Performance and Fiscal Analysis of Distributed Sensor Networks in a Power Plant

Mitchell Colby
Oregon State University
442 Rogers Hall
Corvallis, OR 97331
colbym@engr.orst.edu

Kagan Tumer
Oregon State University
204 Rogers Hall
Corvallis, OR 97331
kagan.tumer@oregonstate.edu

ABSTRACT

As power plants become more complex and system models become difficult or impossible to generate, the plants must be split into subsystems to be controlled independently. Distributed sensor networks show promise for this application, because they are inherently modular and have the ability to preprocess data to help with the control process. We investigate how effectively a distributed sensor network could be implemented in a model power plant, showing that system parameters may be accurately measured and tracked more effectively than traditional power plant sensors. In addition, a distributed sensor network is also shown to be capable of detecting system changes that are undetectable by standard sets of sensors. Finally, a cost analysis is performed to show when a distributed sensor network is a financially viable alternative to a standard set of sensors.

Categories and Subject Descriptors

H.3.4 [**Systems and Software**]: Distributed Systems

General Terms

Algorithms, Experimentation,Economics

Keywords

Multiagent learning, Coordination

1. INTRODUCTION

One of the most significant challenges to developing efficient energy sources is addressing how to control and optimize power plants. As power plants become more and more complex, and models for such plants become difficult or impossible to generate, a distributed control strategy will become necessary [12]. Rather than a central controller making decisions for an entire power plant, subsystems within the plant must be defined and controlled independently, while maintaining effectiveness for the entire plant.

Sensors are becoming smaller, less expensive, more computationally powerful, and more capable of operating in harsh environments [4]. Traditionally, sensors in power plants simply record data and communicate this data with a central controller, which makes system level decisions [5]. In order for the control to be decentralized, it is helpful if sensors can record data, preprocess data, and handle data requests.

Given the complexity of power plants and the power of new sensors, a distributed sensor network is a natural system for a power plant. Increasing the number of sensors in a system provides many benefits. First, with a large system of sensors, the network can compensate for sensor failures. Secondly, with the ability of sensors to preprocess data, sensor networks can give much more useful data relating to system wide performance than a smaller set of sensors. Finally, sensors capable of interacting with other sensors can give system level information that is not available from simply aggregating sensor information.

This paper demonstrates that a distributed sensor network is preferable to a standard set of sensors in a power plant. Five key components provide evidence for this claim:

1. A distributed sensor network can track system parameters such as temperature and pressure
2. A distributed sensor network is robust to sensor failures.
3. A distributed sensor network with low accuracy sensors can measure system parameters more accurately than a single high accuracy sensor
4. A distributed sensor network is capable of detecting changes in the system that go unnoticed by standard sensors.
5. A distributed sensor network is competitive with a standard sensor implementation in terms of fiscal cost

For this work, a multiagent system controls a distributed sensor network in order to measure enthalpy, and track changing enthalpy profiles. We use a modified version of the Defect Combination Problem, as well as the Difference Reward, to give feedback to reinforcement learning agents. The contribution of this paper is two-fold. We show that:

1. The Difference Reward can be used to train reinforcement learning agents controlling a distributed sensor network in a model power plant to achieve the first four objectives outlined above
2. A distributed sensor network which meets or exceeds the performance thresholds of a traditional sensor implementation can be developed for a competitive cost.

Ultimately, we show that a sensor network using a large number of relatively cheap and less accurate sensors can outperform a costly more accurate sensor for measuring the same parameters. The rest of this paper is organized as follows. Section 2 gives the background information and related work. Section 3 explains the modified Defect Combination Problem and how it is implemented in this work.

Section 4 explains the experiments conducted. Section 5 gives the experimental results. Finally, Section 6 gives a discussion of the results and presents possibilities for future research.

2. BACKGROUND

The following sections introduce the Difference Reward, the physics governing the model power plant, the Defect Combination Problem, and related work.

2.1 Difference Reward

The Difference Reward is defined as [8]:

$$DU_i = G(z) - G(z_{-i}) \tag{1}$$

where $G(z)$ is the system reward, and $G(z_{-i})$ is the system reward without the effects of agent i. Intuitively, this reward function tells an agent its specific impact on the system utility, and is thus aligned with the system reward. Further, as this reward depends only on the actions of agent i, there is less noise introduced to the signal from other agents, making the learning process easier for each individual agent.

By being factored and aligned with the global system reward, the Difference Reward gives a private utility which allows individual agents to improve overall system performance. This is an attractive feature which becomes more valuable as the number of agents and complexity of a system increases [9, 10].

