MCGILL UNIVERSITY

ECSE 506 Term Project

Discrete Time Stochastic Adaptive Control

Marc-Etienne Brunet | 260328066 & Hadrien Baillargeon Legault |260236772

4/12/2012

This report discusses some topics in discrete time stochastic control, specifically a paper by Graham C. Goodwin, Peter J. Ramadge and Peter E. Caines published in 1981.

Glossary and Notation

State space representation of system:

$$x_{t+1} = \mathcal{A}x_t + \mathcal{B}u_t + \mathcal{K}w_t$$

$$y_t = \mathcal{C}x_t + w_t$$

ARMAX representation:

$$A(q^{-1})y_t = q^{-d}[B(q^{-1})]u_t + [C(q^{-1})]w_t$$

$$A(q^{-1}) = 1 + a_1 q^{-1} + \ldots + q^{-n}$$

$$B_{ij}(q^{-1}) = b_{ij}^0 + b_{ij}^1 q^{-1} + \ldots + b_{ij}^m q^{-m}$$

$$C_{ij}(q^{-1}) = c_{ij}^0 + c_{ij}^1 q^{-1} + \ldots + c_{ij}^l q^{-l}$$

Glossary:

x_t: System state vector

u_t: Input vector

y_t: Output vector

y_t^*: Desired output

q^{-1} : Unit delay operator, i.e. $q^{-1}y_t = y_{t-1}$

s : Dimension of output y_t

r: Dimension of input u_t

k: Maximum delay, i.e. $\max(n, m+d, l)$

x_0 : Initial conditions

θ_t: Estimate of algorithm system parameters (Note: $\hat{\theta}_t$ in the paper)

θ^*: Actual algorithm system parameters (Note: θ_0 in the paper)

ϕ_t: Algorithm state vector

Variable Simplifications:

$$v_t \triangleq y_t - \mathbb{E}[y_t|\mathcal{F}_{t-d}]$$

$$e_t \triangleq y_t - y_t^*$$

$$z_{t-d} \triangleq \mathbb{E}[y_t|\mathcal{F}_{t-d}] - y_t^* = e_t - v_t$$

Special Notation:

$$x_{1:t} \triangleq \{x_1, x_2, \ldots x_t\}$$

$$a.s. \triangleq almost\ surely$$

Contents

Introduction

In 1981, Graham C. Goodwin, Peter J. Ramadge and Peter E. Caines published a paper entitled "Discrete Time Stochastic Adaptive Control" in the Society for Industrial and Applied Mathematics' (SIAM) journal Control and Optimization. Their paper established the global convergence of a stochastic adaptive control algorithm for discrete time linear systems. This algorithm could "learn" the system dynamics and asymptotically perform as well as could be achieved if the system parameters were known.

This report aims to summarize and discuss some of the results of their paper. It starts by giving the reader some background on adaptive control, formulating the general problem and explaining how to move from a state-space to an ARMAX model. It then, in our opinion, clarifies the resulting SISO unit-delay algorithm by rectifying a few typos and tying up some loose ends. Finally, it walks the reader through the proof of global convergence, concludes with a few general remarks and includes some simulation results in our Appendix C.

Problem Formulation

Background

The general discrete adaptive control problem can be formulated as follows:

$$x_{t+1} = f(x_{1:t}, \theta_{1:t}, u_{1:t}, w_{1:t})$$

$$y_t = g(x_{1:t}, \theta_{1:t}, u_{1:t}, w_{1:t})$$

Where x_t and y_t are the state of the system and observation available to the decision maker at time t respectively, and evolve according to past state values x, control inputs u, and some noise process w. The dynamics of this evolution are known, excluding the set of (possibly time dependant) parameters θ.

As a result, the decision maker must choose a control policy with an incomplete understanding of the governing dynamics, and is thus faced with the added complexity of learning the values of the parameters θ. Depending on the specifications of the control objective: regulation, tracking, cost minimization… etc, and time horizon: finite or infinite, the decision maker must determine the control inputs, faced with a tradeoff of between trying to meet the current required output or better understanding the dynamics in hopes of improving future performance.

