Chapter 12

APPLICATION-LEVEL POWER AWARENESS

Jeffrey A. Barnett
Northrop Grumman Corporation
Automation Sciences Laboratory
El Segundo, California
jbarnett@nrtc.northrop.com

Abstract Optimizing resource allocation to best meet system goals is the essence of good engineering. It is the unifying principal in the design of resource-constrained systems such as aircraft, farming complexes, and even military forces. Well-designed systems accomplish tasks using minimal resources. They also dynamically adapt their methods and goals to best use available resources. This is as evident in artificial systems as it is in natural systems. Below, two examples one a design time problem and the other a dynamic allocation problem–are developed to elucidate the sort of engineering that must be applied to resource management. The first example is a design tradeoff between the efficient use of electric power and the performance of an aircraft. The second example is dynamic allocation of a limited energy budget between cooperating sensors to maximize the quality of the combined measurement.

Keywords: Power aware design, dynamic energy management

Introduction

The most successful of nature's systems are those that have adaptations to optimally use available resources. The form factors of birds and fish, for example, are both optimized for efficient motion. Since drag, density, and buoyancy in the air and in water are different, their shapes are quite different. Thus, we observe that design optimizations are specific to environments.

Dynamic adaptation of behavior is another aspect of energy optimization. Ordinarily, animals select prey that maximize the expected return over their investment to catch it. However, in times of great need, minimal margins are readily accepted and in times of plenty, indifferent choices are made. Plants adapt their processes to such factors as the availability of sunlight, moisture,

In B. Melham and R. Graybill (eds.) *Power Aware Computing* Kluwer (2002)

and nutrients. Even the reproductive decisions of some species are affected by environmental cues and the state of the individual at the time. Nature has indeed developed energy-aware systems that in both design and behavior exhibit resource awareness.

Current events highlight the need for energy-aware artificial systems too. Our petroleum resources are rapidly dissipating and utility companies cannot produce enough electricity to meet all of our demands all of the time. If these were the only drivers, conservation remedies, new energy sources, and normal engineering progress would hold the problems in check.

However technology, particularly the modern computer, promises new types of systems for applications that will operate in severely energy-constrained environments. Extended duration space exploration is one well-known example. Another is the construction and deployment of unattended sensor networks comprising 1000's of miniaturized elements. Each element must be put in place with sufficient power, e.g., from a battery, to last its entire lifetime and support sensing, computing, and communications. Elements, in addition, may scavenge energy from sunlight, natural vibrations, or even chemical interactions with the environment. Energy management and optimization clearly will be of overriding importance in the design of these systems as will the need to dynamically adapt their behavior by considering the urgency of goals and the available resources.

Below, two examples of power-aware applications are described. The first is aircraft design where energy conservation contributes to overall performance by reducing battery and generator weight to enable longer flight times. Section 1 develops the mathematical relationship linking power efficiency to performance. That relationship is used to examine tradeoffs in terms of two different aircraft. Section 2 considers an application where multiple sensors develop joint measurements while sharing an energy budget. The goal is to minimize the variance of the resulting measurements while using the least energy possible and still meet system objectives. Thus, the second example is dynamic decision-making in the management of resources.

Section 3 summarizes the results and relates them to similar observations in other fields. The goal of this paper is not to present the development of specific new technology. Rather, it is too provide examples of how engineers can reason about and construct application-specific systems that will operate in energy-constrained environments.

1. A Design-Time Optimization for Aircraft

The value of power awareness and energy reduction technologies are determined by the specifics of an application domain and the milieu in which its systems will operate. Intelligent power management is one of the key technology enablers to successful deployment of deep space probes, networks comprised of miniature battery-powered sensors, and the Land Warrior. Below, a simple relation between energy efficiency and maximum flight time (endurance) for aircraft is derived. That relation is used to model the value of efficient design as measured by increased endurance: Sometimes it is very valuable and sometimes it is not. The point of this exercise is not to solve aircraft problems per se. Rather, it is to introduce the sort of domain-specific analysis that is necessary to make design-time tradeoffs for an application.

1.1 Domain Considerations

Weight is the enemy of air vehicle performance. Heaviness directly decreases maneuverability, handling, range, and endurance, where endurance is maximum flight time. Since savings in the use of electrical energy will reduce the weight of the generator, the battery, or both, power awareness will increase the missions that a given aircraft can undertake and complete. Another benefit of weight reduction is that the weight budget for the fuel used to power the air vehicle can be increased when the battery or generator weight is decreased.