2.2 Rankine Cycle

Rather than attempting to design a sensor network to operate in a complex power plant with no system model, we develop a sensor network in a well known power generation system: a vapor power Rankine cycle. In a Rankine cycle, the working fluid passes through a boiler and becomes saturated vapor. Next, the fluid goes through the turbine, which results in an energy output which is used to produce electricity. The fluid then passes through the condenser and becomes a saturated liquid. Finally, the fluid is passes through a pump and returns to the boiler, completing the cycle [5]. The Rankine cycle is shown in Figure 1. For the purposes of this analysis, we make the following assumptions:

A1. Each component of the cycle is considered to be a control volume

A2. All processes of the working fluid are internally reversible

A3. The turbine and pump operate adiabatically

A4. Kinetic and potential energy effects are negligible

A5. Saturated vapor enters the turbine. Condensate exits the condenser as a saturated liquid

A6. The working fluid is assumed to be water

As seen in Figure 1, there are four distinct states in the Rankine cycle, each of which lies between two of the components. The system performance is related to the enthalpy h_i at each plant state i by the following relations:

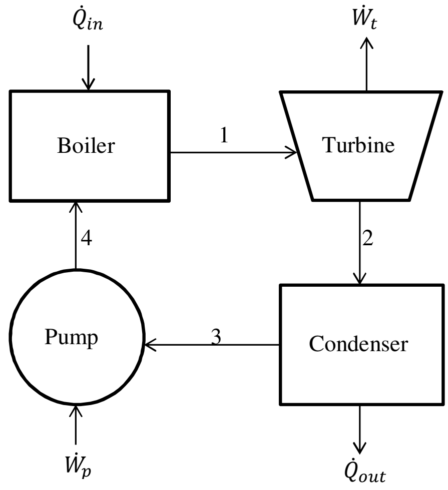

Figure 1: A vapor power Rankine cycle. The working fluid travels through a boiler, turbine, condenser, and pump in succession. The work output of the turbine is used to generate electricity.

$$\frac{\dot{W}_t}{\dot{m}} = h_1 - h_2 \tag{2}$$

$$\frac{\dot{Q}_{out}}{\dot{m}} = h_2 - h_3 \tag{3}$$

$$\frac{\dot{W}_p}{\dot{m}} = h_4 - h_3 \tag{4}$$

$$\frac{\dot{Q}_{in}}{\dot{m}} = h_1 - h_4 \tag{5}$$

where $\dot{m}$ is the mass flow rate of the working fluid, $\dot{W}_t$ is the work output of the turbine, $\dot{Q}_{out}$ is the heat output of the condenser, $\dot{W}_p$ is the work input to the pump, and $\dot{Q}_{in}$ is the heat input to the boiler. In order to find the enthalpy at each of these states, the pressure and temperature must be determined, requiring a sensing policy to be developed. Such a policy was studied in the Defect Combination Problem.

2.3 Defect Combination Problem

The Defect Combination Problem (DCP) assumes that there exists a set of imperfect sensors X which have constant attenuations due to manufacturing defects or imperfections [1]. Each of the sensors $x_i \in X$ has an associated attenuation a_i in its reading. Thus, if sensor x_i is taking a measurement of A, it measures $A + a_i$. The DCP involves choosing a subset of the X sensors such that the aggregate attenuation of the combined readings is minimized:

$$G = \frac{\left|\sum_{i=1}^{N} n_i a_i\right|}{\sum_{i=1}^{N} n_i} \tag{6}$$

where G is the aggregated attenuation of the combined sensor readings, and $n_i \in \{0, 1\}$ is an indicator function based on whether the sensor chooses to be "on" or "off."

Early work on the DCP defined each sensor as an agent, which then chose whether to be on or off [8]. Each sensor learns from feedback provided by the Difference Reward.

The rewards of each agent are thus factored with respect to the global reward, such that agents acting to improve their private rewards also act to minimize the aggregate system attenuation. By using the Difference Reward to provide individual feedback for each agent, the aggregate error was much lower than the error resulting from global rewards. Private rewards eliminate the need for centralized control.

2.4 Related Work

Traditional Power Plant Control.

Traditionally, power plants are controlled by some centralized control system, which receives information from sensors distributed across the plant. This type of control typically requires a thermodynamic model of the system. Processes taking place in power plants are sufficiently complex that idealizations are necessary to create these models [5]. As power plant designs become more complex, system models become increasingly difficult to generate, and idealizations lead to inaccurate models. This results in traditional control methods becoming difficult to implement in new power generation systems. In a complex plant with no system model, a decentralized control system with a distributed sensor network becomes an attractive option. Subsystems of the plant may be accurately modeled (or models may be learned), eliminating the need to derive a full system model for centralized control. An adaptive control strategy and distributed sensor network implementation becomes more desirable as power plant complexity increases [3].

Organization for Area Surveillance.

Previous work on distributed sensor networks demonstrated that a sensor network can self-organize when sensors are deployed with no *a priori* information regarding the environment [6]. The sensors use a max-sum algorithm to coordinate their sense-sleep schedule in order to maximize the effectiveness of the entire sensor network. This research is very attractive in the context of a power plant, because the complexities associated with the physics model result in *a priori* optimal sensor locations being impossible to obtain. Furthermore, as the optimal network policy may change with system parameters, a distributed sensor network and in a power plant should be adaptive [7].

Detection of False Readings.

Previous work on analyzing sensor readings yielded an approach to find false sensor readings in sensor networks [2]. This is an essential step to take in a sensor network, because taking actions on false sensor readings can be extremely costly. One issue with this work is that the world dynamics are explicitly modeled. In a complex power plant, this type of model may be impossible to generate. Thus, rather than attempting to explicitly detect false readings, the distributed sensor network in the power plant has redundant sensors which act to minimize aggregate attenuation.