There are two major approaches to adaptive control: implicit and explicit. In the explicit case an estimate of the missing parameters is computed online through estimation and the control policy uses these estimates to determine the control input. In the implicit case, the control policy is updated without explicitly computing the missing system parameters. In the paper being considered, the authors' algorithm takes the latter approach; while a set of parameters is estimated, they represent parameters in the control policy and not explicitly the missing system parameters.

Intense research in adaptive control methods was motivated by aircraft control in the 1950's. Aircrafts have highly non-linear dynamics and so would be linearized about several different operating points depending on flight conditions. Control engineers were looking for innovative ways to deal with the resulting parameter changes and turned to adaptive control techniques [Robust Adaptive Control].

This particular paper looks at systems where the governing dynamics are linear. The authors however, prove their results for a class of systems in auto-regressive moving average (ARMA) form. This may be unfamiliar to those used to working with a state space representation (as was our case), and so the authors included a transformation that demonstrates the equivalence of the two models (Appendix B). The process is very similar to changing a state space model into its s-domain transfer function representation and will be discussed presently.

State-Space to ARMA

The authors start with a fairly general model for a discrete time stochastic linear system:

$$x_{t+1} = \mathcal{A}x_t + \mathcal{B}u_t + \mathcal{K}w_t$$

$$y_t = \mathcal{C}x_t + w_t$$

Where x_t, y_t, u_t, w_t are the system state, output, control input and a stochastic process on $(\Omega, \mathcal{A}, P)$ respectively, and $\mathcal{A}, \mathcal{B}, \mathcal{K}, \& \ C$ are matrices of appropriate dimensions. As we are dealing with the discrete case, t takes on discrete values. It can be shown that this system can be put in the equivalent auto-regressive moving average (ARMA) form:

Equation 1

$$A(q^{-1})y_t = q^{-d}[B(q^{-1})]u_t + [C(q^{-1})]w_t$$

$$A(q^{-1}) = 1 + a_1 q^{-1} + \ ... + q^{-n}$$

$$B_{ij}(q^{-1}) = b_{ij}^0 + b_{ij}^1 q^{-1} + \ ... + b_{ij}^m q^{-m}$$

$$C_{ij}(q^{-1}) = c_{ij}^0 + c_{ij}^1 q^{-1} + \ ... + c_{ij}^l q^{-l}$$

Where $q^{-1}$is the unit delay operator, A is the characteristic polynomial of $\mathcal{A}$, and $[B(-)]$ and $[C(-)]$ are (in the most general case) matrices of polynomials. (Please carefully note that, with the exception of A, there is no simple relationship between script and plain text variables, ex. $\mathcal{B}$ and B.)

By recursively applying the state update equation, it is found that a future state (at time $t + k$) can be related to an earlier state (at time t) by:

Equation 2

$$x_{t+k} = \mathcal{A}^k x_t + \sum_{i=1}^{k} \mathcal{A}^{i-1}(\mathcal{B}u_{t+k-i} + \mathcal{K}w_{t+k-i})$$

Letting A have characteristic polynomial:

$$A(\lambda) = \lambda^n + a_1\lambda^{n-1} + \cdots + a_n$$

The transformation hinges on the Cayley-Hamilton theorem, which shows every matrix solves its own characteristic polynomial. These results are used to solve for $\mathcal{A}^n = (-a_1\mathcal{A}^{n-1} - a_2\mathcal{A}^{n-2} \ldots - a_n I)$. Thus if we take k to be n in (Equation 2) we can sub in these results getting:

Equation 3

$$x_{t+n} = -\sum_{j=1}^{n} a_j \mathcal{A}^{n-j} x_t + \sum_{i=1}^{n} \mathcal{A}^{i-1}(\mathcal{B}u_{t+n-i} + \mathcal{K}w_{t+n-i})$$

Further, if we rearrange (Equation 2) to describe $\mathcal{A}^{n-j}x_t$ for $j = \{1, 2 \ldots n-1\}$ we can sub this back into (Equation 3) and after some index manipulation we get:

Equation 4

$$x_{t+n} = -\sum_{j=1}^{n} a_j x_{t+n-j} + \sum_{j=1}^{n-1} a_{n-j} \sum_{l=1}^{j} \mathcal{A}^{l-1}\left(\mathcal{B}u_{t+j-l} + \mathcal{K}w_{t+j-i}\right)$$
$$+ \sum_{i=1}^{n} \mathcal{A}^{i-1}(\mathcal{B}u_{t+n-i} + \mathcal{K}w_{t+n-i})$$

We can then multiply by $\mathcal{C}$, add w where needed and use $y = \mathcal{C}x + w$ to obtain a linear equality involving $y_{t:(t+n)}, u_{t:(t+n)}$ and $w_{t:(t+n)}$. This equality is of ARMA form and, with some attention to detail, the coefficients of the entries of $[B(q^{-1})]$ and $[C(q^{-1})]$ can be computed in terms of $\mathcal{A}$, $\mathcal{B}$, and $\mathcal{C}$. Many of these details can be found in the paper's appendix B. With this equivalence established the authors prove the remainder of the results for systems of this form.

Specific Objectives

The objective of this paper is to prove the global convergence of a class of adaptive control algorithms for the aforementioned ARMA system (Equation 4). By global convergence the authors mean that for all initial system and algorithm states the algorithm will:

R1: ensure: $\limsup_{N\to\infty} \frac{1}{N}\sum_{t=1}^{N}(y_t)^2 < \infty$

R2: ensure: $\limsup_{N\to\infty} \frac{1}{N}\sum_{t=1}^{N}(u_t)^2 < \infty$

R3: whenever it exists, minimize: $\lim_{N\to\infty} \frac{1}{N}\sum_{t=1}^{N} \mathbb{E}[(y_t - y_t^*)^2 | \mathcal{F}_{t-d}]$

In the general case it is assumed that the dimension of the output y is s, while the dimension of the input u is r. The deterministic sequence $\{y_t^*\}$ what we are trying to track. The process $\{x_0, w_1, w_2 \ldots\}$ is defined on the underlying probability space $(\Omega, \mathcal{F}, P)$. We define $\mathcal{F}_0$to be the sigma-algebra generated by the initial conditions $x_0 = \{y_{(1-n):0}, u_{(1-m-d):0}, w_{(1-l):0}\}$, $\mathcal{F}_t$ to be the sigma algebra-generated by $\{x_0, w_1, \ldots, w_t\}$, and thus $\mathcal{F}_0 \subset \mathcal{F}_1 \ldots \subset \mathcal{F}_t$ is a filtration.

The authors also require the following conditional independence and variance conditions:

$$\mathbb{E}[w_t|\mathcal{F}_{t-1}] = 0 \quad ; t \geq 1 \;\; a.s.$$

$$\mathbb{E}[w_t w_t^T|\mathcal{F}_{t-1}] = Q \quad ; t \geq 1 \;\; a.s.$$

Note that these are stronger conditions than simply requiring the unconditional expectation and variance of w_t to be zero and Q (which follows as a result of the above).

Finally we require u_t to be measurable with respect to the sigma-algebra generated by the initial conditions x_0 and outputs $\{y_{1:t}\}$, and note that this is in general smaller than $\mathcal{F}_t$.

Existence of a solution for the SISO unit-delay case

Particularities of the Objectives

In the single input single output (SISO) case, (where the dimension of the input and output vectors are one), the authors motivate the objective by first exploring the limit in (R3). Before we discuss how they proceeded we would first like to explore the meaning of the term ourselves.

At any given instance we seek to minimize the expected square of the error:

$$\mathbb{E}[(y_t - y_t^*)^2]$$

By the towering property of the expectation this can be rewritten as:

$$\mathbb{E}\big[\, \mathbb{E}[(y_t - y_t^*)^2|\mathcal{F}_{t-d}]\big]$$

Where the inner expectation is conditioned on the sigma-algebra generated by all the information that could possibly be available at the time of deciding on a control input. (Recalling that from (Equation 1) the control input can affect the output no sooner than d time steps later.)

