Quiz 2: Take Home Quiz on Algorithm Design
Due: Wednesday, April 4, 2007

I. MINIMUM ENERGY DATA FUSION (50 POINTS)

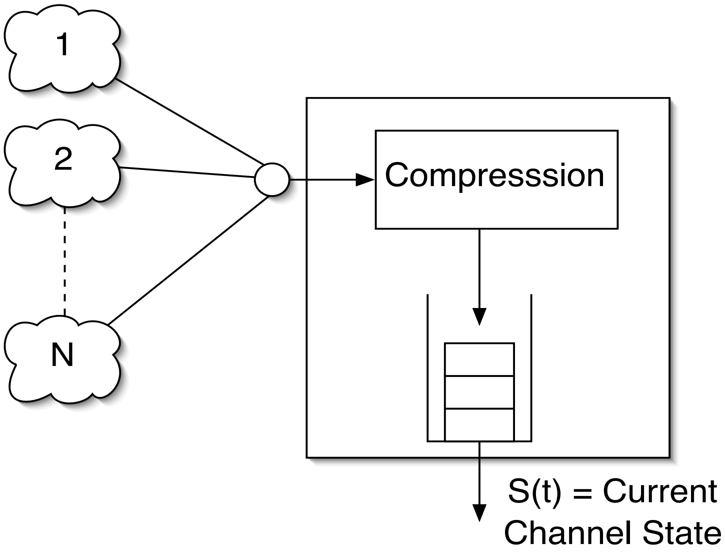

Fig. 1. Multiple sensors sending data to a single wireless link.

Consider a wireless sensing unit operating in slotted time. There are N different sensors that send data to the system every timeslot. If an individual sensor sends data during a timeslot, this data is in the form of a fixed length packet of size b bits, containing sensed information. Let $A(t)$ represent the number of sensors that send packets during slot t, so that $A(t) \in \{0, 1, \ldots, N\}$. These $A(t)$ packets may be *correllated*, and hence it may be possible to compress the information within the $A(t)$ packets (consisting of $A(t)b$ bits) into a smaller data unit for transmission over the wireless link. This is done via a *compression function* $\Psi(a, k)$ defined as follows. There are $K+1$ compression options, options $\{0, 1, \ldots, K\}$, where option 0 represents no attempted compression, and options $\{1, 2, \ldots, K\}$ represent more and more powerful attempts to compress the data. The function $\Psi(a, k)$ takes input $a \in \{0, 1, \ldots, N\}$ (representing the number of packets to be compressed) and compression option $k \in \{0, 1, \ldots, K\}$, and generates a *random variable* output R, representing the total size of the data after compression. The average compressed output $m(a, k)$ and the power expenditure ϕ_k associated with compression option $k \in \{0, 1, \ldots, K\}$ are shown in the table below. It is assumed that the values $\{m(a, k)\}$ and $\{\phi_k\}$ are known.

k	$\Psi(a,k)$	$\mathbb{E}\{\Psi(a,k) \mid a,k\}$	$P_{compress}$
0	ab	ab	$\phi_0 = 0$
1	Random	$m(a,1)$	ϕ_1
2	Random	$m(a,2)$	ϕ_2
$\cdots$	$\cdots$	$\cdots$	$\cdots$
K	Random	$m(a,K)$	ϕ_K

Note that the random outcome of $\Psi(a, k)$ satisfies:

$$\Psi(0, k) = 0 \qquad \text{for all } k \in \{0, 1, \ldots, K\}$$
$$b \leq \Psi(a, k) \leq ab \qquad \text{for } a \in \{1, \ldots, N\} \text{ and all } k \in \{0, 1, \ldots, K\}$$