Reducing electric power consumption also decreases heat production. Therefore, more equipment can be packed into the same or smaller form factor. Since decreasing size increases maneuverability, range, and endurance, we have an additional argument for power awareness. Further, obvious benefits accrue when power-aware avionics can compensate for unusual situations such as generator loss. This example makes the point that power awareness is more than energy minimization: it's about doing the best you can with the resources that are actually available.

Concerns about the effects of electric energy consumption on aircraft performance are not new. For example, space and heat considerations are discussed in [2] along with their impacts on operational costs. Power awareness is even more important in the modern world. One example is the Helios, a product of AeroVironment Corporation, with the goal of indefinite sustained flight powered by scavenging solar energy. Other extreme examples are smart munitions that either have no propulsion system or only use propulsion in burst modes. Since there is no engine-powered generator, all electric energy must be supplied by batteries to support sensors, computation, communications, and control surface actuators. Batteries are a limiting resource that produce unwanted heat, consume precious space, and add weight.

The standard relation approximating the endurance achievable by an aircraft as a function of its weight and the amount of fuel it carries is derived next. The following discussion extends that relation to include the effects of electric power utilization efficiency, hence battery weight, on endurance. Finally, that relation is used to analyze the effects of power efficiency on endurance for two example aircraft.

1.2 Endurance is a Function of Weight and Fuel

The relation between endurance and an air vehicle's weight and the amount of fuel it carries is derived here. A reasonable first-order approximation is used to simplify the mathematics: instantaneous fuel consumption rate is linearly related to total vehicle weight for a constant velocity and altitude. The next section extends the relation to consider the effects of electric energy conservation on weight, hence, its relation to endurance. See [7] for justification of the assumptions made here as well as derivations of more detailed relationships when other variables are considered.

Let $f = f(t)$ be the amount of fuel remaining at flight time t and let $f_0 = f(0)$ be the amount of fuel loaded on an aircraft that weights W_d dry (fuel weight is zero). The assumption is that the amount of fuel, measured by weight, necessary per second of flight time per unit of aircraft weight is a constant. Thus,

$$f' = -c(W_d + f)$$

where $f' = df/dt$ and $c > 0$ is the fuel consumption constant. The solution to this simple ode, subject to the constraint that $f(0) = f_0$, is

$$f(t) = W_0 e^{-ct} - W_d,$$

where $W_0 = W_d + f_0$ is the initial weight of the aircraft at $t = 0$.

The endurance, E, is defined to be the maximum flight time. Clearly, $f(E) = 0$ is the condition to find that unique value of E. Therefore,

$$E = \frac{1}{c} \times \log\left(\frac{W_0}{W_d}\right). \tag{12.1}$$

An explanation, in part, for this, perhaps unexpected result, is that in order to increase endurance, more fuel is obviously needed. But, additional fuel is necessary to carry the original extra fuel, etc. Thus, weight has a nonlinear effect on endurance. Note, this derivation is a reasonable approximation for many land as well as airborne vehicles.

1.3 Endurance is a Function of Energy Conservation

In order to extend the above derivation to account for a battery and, hence, to account for the efficiency with which that battery is used, a few details must be added to our simple model. Let W_S be the structural weight of the aircraft and W_B be the weight of its battery. Thus,

$$\begin{aligned} W_d &= W_S + W_B \\ W_0 &= W_S + W_B + f_0. \end{aligned}$$

In order to focus on the relevant tradeoff for this discussion—battery weight versus fuel—define $L = W_B + f_0$ as the total weight budget for fuel plus battery and β as the battery weight needed for one unit of mission time. Thus, $W_B = \beta E$. β is interpreted as the efficiency of the *energy consumers* not of the *battery*, though the analyses are practically identical in either case. Now,

$$\begin{aligned} W_d &= W_S + \beta E \\ W_0 &= W_S + L, \end{aligned}$$

where the first equation assumes that there is just enough battery to last for the whole mission. In other words, L is split between fuel and battery weight so that both resources expire simultaneously. There is one important difference between fuel and battery that makes these equations non-symmetric: when fuel is consumed, its weight decreases but no such reduction will be observed for batteries. Substituting the formulas into (12.1) yields

$$E = \frac{1}{c} \times \log\left(\frac{W_0}{W_S + \beta E}\right).$$

What we desire is an expression for E as a function of β so that the payoff of efficient energy utilization can be measured in terms of endurance enhancement. A quick inspection shows that an elementary expression isn't available. However, it is straightforward to express β as a function of E:

$$\beta = \frac{W_0 - W_S \exp(cE)}{E \exp(cE)}. \tag{12.2}$$

Clearly, the value of E must be positive and is bounded from above by the case where $f_0 = L$ and, hence, $\beta = 0$. Thus,

$$0 < E < \frac{\log(W_0) - \log(W_s)}{c}. \tag{12.3}$$

Equation 12.2 and the bounds on E are used below to numerically generate and graph (β, E) pairs for some examples.