3. DCP FOR POWER PLANTS

We apply a slightly modified version of the DCP to a Rankine cycle power plant. There is a set of sensors X_s at each of the four plant states, where $s \in \{1, 2, 3, 4\}$ is the state of the power plant (see Figure 1). Each sensor $x_{s,i} \in X_s$ is capable of measuring pressure and temperature, the two parameters needed to determine the enthalpy of the working fluid. Sensor $x_{s,i}$ also has a temperature attenuation $t_{s,i}$, and a pressure attenuation $p_{s,i}$, an effective temperature operating range $[T_{s,i,min}, T_{s,i,max}]$, and an effective pressure operating range $[P_{s,i,min}, P_{s,i,max}]$. If a sensor measures values outside of its effective operating range, then the measurement error increases exponentially with the distance from the value to the effective operating range bounds.

Each sensor is considered to be an agent. First, each agent must decide whether to be on or off. If an agent decides to be on, then it must determine whether it will measure temperature, pressure, or both temperature and pressure. The goal of the agents is to collectively take actions which minimize the aggregate system attenuation. The aggregate attenuation for temperature at state s is defined as:

$$g_{T,s} = \frac{\sum_{i=1}^{N_s} n_{s,i} t_{s,i}}{\sum_{i=1}^{N_s} n_{s,i}} \tag{7}$$

where N_s is the number of sensors in state s, and $n_{s,i} \in \{0, 1\}$ denotes whether sensor $x_{s,i}$ is on or off. Similarly, the aggregate attenuation for pressure at state s is defined as:

$$g_{P,s} = \frac{\sum_{i=1}^{N_s} n_{s,i} p_{s,i}}{\sum_{i=1}^{N_s} p_{s,i}} \tag{8}$$

From equations 7 and 8, if state s has a true temperature of T_s and true pressure P_s, then the measured values will be:

$$T_{s,sensed} = T_s + g_{T,s} \tag{9}$$

$$P_{s,sensed} = P_s + g_{P,s} \tag{10}$$

The enthalpy of the working fluid is a thermodynamic property which quantifies the level of energy in that fluid. Enthalpy change in a fluid corresponds to the fluid either absorbing or expelling energy, and is used to determine power levels in a power cycle. Enthalpy data has been found experimentally, and given pressure and temperature, the enthalpy of a fluid can be found with empirical thermodynamic tables. Thus, the measured enthalpy at state s is given by:

$$h_{s,sensed} = f(T_{s,sensed}, P_{s,sensed}) \tag{11}$$

where $h = f(T, P)$ is found using thermodynamic tables, which are based on empirical data. The error in the enthalpy reading at state s is given as:

$$h_{s,error} = h_s - h_{s,sensed} \tag{12}$$

where h_s is the actual enthalpy at state s. It is desirable that we calculate errors in enthalpy reading rather than errors in temperature or pressure readings, because enthalpy values can be obtained from the plant output. For example, we control the mass flow rate of the working fluid $\dot{m}$, and we monitor the power output of the turbine, $\dot{W}_t$, by monitoring the generator, we can determine the true values of h_1 and h_2 with the turbine model. These values can then be compared to the estimated values provided by the sensor network. One key issue with implementing the DCP in practice is that it is assumed that the sensor attenuations are known. This is rarely the case in real-world problems, so we must use information that is actually available in order to determine the accuracy of the sensor measurements. We can determine the enthalpy values at each state by monitoring system outputs, and then use this information to judge the effectiveness of the sensors measuring pressure and temperature. The objective of the entire sensor network is to minimize the total attenuation of enthalpy estimations, given by:

$$G = \sum_{s=1}^{4} |h_{s,error}| \tag{13}$$

Intuitively, this simply seems like four separate instantiations of the DCP. Although this is true mathematically, this does not take into account the coupling between each state and the components of the system. For example, if states 1 or 4 are measured incorrectly, the amount of heat added to the boiler is affected, resulting in wasted energy by adding too much heat or an inefficient cycle by adding too little heat. In addition, a change in enthalpy at one state will affect the enthalpy of all other states. Minimizing the total attenuation at each state will provide the best possible data to the controllers in the power plant, such that the total power output of the Rankine cycle may be maximized. The power output of the Rankine cycle is given by:

$$\dot{E}_{out} = \dot{W}_t - \dot{W}_p - \dot{Q}_{in} - \dot{Q}_{out} \tag{14}$$

In order to demonstrate that a distributed sensor network is a viable alternative to a traditional sensor implementation in the Rankine cycle, certain properties of the sensor network must be demonstrated. First, the sensor network must measure system parameters with at least the same level of accuracy as a traditional sensor implementation. If the sensor network gives more accurate measurements than standard sensors, then system phenomena which go undetected by standard sensors will be detected by the sensor network. As the sensors in a sensor network are necessarily less expensive and accurate than the single sensors they replace, a policy must be devloped such that a set of inaccurate sensors has a lower aggregate attenuation than a single accurate sensor. The sensor network must accurately track changing parameters, in order for the plant controllers to act based on useful data. Finally, the distributed sensor network must be shown to be a cost-effective alternative. In order to demonstrate these properties, a specific model regarding sensor performance and cost must be developed.