Also note that power expenditure for compression is non-decreasing in the level of compression, so that:

$$0 = \phi_0 \leq \phi_1 \leq \phi_2 \leq \ldots \leq \phi_K$$

Let $k(t)$ represent the compression decision on slot t. The resulting queue backlog is thus given by:

$$U(t+1) = \max[U(t) - \mu(t), 0] + R(t) \tag{1}$$

where $R(t) = \Psi(A(t), k(t))$, and where the transmission rate $\mu(t)$ is given by a rate-power function:

$$\mu(t) = C(P_{tran}(t), S(t))$$

where $S(t)$ is the current channel state, and where $P_{tran}(t)$ is the power allocated for data transmission on slot t. Thus, every timeslot t, the number of new packets $A(t)$ is observed and a compression decision $k(t) \in \{0, 1, \ldots, K\}$ is chosen (expending power $P_{compress}(t) = \phi_{k(t)}$). Additionally, every timeslot the channel state $S(t)$ is observed and the transmission power allocation decision $P_{tran}(t)$ is made subject to $0 \leq P_{tran}(t) \leq P_{max}$. The total power expenditure on slot t is thus:

$$P(t) = P_{compress}(t) + P_{tran}(t)$$

The goal is to stabilize the system (and thus deliver all sensed information) while minimizing the total time average power expenditure.

Suppose that $A(t)$ is i.i.d. over timeslots, and that $S(t)$ is i.i.d. over timeslots. Define r_{min} and r_{max} as follows:

$$\begin{aligned} r_{min} &\triangleq \mathbb{E}\left\{\Psi(A(t), K)\right\} \\ r_{max} &\triangleq \mathbb{E}\left\{C(P_{max}, S(t))\right\} \end{aligned}$$

That is, r_{min} is the minimum bit rate delivered to the queueing system, assuming that the maximum compression option $k(t) = K$ is always used. The value of r_{max} represents the maximum possible transmission rate over the wireless link, assuming that $P_{tran}(t) = P_{max}$ is used for all time t. We assume that $r_{min} < r_{max}$.

Thus, there are two reasons to compress data: (i) In order to stabilize the network, we may need to compress (particularly if $\mathbb{E}\left\{A(t)\right\} b > r_{max}$). (ii) We may actually save power if the power used to compress is less than the extra amount of power that would be used transmitting the extra data if it were not compressed.

Optimal Stationary Randomized Compression: Suppose that for each rate r such that $r_{min} \leq r \leq \mathbb{E}\left\{A(t)\right\} b$, there is a stationary randomized compression rule that chooses $k^*(t)$ as a random function of $A(t)$ (and independent of queue backlog), such that $k^*(t) \in \{0, 1, \ldots, K\}$, and:

$$\begin{aligned} \mathbb{E}\left\{\Psi(A(t), k^*(t))\right\} &= r \\ \mathbb{E}\left\{P^*_{compress}(t)\right\} &= h^*(r) \end{aligned}$$

where $h^*(r)$ is the minimum average power required to compress the $A(t)$ input to a bit stream of time average rate r.

Optimal Stationary Randomized Power Allocation: Suppose that for each rate r such that $0 \leq r \leq r_{max}$, there is a stationary randomized power allocation rule that chooses $P^*_{tran}(t)$ as a random function of $S(t)$ (and independent of queue backlog), such that $0 \leq P^*_{tran}(t) \leq P_{max}$, and:

$$\begin{aligned} \mathbb{E}\left\{C(P^*_{tran}(t), S(t))\right\} &= r \\ \mathbb{E}\left\{P^*_{tran}(t)\right\} &= g^*(r) \end{aligned}$$

where $g^*(r)$ is the minimum average power required to yield a time average transmission rate r over the wireless link.

a) Let $L(U) = \frac{1}{2}U^2$. Use (1) to compute a bound on the Lyapunov drift of the form:

$$\Delta(U(t)) \leq B - U(t)\mathbb{E}\left\{C(P_{tran}(t), S(t)) - \Psi(A(t), k(t)) \mid U(t)\right\}$$

make sure to compute the constant B.

b) Design a dynamic control policy that observes $A(t)$ every slot t and chooses a compression option $k(t) \in \{0, 1, \ldots, K\}$, and that observes $S(t)$ every slot t and chooses a power option $P_{tran}(t)$ such that $0 \leq P_{tran}(t) \leq P_{max}$. Your policy should attempt to minimize total average power expenditure. Your policy should be based on a control parameter $V \geq 0$ that affects an explicit tradeoff between average queue congestion and average power expenditure. Your policy should not require knowledge of the channel statistics $Pr[S(t) = S]$ or the traffic statistics $Pr[A(t) = a]$.

c) Show that your policy can make the time average power expenditure arbitrarily close to the minimum possible average power expenditure, with a tradeoff in average delay.