1.4 Aircraft Examples & Analysis

Two simple examples of vehicles designed for quite different mission profiles are analyzed in this section. The first, called Explorer, is a simple unmanned sensor craft that uses a generator to support most applications. It flies

Table 12.1. Parameters defining two aircraft examples.

Example	c	W_S	L	W_0	$E_{\max}$
Explorer	.00155/min	6,900 lb	3,100 lb	10,000 lb	240 min
Missile	.05108/min	15 lb	10 lb	25 lb	10 min

missions of a few hours duration and its battery is primarily for power during emergencies, e.g., for flight recorders or engine restarts after flame outs. The second example, called Missile, is a smart munition that uses batteries to power all onboard electronics. Fuel is only used in burst mode to regain altitude so as to increase target seek time. Its mission times are of the order of a few minutes. Table 12.1 summarizes the relevant parameters for these two hypothetical examples and $E_{\max} = \log(W_0/W_s)/c$ is the maximum endurance from (12.3).

E versus β pairs have been generated using (12.2) and are displayed in Figures 12.1a and 12.1b. Endurance, E, clearly increases as electric power utilization, as measured by β, decreases. There is no way to tell, at this point in the

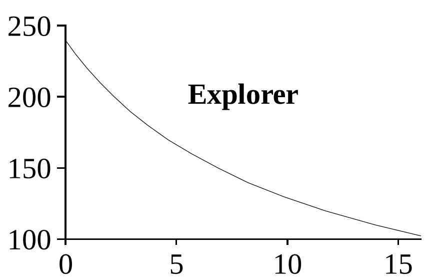

Figure 12.1a. E versus β relation for the Explorer example.

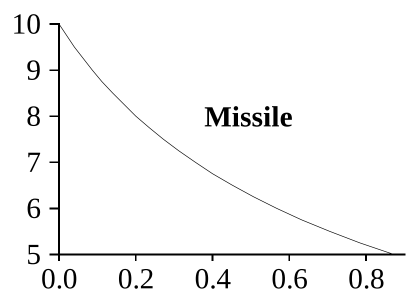

Figure 12.1b. E versus β relation for the Missile example.

development of the two examples, whether power-saving measures would be worthwhile or not. We need to know the current β, the merit of a ΔE measured as an increased fraction of successful missions, and the cost of developing and inserting power-aware technology that can provide the ΔE in order to make that determination.

Such details are beyond the scope of this paper though a simple heuristic analysis is presented. Figures 12.2a and 12.2b suggest the shapes of typical tradeoff curves: $p = p(E)$ is the probability of a successful mission if endurance is E and $C = C(\Delta E, E)$ is the cost to achieve an incremental endurance enhancement given the current value of E. Design tradeoffs are then formed in terms of this and other relevant information.

The Explorer example posits that batteries are not the main source of electric power and it is fair to assume that its designers made various decisions to

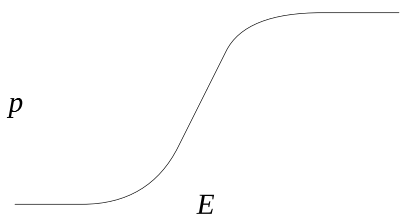

Figure 12.2a. Probability of mission success as a function of endurance.

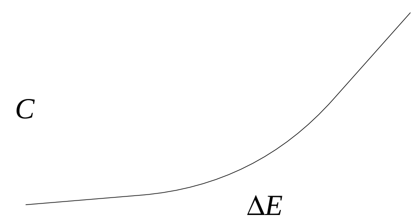

Figure 12.2b. Cost to increase endurance as a function of the desired increment.

maximize endurance. In other words, it is unlikely that the total weight of the generator and battery are a significant fraction of the total vehicle weight. Further, $f_0 \approx L$. Therefore, power-awareness, even if it could reduce the generator weight cannot increase endurance by more than a few minutes.