Sensor Model.

When considering replacing a traditional set of sensors with a distributed sensor network, cost becomes an important issue. In general, the cost of a sensor increases exponentially with its accuracy. In this analysis, the cost of a sensor C is related to its attenuation A by the following relation [11]:

$$C(A) = c_0(1 + c_1 A^{-\beta_1}) \tag{15}$$

Where c_0 is the *fixed* cost of a sensor, and the second term captures the variable cost which decreases as sensor attenuation increases. For the purposes of an optimization problem, we may assume that c_0 is 1 without loss of generality. For a distributed sensor network to be a feasible option to replace traditional sensors, the accuracy of each individual sensor must be lower than that of a single sensor, in order to prevent the sensor network from being more expensive than the standard sensors used. Thus, a distributed sensor network will contain less accurate sensors than a standard implementation of sensors. In particular, each sensor in the network is defined by the values Φ_{min}, Φ_{max}, and a, where:

- Φ_{min} is the minimum value of the sensor's operating range
- Φ_{max} is the maximum value of the sensor's operating range
- a is the attenuation of the sensor while sensing in its operating range

For example, for a temperature sensor, Φ_{min} is the minimum temperature that the sensor can effectively measure, Φ_{max} is the maximum temperature that the sensor can effectively measure, and a is the attenuation of the sensor when sensing temperatures between Φ_{min} and Φ_{max}. If a sensor is measuring values outside of its operating range, the error in its measurement increases exponentially as the parameter being measured moves away from the effective operating range. The attenuation a_{OR} of a sensor measuring values outside of its effective operating range is:

$$a_{OR} = |\Phi - \Phi_{bound}|^2 \tag{16}$$

where Φ is the parameter being measured and Φ_{bound} is the closest bound of the sensor's operating range to Φ.

3.1 Agent Learning

For the enthalpy measurement and tracking, we use a multiagent reinforcement learning algorithm, where each agent keeps an individual value table and updates using the difference reward. The state is the sensed temperature and pressure T_{sensed} and P_{sensed}, and the action is whether the sensor is off, measures temperature, measures pressure, or measures both temperature and pressure. Action selection is completed with an ϵ-greedy policy. Thus, each agent selects the best action with probability $1 - \epsilon$ and a random action with probability ϵ. The reinforcement learning algorithm is standard Q-learning, and is shown in Algorithm 1.

Algorithm 1 Reinforcement Learning Algorithm

Each agent i generates a randomly seeded Q-table Q_i;
$episode = 1$;
while $episode < maxEpisodes$ **do**
  1. Each agent i measures system state $s_i = \{T_{s,i}, P_{s,i}\}$;
  2. Each agent selects an action from Q-table using ϵ-greedy;
  3. Calculate total enthalpy attenuation $G(\vec{z})$;
  4. Calculate difference reward for each agent D_i;
  5. Q-update: $Q_i(s,a) \leftarrow Q(s,a)(1-\alpha) + \alpha D_i$;
  6. $episode = episode + 1$;
end while

4. EXPERIMENTAL SETUP

The following sections detail the experiments conducted in order to show the effectiveness of distributed sensor networks in a power plant. First, we determine if the distributed sensor network can accurately estimate enthalpy by measuring temperature and pressure. Next, we determine if sensor failures in the network can be compensated for. Next, we use the distributed sensor network to track a changing enthalpy profile, to determine if the sensor network can detect changes in the system during transient operation. Finally, we analyze the cost versus performance of a distributed sensor network to determine the cost-effectiveness of the sensor network against traditional sensor implementations.

4.1 Enthalpy Measurement

For the first experiment, a single state of the Rankine cycle is randomly selected in order to determine if the enthalpy value can be accurately estimated with a distributed sensor network. Suppose we have a state with actual temperature, pressure, and enthalpy values of T_{actual}, P_{actual}, and h_{actual}. A set of N sensors is placed in the state, and the sensors can choose to measure temperature, pressure, or both temperature and pressure. Given the sensor attenuations, the measured values of temperature and pressure are:

$$T_{measured} = T_{actual} + g_T \tag{17}$$
$$P_{measured} = P_{actual} + g_P \tag{18}$$

Where g_T and g_P are the aggregate attenuations defined in Equations 7 and 8. Thus, the measured enthalpy of the state based on the measurements is:

$$h_{estimated} = f(T_{measured}, P_{measured}) \tag{19}$$

where $f(T, P)$ is based on empirical thermodynamic data. The objective function to be minimized for single state enthalpy estimation is given by the error in enthalpy estimation:

$$G_1(\vec{z}) = |h_{actual} - h_{estimated}| \tag{20}$$

In this case, h_{actual} is determined by the system model, given the knowledge of the control inputs $\dot{W}_p$, $\dot{Q}_{in}$, and $\dot{Q}_{out}$, as well as the total power output of the plant $\dot{W}_t$. The task of the sensor network is to select a subset of sensors to measure temperature, and a subset of sensors to measure pressure, in order to minimize the error in the enthalpy reading. Each agent is given the difference reward as a learning signal. The difference reward for this experiment is derived as follows:

$$g_{T,-i} = \frac{\sum_{j=1}^{N_s} n_{s,j} t_{s,j}}{\sum_{j=1}^{N_s} n_{s,j}} - \left. \frac{\sum_{j=1}^{N_s} n_{s,j} t_{s,j}}{\sum_{j=1}^{N_s} n_{s,j}} \right|_{n_{s,i}=0} \tag{21}$$
$$T_{measured,-i} = T_{actual} + g_{T,-i} \tag{22}$$
$$g_{P,-i} = \frac{\sum_{j=1}^{N_s} n_{s,j} p_{s,j}}{\sum_{j=1}^{N_s} n_{s,j}} - \left. \frac{\sum_{j=1}^{N_s} n_{s,j} p_{s,j}}{\sum_{j=1}^{N_s} n_{s,j}} \right|_{n_{s,i}=0} \tag{23}$$
$$P_{measured,-i} = P_{actual} + g_{P,-i} \tag{24}$$
$$h_{measured,-i} = f(T_{measured,-i}, P_{measured,-i}) \tag{25}$$
$$G_1(\vec{z}_{-i}) = |h_{actual} - h_{measured,-i}| \tag{26}$$
$$D_1(\vec{z}_{-i}) = G_1(\vec{z}) - G_1(\vec{z}_{-i}) \tag{27}$$

The difference reward for agent i is simply the error in the enthalpy measurement minus the error in the enthalpy measurement if agent i was sensing nothing. Note that by definition, an agent whose action is to sense nothing has a difference reward of zero. In other words, an agent which senses nothing has no effect on the system performance at all.

4.2 Enthalpy Measurement with Sensor Failures

In a distributed sensor network, there is a nonzero probability that some subset of sensors will fail due to lack of proper maintenance or manufacturing defects. It is imperative that sensor failure does not result in significant loss of performance. A set of sensors will be trained to accurately measure system enthalpy, and then some subset of those sensors will fail. The remaining sensors will then continue learning in order to compensate for the loss of sensors. This experiment will give insight into how robust the distributed sensor network is.

4.3 Enthalpy Tracking

The third experiment involves tracking enthalpy as it changes over time. Suppose that the enthalpy at the boiler output is being measured. The heat input to the boiler is generally created by a combustion process with a mixture of fuel and air [5]. During transient startup processes, the heat input to the boiler gradually increases over time, which results in the enthalpy of the working fluid at the boiler output to gradually increase over time. It is imperative that this increase in enthalpy is accurately tracked in order to properly control the system. In the case of rising temperature and pressure, different sensors will need to be on at different times, as the temperature and pressure values enter and leave effective operating ranges of specific sensors. Thus, for accurate enthalpy tracking, different sensors will need to learn to turn on or off at different times in response to changing system parameters as well as the activation levels of other sensors. For the enthalpy tracking experiment, the difference reward is the same as that in the enthalpy measurement experiment (Equation 27). For this experiment, the enthalpy profile as a function of time is given in Figure 2.

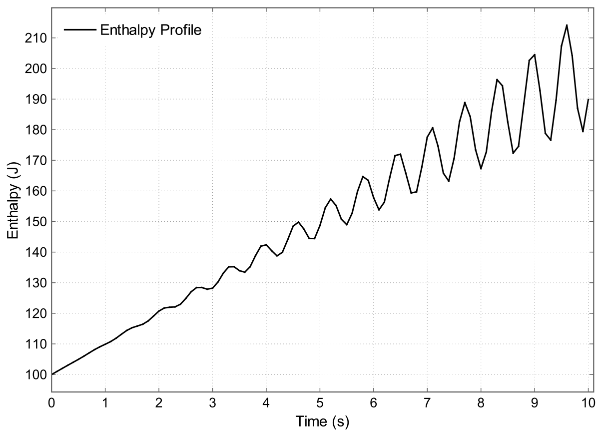

Figure 2: Enthalpy profile to be tracked. The enthalpy rises from h_{min} to h_{max} as time progresses, with oscillations occuring during the rise.

The oscillations seen in Figure 2 are artificially added, in order to make the tracking problem more complex, and would not generally be seen in an actual power plant. Tracking a profile such as that in Figure 2 has one key difficulty given the sensor model used. Each sensor has an effective temperature operating range, and this range is a subset of the range $\Delta T = T_{max} - T_{min}$. Different sensors must learn to turn on and off depending on the state of the system. As an example, consider the case where some sensors are accurate between $0^o C$ and $50^o C$, and other sensors are accurate between $50^o C$ and $100^o C$. In order to accurately track a temperature profile that ranges from $0^o C$ and $100^o C$, different sensors should be activated at different times. More generally, if the network of sensors has different sensors with n effective temperature operating ranges, each of which is a subset of the range $[T_{min}, T_{max}]$, then a policy which maps the temperature of the system to which sensors are activated is necessary for accurate tracking. The same principles hold for tracking pressure. Both temperature and pressure must be tracked accurately in order to properly track the rising enthalpy of the working fluid at the turbine outlet.