Note: This would be an example of a great course project (it could also likely be turned into a research paper). The problem formulation is interesting, and the proposed solution directly applies the theory of the course. You can likely formulate other problems in your own area of interest for your course project.

II. DESIGNING FAIR TRANSMISSION RATES (50 POINTS)

Consider a K-user downlink with time varying channel states and a single server. Only one channel can be served each timeslot. Let $x_k(t)$ be the channel state of link k, representing the transmission rate supportable over link k if the server is allocated to link k during slot t. Let $\boldsymbol{x}(t) = (x_1(t), \ldots, x_K(t))$ represent the vector of channel states. The $\boldsymbol{x}(t)$ vector is assumed to be known at the beginning of each timeslot so that the controller can choose which channel to serve based on this information. Assume that the vector $\boldsymbol{x}(t)$ is i.i.d. over timeslots (with unknown probabilities), and that transmission rates are bounded so that:

$$0 \leq x_i(t) \leq x_{max} \text{ for all } i \in \{1, \ldots, K\} \text{ and all } t$$

Let $\mu_i(t) = x_i(t)1_i(t)$, where $1_i(t)$ is an indicator function representing our control action during slot t. Specifically:

$$1_i(t) = \begin{cases} 1 & \text{if we serve channel } i \text{ during slot } t \\ 0 & \text{otherwise} \end{cases}$$

Define the time average rate $\overline{\mu}_i$ as follows (assuming the limit exists):

$$\overline{\mu}_i \overset{\triangle}{=} \lim_{t \to \infty} \frac{1}{t} \sum_{\tau=0}^{t-1} \mathbb{E}\left\{\mu_i(\tau)\right\}$$

Let Λ represent the set of all possible time average rate vectors $(\overline{\mu}_1, \ldots, \overline{\mu}_K)$ achievable by some server allocation policy.

For each $i \in \{1, \ldots, K\}$, let $g_i(r)$ represent a concave increasing utility curve. We want to design a server allocation rule that every timeslot observes the current channel states $\boldsymbol{x}(t) = (x_1(t), \ldots, x_K(t))$ and chooses which channel to serve. The goal of the policy is to achieve a time average service rate vector $(\overline{\mu}_1, \ldots, \overline{\mu}_K)$ such that each entry of this time average vector is greater than or equal to the optimal solution vector $(r_1^*, \ldots, r_K^*)$ that solves:

$$\begin{aligned} \text{Maximize:} \quad & \textstyle\sum_{i=1}^K g_i(r_i) \\ \text{Subject to:} \quad & (r_1, \ldots, r_K) \in \Lambda \\ & r_i^{min} \leq r_i \leq r_i^{max} \text{ for all } i \in \{1, \ldots, K\} \end{aligned}$$

where $\{r_i^{min}\}_{i=1}^K$ and $\{r_i^{max}\}_{i=1}^K$ represent a collection of minimum service rate requirements and maximum service rate constraints, respectively. It is assumed that $r_i^{min} \geq 0$ for all i, and that:

$$\boldsymbol{r}^{min} \overset{\triangle}{=} (r_1^{min}, r_2^{min}, \ldots, r_K^{min}) \text{ is strictly interior to } \Lambda$$

That is, it is possible to support the minimum rate requirement $\boldsymbol{r}^{min}$. You can also assume that $r_i^{min} < r_i^{max}$ for all $i \in \{1, \ldots, K\}$. Note that there are no queues in this problem. You can create your own queues if you want.

a) Design a server allocation policy to solve the above problem. Your policy can be in terms of a control parameter $V > 0$ that yields a solution within $O(1/V)$ of the optimal solution of the above problem, with a corresponding tradeoff in something else.

b) Show that your policy gets the optimal solution to within $O(1/V)$. Explicitly show the tradeoff and comment on what the tradeoff means in terms of the physical system.

Hint: Recall that for all vectors $\boldsymbol{r} = (r_1, \ldots, r_K) \in \Lambda$, there exists a stationary randomized control policy for choosing $\boldsymbol{\mu}(t)$ as a potentially random function of the observed channel states $\boldsymbol{x}(t)$ such that: $\mathbb{E}\left\{\boldsymbol{\mu}^*(t)\right\} = \boldsymbol{r}$.