The Missile example is completely different. Battery weight is a significant fraction of L and L is a significant fraction of W_0. Once again, conservation can only increase endurance by a few minutes but a few minutes, compared to a current endurance of less than 10 minutes, could make all the difference in whether or not missions are successful.

1.5 Beyond Design-Time Optimization

The results and examples developed above are meant to show how considerations of power utilization can impact overall system performance. In particular, the effects of energy conservation on maximum flight time are demonstrated. Energy conservation, however, is only one aspect of power awareness. The more general problem is to dynamically determine how to optimize mission goal achievement as a function of the resources actually available. This is a much more difficult problem and is even more of an application-dependent issue than is straightforward energy minimization.

The Missile example suggest a plethora of interesting power management problems. For example, given a projected time to target impact maximize the probability of a successful strike. The problem is how best to use the remaining battery energy during the known flight time. Should sensor resource allocations be increased or should communications and computation be used to receive and fuse an ally's sensor data with onboard information? The answer depends on the resources available as well as the relative quality and costs of the two solutions. In other words, real power awareness depends on the ability to make dynamic, situation-dependent as well as application-dependent decisions. The sort of design-time tradeoffs analyzed herein are important too, but they are only a first step toward realization of fully power-aware systems.

2. Dynamic Energy Allocation for Cooperating Sensors

The problem considered in this section is how to split a limited energy budget between several sensors whose measurements will be fused. The objective is to minimize the variance of that joint measurement. A simple fusion model is defined, measurement variance is related to energy expenditures, and a criterion for an optimal allocation is derived. An extended example of a dynamic allocation problem and its solution is presented in three parts: (1) a specific model of sensor variance as a function of energy allocation is introduced, (2) the optimality problem for two sensors described by that particular model is solved, and (3) a numerical example is used to illustrate an optimal policy and the sort of decision-making that it engenders. This note is meant as an example of how one might formulate and solve dynamic energy management problems and not as a definitive practical result.

2.1 Fusing Sensor Measurements

If sensors $1,\ldots,n$ make measurements $m_1,\ldots,m_n$ that are statistically independent, the joint estimate m is taken to be

$$m = \frac{w_1 m_1 + \cdots + w_n m_n}{w_1 + \cdots + w_n}, \tag{12.4}$$

where $w_i = v_1 \cdots v_n / v_i$ and v_i is the variance of measurement m_i. The variance, v, of the joint estimate is

$$v = \frac{1}{\frac{1}{v_1} + \cdots + \frac{1}{v_n}}. \tag{12.5}$$

The function, $v = v(v_1,\ldots,v_n)$, has the following intuitive properties:

Less is better: $\dfrac{\partial v}{\partial v_i} \geq 0$.

Progress: $v(v_1,\ldots v_n) \leq v_i$.

As good as it gets: $v(v_1,\ldots,v_n) = 0$ if any $v_i = 0$, a corollary of progress.

No pain, no gain: $v(v_1,\ldots,v_n) = v(v_1,\ldots,v_{i-1},v_{i+1},\ldots,v_n)$ if $v_i = \infty$.

Unit independence: $v(rv_1,\ldots,rv_n) = rv(v_1,\ldots,v_n) = 0$.

Some justification is provided in [6] for the formulas assumed here as well as an introduction to the general topic of fusion using Kalman filters. Note, fusion formulas such as (12.4) and (12.5) are more complicated if the sensor measurements are statistically dependent or if multiple features are measured as is often the case.

2.2 Energy Allocation Determines Sensor Performance

Modern sensors offer many controls that trade energy for performance. Examples are raw power, pulse width, number of pulses per measurement, frequency of measurement, and even spectrum choices for hyperspectral sensors. Some passive sensors can vary the number of pixels or the parts of the spectrum that are processed. The more energy used, the higher the quality of the measurement. Below, quality is grossly summarized by the resulting variance of the measurement. Specific issues such as precision, etc., are assumed to be sufficiently captured by the variance estimate.

Let $v_s(e)$ represent the variance expected when sensor s makes a measurement using energy $e \geq 0$. Such a function should satisfy these criteria:

Positivity: $v_s(e) > 0$.

More is better: $\dfrac{dv_s}{de} \leq 0$.

No free lunch: $v_s(0) = \infty$.

The astute reader may be troubled by the *no free lunch* property. Isn't there always some *a priori* information available so that the initial variance would be finite? For example, an angle measurement is always in the interval $[0, 2\pi)$. The way to analyze the situation is to assume that the variance of the *a priori* estimate is v_0 and to combine it with $v_1, \ldots, v_n$ to get the combined variance, $v = (v_0^{-1} + v_1^{-1} + \cdots + v_n^{-1})^{-1}$. Since the optimality criterion derived in the next section is the same whether v_0 is accounted for or not, the math is simplified by ignoring the presence of *a priori* estimates.