4.4 Cost Analysis

In order to show that a distributed sensor network can have performance similar to standard sensor implementations while remaining cost-competitive, we repeat the enthalpy measurement experiment while varying the value of β_1 in the sensor cost model (Equation 15). The single sensor is assigned an atteuation a_{single}, and the sensors in the distributed sensor network will be assigned an average attenuation $a_{network} > a_{single}$. The single sensor attenuation will be used to develop a cost of the sensor based on the cost model and the value of β_1. Given the value of β_1 and the attenuation of network sensors $a_{network}$, the size of the distributed sensor network will be determined such that the sensor network has an equivalent fiscal cost as the single sensor. For each value of β_1, the aggregate error of the distributed sensor network is compared to the error of the single sensor. The purpose of this experiment is to show the difference in performance between the standard sensor implementation and a distributed sensor network as a function of the sensor cost model. This will give insight as to when distributed sensor networks are financially feasible for implementations in power plants.

5. RESULTS

The following sections give the results for each experiment. For all reinforcement learning processes, the exploration parameter ϵ is set to 0.05. Each experiment was run for 50 statistical runs, and the error bars reported are the standard error in the mean $(\sigma/\sqrt{(N)})$, where σ is the sample standard deviation of the dataset and N is the number of statistical runs.

5.1 Enthalpy Measurement

For the enthalpy measurement experiment, a network of sensors was developed using the sensor cost model (Equation 15), where c_1 and β_1 were set to 2. A sensor network was created with 200 sensors that had attenuations drawn from a normal distribution with a mean of 0.0 and a standard deviation of 0.5. This network is compared with a single sensor with an attenuation drawn from a normal curve with a mean of 0.0 and a standard deviation of 0.01. Given the cost model for the sensors, the single sensor is 100 times more expensive than the network of sensors. The experiment ran for 3000 learning episodes. The aggregate attenuations for the sensor network compared with the single sensor are shown in Figure 3.

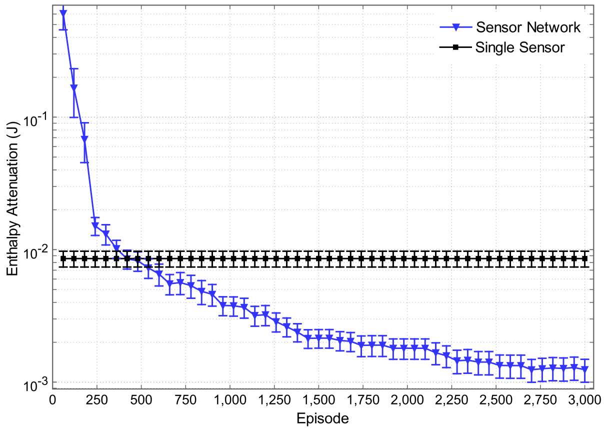

Figure 3: Comparison of a single sensor and a sensor network. After training, a network of 200 sensors with high error are more accurate than a single sensor with low error

The reason that the mean attenuation for the single sensor is not zero is that the attenuation reported is the magnitude of attenuation. In terms of performance, a sensor with an error of δ is equivalent to a sensor with an error of $-\delta$. The sensor network learned a policy with much lower error in the enthalpy reading. This means that a distributed sensor network with fairly inaccurate sensors can learn a policy such that enthalpy measurements are more accurate than a single accurate sensor, in addition to being cheaper than the single sensor. Because the error in measurement is lower for the distributed sensor network, this network is capable of detecting minute parameter changes that are not detectable by the single sensor. For example, if there was a slight leak in a pipeline that transports the working fluid, and the drop in pressure associated with that leak was smaller than the single sensor could detect, the distributed sensor network could still detect that leak. Thus, the distributed sensor network is capable of detecting problems in the system that a single sensor could not. This is especially important in the context of leaks in the fluid pipelines, because these leaks need to be detected before they grow and cause catostrophic failures in the system.

5.2 Enthalpy Measurement with Sensor Failures

The sensor failure experiment was run with the same sensor network utilized in the enthalpy measurement experiment. After 3000 training episodes, 40 of the 200 sensors failed, leaving 160 functional sensors. The sensors then retrained for 2000 episodes, and the resulting performance was determined. This experiment was repeated for 500 statistical runs, with the error in the mean being reported $(\sigma/\sqrt{N})$. The results for the sensor failure experiment are shown in Figure 4.

As seen in Figure 4, the distributed sensor network is able to regain over 99% of lost performance after 20% of the sensors fail. This demonstrates that the distributed sensor network is robust to sensor failure. This is an extremely important property for a sensor network operating in a power plant, because extreme temperatures and pressures create a harsh

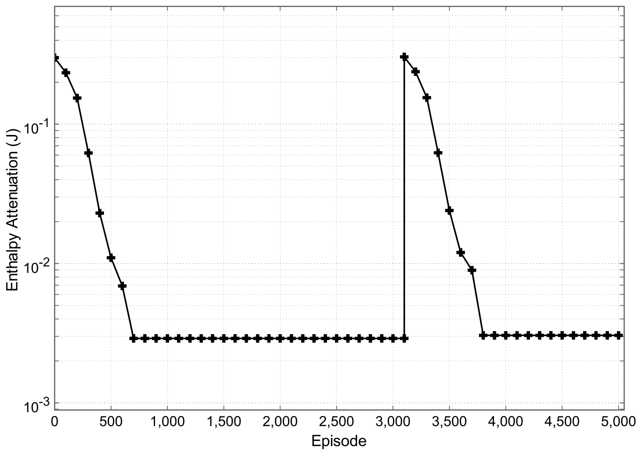

Figure 4: Sensor failure results. A system with 200 sensors is trained and 40 sensors fail at training episode 3000. over 99% of system performance is regained. The error bars are smaller than the plot symbols, and are thus omitted.

environment which is more difficult for sensors to survive in. It is important to note that in a power plant, network training would most likely occur in an offline manner, such that the plant would need to be shut down if a large portion of sensors failed.

