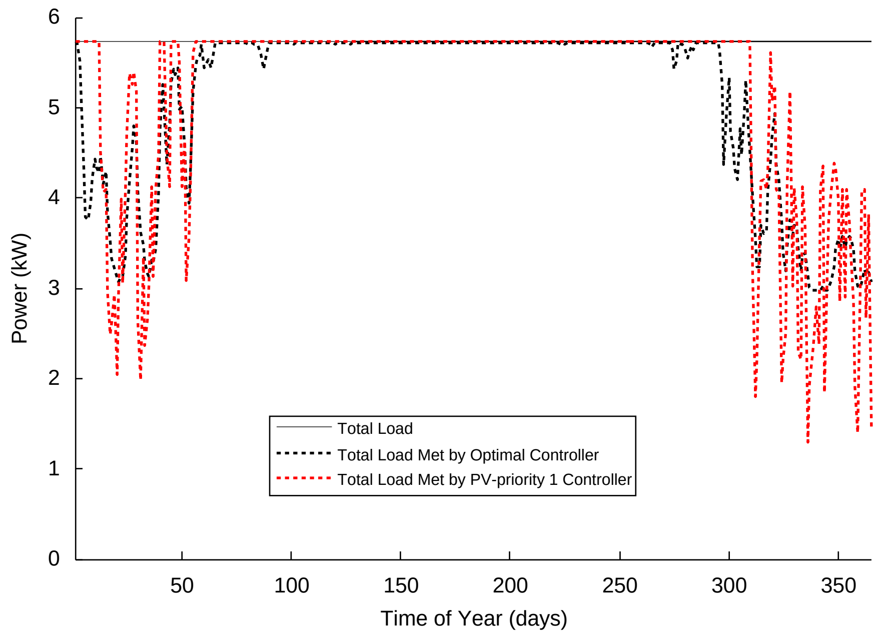

Fig. 12. Performance of the PV-priority 1 controller and optimal controller (trained using data from Caribou, ME) in meeting the total load in the city of Springfield, MO.

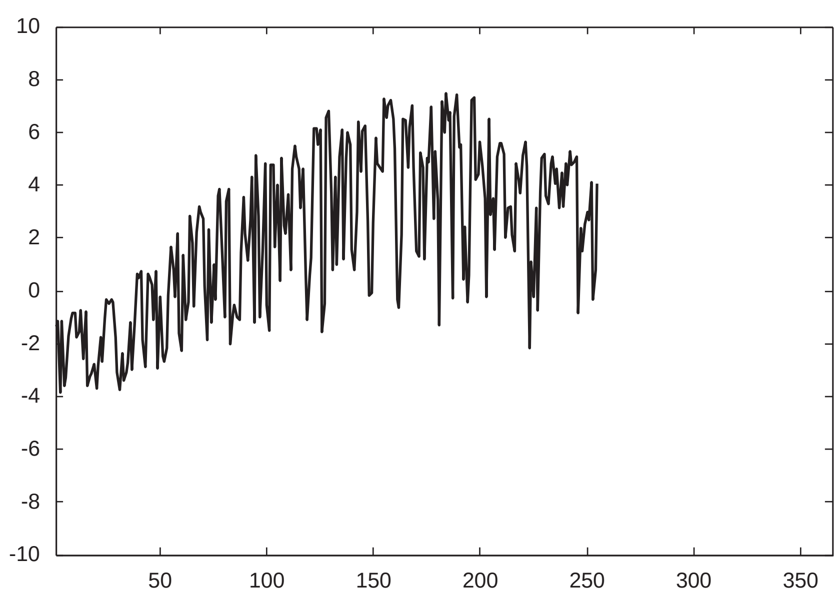

demanded load and satisfied load on a daily basis, so the demanded load curve is
varies on
9ontroller is able to 9ompletely outperform the PV-priority 1 9ontroller in powering the critical load but does so at the expense of powering the load.

In Fig. 13, the daily energy is for the same
This surplus (or deficit) is calculated by the demanded
energy of the loads from the energy received form the PV arrays. This is
9ontroller. If there is always a surplus, then the 9ontroller has an easier job of supplying energy to the loads, but there is always 0
deficit, then the 9ontroller would have 02588(more)-587(difficult)-586(time)]TJ0-1.316TD[(supplying)-332(energy)-3 the without dropping th

Interestingly, optimal PVtrollerrained using the Caribou,
ME data seems to always equal or performetter than any other
9ontroller except for the PV-priority 2 9ontroller. Also, i seems that the performance of the PV-priority 1 controller is generally slightly better than the optimal PV 9ontroller trained using data from the Phoenix, AZ especially as the becomes more
demanding andessolarnsolation available.

Fig. 14 shows the battery state of charge for Springfield, MO for period2642(of)-644(late)-649(fall)-643(and)-643(early)-648(winte than the 9onventional PV-priority 1 9ontroller. Additionally, it meets much the load but a little less the
9ritical load. It is also to state of note that this
9ontroller changed behavior as the battery charge increased. When it was it less of the
load and focus on the 9ritical load, and when it was more fully

charged it would attempt to power both loads. This is in sharp contrast to the PV-priority controller 1 which always tried to power both loads, leading to a much lower average battery

Furthermore, it can be seen that as the load increases, the margins by which the Caribou-trained optimal controller and PV-priority 2 controller outperforms the others increases. This is time to ensure that the critical load demand is met primarily all the time and then the non-critical load demand. The battery state charge is also maintained as high as possible to ensure energy supply to the critical loads during nights and the winter months. This in turn provides the benefit of extended battery life. Results have been presented for six different cities with different solar radiation profiles using five controllers: 2 PV-priority controllers