2.3 Minimizing Variance Through Energy Allocation

Given a total energy budget, e, to be used to obtain independent measurements $m_1, \ldots, m_n$, the obvious objective is to divide it so that the resulting variance is minimized. Assume that $e_i \geq 0$ is used for the measurement made by sensor i. The problem is to minimize (12.5) subject to the constraint that $e_1 + \cdots + e_n = e$. Simply use Lagrangian multipliers as follows:

$$v(e_1, \ldots, e_n, \lambda) = \frac{1}{\frac{1}{v_1} + \cdots + \frac{1}{v_n}} + \lambda \left(e - \sum e_i\right).$$

Necessary minimization criteria follow from $dv/de_i = 0$:

$$\begin{aligned} 0 &= \frac{dv}{de_i} \\ &= \frac{-v_i'/v_i^2}{\left(\frac{1}{v_1} + \cdots + \frac{1}{v_n}\right)^2} - \lambda \\ \frac{v_i'}{v_i^2} &= -\lambda \left(\frac{1}{v_1} + \cdots + \frac{1}{v_n}\right)^2, \end{aligned}$$

where $v_i' = dv_i/de_i$. Since λ does not depend on i, the entire right hand side is independent of i. Therefore,

$$\frac{v_i'}{v_i^2} = \frac{v_j'}{v_j^2}, \tag{12.6}$$

for all $1 \le i, j \le n$. Standard dynamic programming techniques can be used to develop optimal allocations when there are more than two sensors.

2.4 Applying the Theory

This section develops an example of variance minimization through energy allocation. The problem of achieving a given result quality with minimum resource investment is addressed in the following section. Both problems are examples of dynamic energy management. The discussion proceeds in three parts: (1) a specific, parameterized sensor model that defines variance as a function of energy is introduced, (2) the two-sensor minimization problem is solved, and (3) the optimal allocation policy and the variance achieved are analyzed for particular sensor instances.

2.4.1 A Parameterized Sensor Model. The form of v_i assumed in the optimization example developed in the rest of this section is

$$v_i(e) = a_i + b_i/e, \tag{12.7}$$

where $a_i, b_i > 0$. Such forms obviously possess the properties required above for variance functions. For future reference, note that

$$\frac{dv_i}{de} = \frac{-b_i}{e^2}. \tag{12.8}$$

The parameters a_i and b_i have the following interpretations in this formulation: a_i is the limit of the sensor's operation, i.e., $a_i = v_i(\infty)$, and b_i is the difference between the sensor's variance using one unit of energy and its best possible performance, i.e., $b_i = v_i(1) - v_i(\infty)$.

2.4.2 Solving the Two-Sensor Problem. Assume two independent sensors whose behaviors are defined by (12.7) and a total energy budget of e. Let energy $0 \le x_1 \le e$ be allocated to the first sensor and $x_2 = e - x_1$ be allocated to the second. Then from (12.7) and (12.8),

$$\begin{aligned} v_1(x_1) &= a_1 + b_1/x_1 \\ v_2(x_2) &= a_2 + b_2/(e-x_1) \\ v_1'(x_1) &= -b_1/x_1^2 \\ v_2'(x_2) &= -b_2/(e-x_1)^2. \end{aligned}$$

Now use criterion (12.6) to find the optimal allocation as follows:

$$\begin{aligned} \frac{v_1'}{v_1^2} &= \frac{v_2'}{v_2^2} \\ \frac{-b_1/x_1^2}{(a_1+b_1/x_1)^2} &= \frac{-b_2/(e-x_1)^2}{(a_2+b_2/(e-x_1))^2} \\ \frac{b_1}{(a_1x_1+b_1)^2} &= \frac{b_2}{(a_2(e-x_1)+b_2)^2} \\ \sqrt{b_1}\,(a_2(e-x_1)+b_2) &= \sqrt{b_2}\,(a_1x_1+b_1). \end{aligned}$$