5.3 Enthalpy Tracking

For the enthalpy tracking experiment, an oscillatory rise in enthalpy from h_0 to h_{final} was tracked. There were four sets of 200 sensors, each of which had an effective operating range of one-fourth the total change in pressure and temperature. Thus, sensors needed to learn to turn on or off depending on the system state. The sensors were allowed to train on random states ($s = \{temperature, pressure\}$) for 3000 episodes using the Q-learning policy, and then the learned policy was used to control the sensors while monitoring the temperature and pressure profiles. The actual enthalpy compared to the measured enthalpy as a function of time is shown in Figure 5.

As seen in Figure 5, the enthalpy profile is accurately tracked. The distributed sensor network learned an adequate policy which determines when each sensor should turn on and off dependent upon the temperature and pressure of the system. This is critical to proper control of a power plant, especially during transient operation. The error in measurement at each time step is also of interest. As the temperatures and pressures go from the effective operating range of one set of sensors to another, it is important that jumps in error do not occur as one set of sensors shuts off and another set of sensors turns on. The error in enthalpy measurement as a function of time is shown in Figure 6.

As seen in Figure 6, the error in the enthalpy measurement does not drastically increase at any point, indicating that the transition of one set of sensors being activated to another set being activated does not cause discontinuous jumps in enthalpy estimation, an important property with regards to system control.

5.4 Cost Analysis

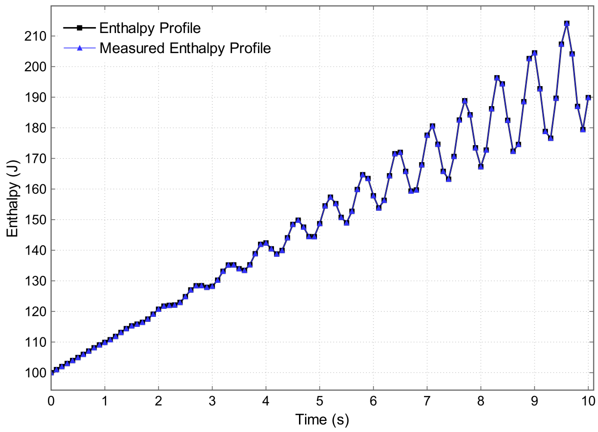

Figure 5: Temperature profile versus the sensed temperature profile. There is a negligible difference between the two profiles

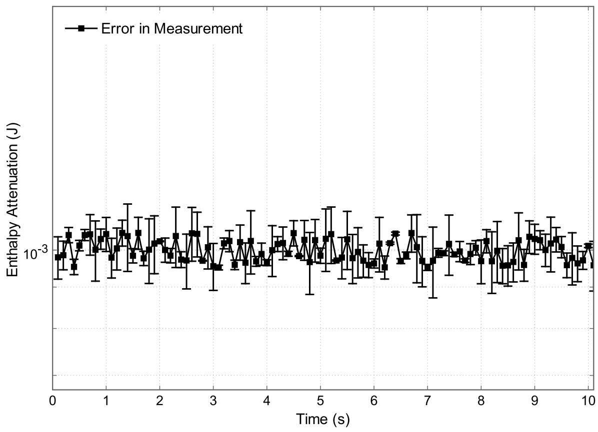

Figure 6: Sensing error as a function of time during enthalpy tracking. Although the error varies with time, it never exceeds $0.0011J$

An analysis of the cost model was completed to analyze the financial feasability of the sensor network. The variable β_1 in the cost model (Equation 15) dictates how sensor prices vary as a function of their attenuation. Given a specific attenuation for a sensor, the cost increases with β. The performance of a single sensor and an equivalent priced sensor network was tested as a function of β. The results for this analysis are shown in Figure 7.

As seen in Figure 7, as long as $\beta_1 \geq 0.621$, the equivalently priced distributed sensor network outperforms the single sensor. As β_1 continues to increase, the distributed sensor network outperforms the single sensor by a wider margin. Thus, for the sensor cost model given in Equation 15, we have developed a threshold which tells us when a distributed sensor network is a financially viable alternative to the traditional sensor implementation in a power plant.