The solutions to this linear equation for the optimal x_1^* and $x_2^* = e - x_1^*$ are

$$x_1^*(e) = \frac{a_2\sqrt{b_1}e + b_2\sqrt{b_1} - b_1\sqrt{b_2}}{a_1\sqrt{b_2} + a_2\sqrt{b_1}} \tag{12.9a}$$

$$x_2^*(e) = \frac{a_1\sqrt{b_2}e + b_1\sqrt{b_2} - b_2\sqrt{b_1}}{a_1\sqrt{b_2} + a_2\sqrt{b_1}}. \tag{12.9b}$$

Since x_1^* and x_2^* must be nonnegative, the numerators must be nonnegative. Thus, there are two validity constraints entailed by (12.9a) and (12.9b):

$$\begin{aligned} e &> \frac{\sqrt{b_1b_2} - b_2}{a_2} \\ e &> \frac{\sqrt{b_1b_2} - b_1}{a_1} \end{aligned}$$

If the first constraint is violated, the optimal allocation is $x_1^* = 0$ and $x_2^* = e$. If the second is violated, the optimal allocation is $x_1^* = e$ and $x_2^* = 0$. In the first case the resultant minimum fused variance, v^*, is $v^*(e) = a_2 + b_2/e$ and in the second the minimum variance is $v^*(e) = a_1 + b_1/e$. When there is an "interior" solution, v_1^* and v_2^* are calculated by substituting (12.9a) and (12.9b) into (12.7) to get

$$v_1^*(x_1) = \frac{a_1 a_2 e + a_1 b_2 + a_2 b_1}{a_2 e + b_2 - \sqrt{b_1 b_2}}$$
$$v_2^*(x_2) = \frac{a_1 a_2 e + a_1 b_2 + a_2 b_1}{a_1 e + b_1 - \sqrt{b_1 b_2}}.$$

These values are substituted into (12.5) and simplified to calculate the resultant minimum total variance as

$$v^*(e) = \frac{a_1 a_2 e + a_1 b_2 + a_2 b_1}{(a_1 + a_2)e + (\sqrt{b_1} - \sqrt{b_2})^2}.$$

Table 12.2 summarizes these results for the optimal allocation policy: $x_1^*(e)$ is

Table 12.2. Formulas for optimal energy allocations and resulting variances.

	$0 \le e < \frac{\sqrt{b_1 b_2} - b_2}{a_2}$	$0 \le e < \frac{\sqrt{b_1 b_2} - b_1}{a_1}$	*Otherwise*
$x_1^*(e) =$	0	e	$\frac{a_2\sqrt{b_1}e + b_2\sqrt{b_1} - b_1\sqrt{b_2}}{a_1\sqrt{b_2} + a_2\sqrt{b_1}}$
$v_1^*(e) =$	∞	$a_1 + b_1/e$	$\frac{a_1 a_2 e + a_1 b_2 + a_2 b_1}{a_2 e + b_2 - \sqrt{b_1 b_2}}$
$x_2^*(e) =$	e	0	$\frac{a_1\sqrt{b_2}e + b_1\sqrt{b_2} - b_2\sqrt{b_1}}{a_1\sqrt{b_2} + a_2\sqrt{b_1}}$
$v_2^*(e) =$	$a_2 + b_2/e$	∞	$\frac{a_1 a_2 e + a_1 b_2 + a_2 b_1}{a_1 e + b_1 - \sqrt{b_1 b_2}}$
$v^*(e) =$	$v_2^*(e)$	$v_1^*(e)$	$\frac{a_1 a_2 e + a_1 b_2 + a_2 b_1}{(a_1 + a_2)e + (\sqrt{b_1} - \sqrt{b_2})^2}$

the optimal allocation of total energy, e, to the first sensor and $v_1^*(e) = v_1(x_1^*(e))$ is the variance it achieves. The functions x_2^* and v_2^* are similarly defined for the second sensor. Finally, $v^*(e) = v(v_1^*(e), v_2^*(e))$ is the minimized variance of the fused measures given total energy budget e.

2.4.3 Applying the Results. This section develops a simple numerical example with two sensors. The resulting allocation policy and variance as a function of energy curves exhibit some interesting behaviors. Formulas for the optimal allocation policy and the resulting variances, derived from Table 12.2 using the constants $a_1 = 4$, $b_1 = 8$, $a_2 = 3$, and $b_2 = 18$, are shown in Figure 12.3a.

The optimal allocation of energy to the sensors is depicted in Figure 12.3b. Initially when $0 \le e \le 1 = (\sqrt{b_1 b_2} - b_1)/a_1$, all of the available energy is

$$x_1^*(e) = \begin{cases} e & e \leq 1 \\ \frac{e+2}{3} & e > 1 \end{cases} \qquad v_1^*(e) = \begin{cases} 4+8/e & e \leq 1 \\ \frac{4e+32}{e+2} & e > 1 \end{cases}$$

$$x_2^*(e) = \begin{cases} 0 & e \leq 1 \\ \frac{2e-2}{3} & e > 1 \end{cases} \qquad v_2^*(e) = \begin{cases} \infty & e \leq 1 \\ \frac{3e+24}{e-1} & e > 1 \end{cases}$$

$$v^*(e) = \begin{cases} 4+8/e & e \leq 1 \\ \frac{12e+96}{7e+2} & e > 1 \end{cases}$$

Figure 12.3a. Optimal allocation and resulting variances for the example where $a_1 = 4$, $b_1 = 8$, $a_2 = 3$, and $b_2 = 18$.

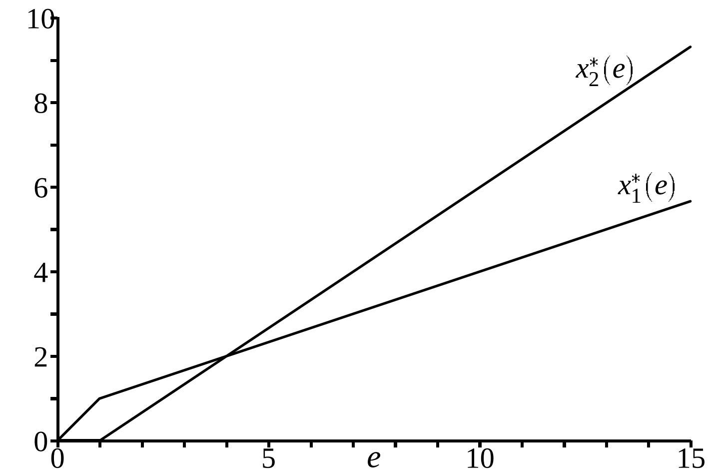

Figure 12.3b. Optimal resource allocation for the example.

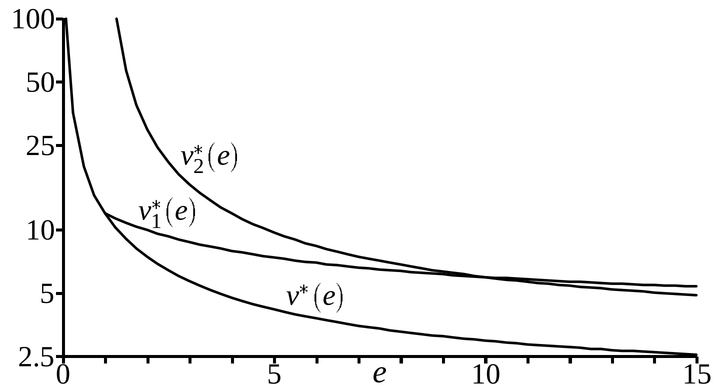

Figure 12.3c. Optimal variances for the example.

allocated to the first sensor; when $e > 1$, it is split between the two sensors. The first sensor continues to receives a larger allocation, i.e, $x_1^*(e) > x_2^*(e)$, while $0 < e < 4$ but there is a role reversal when $e > 4$ where $x_1^*(e) < x_2^*(e)$. When $e = 4$, $x_1^*(e) = x_2^*(e) = 2$.

The variances achieved using the optimal allocation policy are shown in Figure 12.3c where a log-valued vertical distance scale has been employed to improve visualization. The first sensor gives better results (has a lower variance) when a relatively small amount of energy is available but its limiting behavior is not as good as the second sensor. The variance of the fused measure, v^*, coincides with v_1^* when $0 \leq e \leq 1$. From then on it is lower. The contribution from the first sensor has lower variance than the second, i.e., $v_1^*(e) < v_2^*(e)$, while $0 < e < 10$. However, $v_1^*(10) = v_2^*(10) = 6$ and the second sensor's contribution finally dominates when $e > 10$.

This example was deliberately contrived to exhibit crossovers in both the optimal resource allocation policy and the quality of the results as shown in Figures 12.3b and 12.3c. The purpose of this ploy was two fold: (1) to indicate that even simple models can lead to complex or unexpected behavior and (2) to make the point that straightforward analytic methods can usually derive or approximate the actual optimal policy to be applied dynamically.