6. DISCUSSION

This paper demonstrates that there are multiple benefits to installing distributed sensor networks to replace stan-

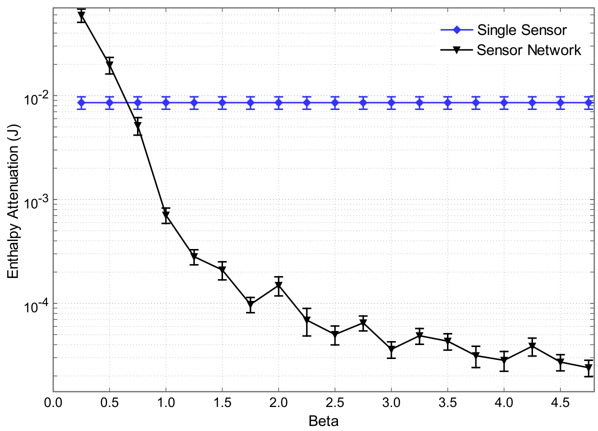

Figure 7: Performance of single sensor versus equivalently priced distributed sensor network as a function of cost variable β_1. As long as $\beta_1 \geq 0.621$, the distributed sensor network outperforms the single sensor

dard sets of sensors in power plants. First, a distributed sensor network of inexpensive and fairly inaccurate sensors can measure and track system parameters more accurately than single, expensive sensors. Because the sensor network has more accurate aggregate measurements, small changes in system parameters that go undetected by standard sensors can be detected by distributed sensor networks. Small system changes such as diminutive fluid leaks which would previously go undetected can now be detected. This is important from the standpoint of plant health, because it is desirable that these problems be detected early before they grow into larger problems. The sensor network is also robust to sensor failures, with 20% failures resulting in less than 1% loss in performance. This is extremely important in harsh environments where sensor failure probabilities are increased. Finally, a distributed sensor network is shown to be cost-competitive with standard sensor implementations in a power plant, and the parameters for determining when distributed sensor networks are financially viable were defined. This method could be extended to any sensor cost model, in order to determine the cost-effectiveness of specific types of sensors.

The experiments in this paper show that a distributed sensor network is a cost-effective and functional alternative to standard sensor implementations in power plants. By demonstrating that a distributed sensor network is worth implementing in a power plant, work can begin on specific implementations in a model power plant. More specifically, future work will involve installing a distributed sensor network across an entire power plant, in order to test how such a network may be used in conjunction with plant control. Another important area to research involves determining how a distributed sensor network could communicate such that a controller could be informed of disturbances in one plant state and compensate before the disturbances propogate through all other plant states. We also plan to analyze more complex sensor models, where attenuation is not constant, but either has a measure of randomness or changes over time. Finally, we will construct or more realistic plant model (e.g. without the adiabatic assumptions) in order to test the sensor network on a more complex and realistic domain.

7. ACKNOWLEDGEMENTS

This work was partially supported by DOE NETL grant DE-FE0000857.

8. REFERENCES

[1] D. Challet and N. F. Johnson. Optimal combinations of imperfect objects. *Physical Review Letters*, 89(2):028701, 2002.

[2] A. Farinelli and P. Scerri. Low-overhead cooperative detection of false sensor readings, 2005.

[3] C.-C. Hsu and C.-T. Su. An adaptive forecast-based chart for non-gaussian processes monitoring: With application to equipment malfunctions detection in a thermal power plant. *Control Systems Technology, IEEE Transactions on*, 19(5):1245 –1250, sept. 2011.

[4] A. Hussey, A. Nasipuri, R. Cox, and J. Sorge. Feasibility of using a wireless mesh sensor network in a coal-fired power plant. In *IEEE SoutheastCon 2010 (SoutheastCon), Proceedings of the*, pages 384 –389, march 2010.

[5] M. J. Moran and H. N. Shapiro. *Fundamentals of Engineering Thermodynamics*. John Wiley and Sons, Inc., fifth edition, 2004.

[6] A. Rogers, A. Farinelli, and N. R. Jennings. Self-organising sensors for wide area surveillance using the max-sum algorithm. *Sensors Peterborough NH*, pages 84–100, 2010.

[7] R. Swartz, J. Lynch, and C.-H. Loh. Near real-time system identification in a wireless sensor network for adaptive feedback control. In *American Control Conference, 2009. ACC '09.*, pages 3914 –3919, june 2009.

[8] K. Tumer. Designing agent utilities for coordinated, scalable and robust multiagent systems. In P. Scerri, R. Mailler, and R. Vincent, editors, *Challenges in the Coordination of Large Scale Multiagent Systems*, pages 173–188. Springer, 2005.

[9] K. Tumer and A. K. Agogino. Improving air traffic management with a learning multiagent system. *Intelligent Systems*, 24(1), Jan/Feb 2009.

[10] K. Tumer, A. K. Agogino, and Z. Welch. Traffic congestion management as a learning agent coordination problem. In A. Bazzan and F. Kluegl, editors, *Multiagent Architectures for Traffic and Transportation Engineering*, pages 261–279. Lecture notes in AI, Springer, 2009.

[11] H. Varian. *Intermediate Microeconomics*. W.W. Norton and Co., 1999.

[12] A. Venkat, I. Hiskens, J. Rawlings, and S. Wright. Distributed mpc strategies with application to power system automatic generation control. *Control Systems Technology, IEEE Transactions on*, 16(6):1192 –1206, nov. 2008.