2.5 Related Problems

The area of dynamic sensor management and scheduling comprises many related problems. For example, the problem may be to develop an estimate that is good enough, i.e., to use the least energy to achieve a specific variance. This problem can be solved using formulas such as those in Table 12.2. The necessary minimum energy, e^*, is found from the inverse of v^*. This is straightforward because v^* is a necessarily monotonic function. The appropriate allocations are then calculated as the values of $x_1^*(e^*)$ and $x_2^*(e^*)$.

In a more realistic scenario, the analyst would need to account for the fact that measurements are not statistically independent, as assumed above, consider additional tradeoffs, and deal with measurement sources whose analytic descriptions are different than one another, e.g., it might be more economical to use communications to obtain an ally's measurement and fuse it than to use some or all of the local sensors.

Problems related to those above are discussed in [4] and [5] . Since some sensors can be time multiplexed amongst several applications, more applications can be serviced if the fraction of the sensor resources needed for each is reduced. A combination of filtering (fusion) and modelling is used. The models, formed from filter output parameters, predict target behavior between sensor measurements. The models also predict the variance increase as time of model use increases. That information is used dynamically to minimize the product of measurement duration and update rate for each application while maintaining track quality. Such techniques could be adapted by power-aware sensor schedulers. A benchmark facility to compare algorithms of this sort is avaible in [1].

Dynamic reasoning about the use of scarce resources is becoming more important as time goes by. It is an area that will pose many interesting science and engineering challenges in the future.

3. Afterword

Much of our understanding of the world around us is gained through hypotheses about how systems manage energy. Physicists assume least energy principals to form laws about the universe in the large and interactions of subatomic particles in the small. Biologists explain life, in part, by showing how organisms scavenge energy from the environment and store it for their own use at later times.

So concerns about power-aware systems are certainly not new. What is different today is that engineers are proposing and building systems whose functioning is all about energy management. Deep-space exploration and large unattended miniature sensor networks are two examples mentioned above. In these and other systems of the same ilk, power awareness is a primary design and operational principal, not just another support technology.

It may surprise the reader to know that Sigmund Freud proposed an energy minimization principal, as part of a cognitive economy model, to partially explain the workings of the human mind. He described dreaming and humor processes through a minimization of a quantity called *psychical energy*. I will close with a quote [3] that seems particularly germane to our current topic:

> *"I may perhaps venture on a comparison between psychical economy and a business enterprise. So long as the turnover in the business is very small, the important thing is that outlay in general shall be kept low and administrative costs restricted to the minimum. Economy is concerned with the absolute height of expenditure. Later, when the business has expanded, the importance of the administrative cost diminishes; the height reached by the amount of expenditure is no longer significant provided that the turnover and profits can be sufficiently increased. It would be niggling, and indeed positively detrimental, to be conservative over expenditures on the administration of business. Nevertheless, it would be wrong to assume that when expenditure was absolutely great there would be no room left for the tendency to economy."*

Acknowledgments

This effort is sponsored by Defense Advanced Research Projects Agency (DARPA) through the Air Force Research Laboratory, USAF, under agreement number F30602-00-1-0511. The opinions expressed are those of the author and do not necessarily reflect the opinions or conclusions of any other individual or agency.

References

[1] W. D. Blair, G. A. Watson, T. Kirubarajan, and Y. Bar-Shalom, Benchmark for radar allocation and tracking in ECM, *IEEE Transactions on Aerospace and Electronic Systems*, 34 (4) (1998).

[2] W. W. Davis, *Cargo Aircraft*, Pitman Publishing, 1946.

[3] S. Freud, *Jokes and their Relation to the Unconscious*, Leipzig and Vienna: Deuticke 1905. J. Strachey, translator, Norton Library, 1963.

[4] S. M. Hong and Y. H. Jung, Optimal scheduling of track updates in phased array radars, *IEEE Transactions on Aerospace and Electronic Systems*, 34 (3) (1998).

[5] T. Kirubarajan, Y. Bar-Shalom, W. D. Blair, and G. A. Watson, Management and tracking benchmark with ECM, *IEEE Transactions on Aerospace and Electronic Systems*, 34 (4) (1998).

[6] R. M. du Plessis, *Poor Man's Explanation of Kalman Filtering or How I Stopped Worrying and Learned to Love Matrix Inversion*, Rockwell International, 1967. (Reprinted by Taygeta Scientific Incorporated, 1996; can be purchased at `http://www.taygeta.com`.)

[7] D. P. Raymer, *Aircraft Design: A Conceptual Approach*, AIAA Educational Series, (Ed.) J. S. Przemieniecki, AIAA Press, 1992.