We now change the exterior expectation to a sample mean in order to capture the time evolution, getting the term in (R3):

$$\lim_{N\to\infty} \frac{1}{N} \sum_{t=1}^{N} \mathbb{E}[(y_t - y_t^*)^2|\mathcal{F}_{t-d}]$$

Motivation of Minimal Variance Control Policy

The authors show that the term inside the above sum (and (R.3)) is bounded below by a constant γ^2. To do so they first expand the quadratic, then use the linearity of the expectation and the fact that y_t^* is deterministic (and therefore $\mathcal{F}_{t-d}$ measurable) to get:

$$\mathbb{E}[(y_t - y_t^*)^2|\mathcal{F}_{t-d}] = \mathbb{E}[(v_t)^2|\, \mathcal{F}_{t-d}] + z_t^2$$

Where they have defined $v_t = y_t - \mathbb{E}[y_t|\mathcal{F}_{t-d}]$,the innovation, and $z_t = \mathbb{E}[y_t|\mathcal{F}_{t-d}] - y_t^*$. They then show that $(v_t)^2 = \gamma^2$ a constant, and that the algorithm brings $\lim_{N\to\infty} \frac{1}{N}\sum_{t=1}^{N}(z_t)^2$ to zero almost surely, which proves that in the limit the solution is optimal.

To do so the authors begin by using a lemma in the paper's Appendix C, factoring $C(q^{-1})$ as:

$$C(q^{-1}) = F(q^{-1})A(q^{-1}) + q^{-d}G(q^{-1})$$

The details of which are omitted from the paper. However it is worth considering the degree of the polynomials; from our understanding this decomposition (in the way that its used) seems to require $l + 1 < n + d$.

This decomposition is absolutely critical in the proof since when subbed into (Equation 1) we get:

Equation 5

$$C(q^{-1})(y_t - F(q^{-1})w_t) = q^{-d}G\big(q^{-d}\big)y_t + q^{-d}F(q^{-1})B(q^{-1})u_t$$

Where every term on the right side of the equation is delayed by at least d time steps. This allows for two important manipulations.

Firstly, along with (Equation 4) which we recall describes y_t in terms of past values, we can write:

$$y_t - \sum_{k=0}^{d-1} f_k\, w_{t-k} = func(past) + \beta u_{t-d}$$

Where $func(past)$ is $\mathcal{F}_{t-d}$ measurable. So taking the conditional expectation of both sides, (and recalling the zero mean noise requirements), we get:

Equation 6

$$\mathbb{E}[y_t|\mathcal{F}_{t-d}] = func(past) + \beta u_{t-d}$$

This in turn gives:

$$v_t \triangleq y_t - \mathbb{E}[y_t|\mathcal{F}_{t-d}] = \sum_{k=0}^{d-1} f_k\, w_{t-k} \quad a.s.$$

From this the authors easily show:

Equation 7

$$\gamma^2 \triangleq \mathbb{E}[(v_t)^2|\mathcal{F}_{t-d}] = Q^2 \sum_{k=0}^{d-1}(f_k)^2 \quad a.s.$$

We also note from (Equation 6) that $\mathbb{E}[y_t|\mathcal{F}_{t-d}]$ can be made to take on any value by appropriate choice of control input. Note that from the conditions in (Eq.4), β is guaranteed to be non-zero (since by factoring out the minimum delay d, $B(q^{-1})$ has a constant term).

These two results (Equation 6) & (Equation 7) really serve to motivate the search for the adaptive control strategy and the proof of its convergence.

The SISO unit-delay Algorithm

Note that this explanation differs slightly from that of the authors, includes some minor corrections, uses a slightly different notation, and is re-indexed to start at $t = 1$.

Setup

Define algorithm state vector (using x_0):

$$\phi_0 \triangleq \left[\underbrace{y_0, \dots y_{1-n}}_{n\ entries}, \underbrace{u_0, \dots, u_{-m}}_{m+1\ entries}, \underbrace{-y_0^*, \dots, -y_{1-l}^*}_{l\ entries}\right]^T$$

Define additional scalar state variable:

$$r_0 \triangleq 1 + \phi_0^T \phi_0$$

Define θ_0, a vector of the same dimensions as ϕ_0 that estimates the algorithm system parameters, either arbitrarily or with whatever information is available.

Updating

Now for $t = 1, 2, 3, \dots$

1) Determine y_t

2) Update parameter estimate according to: (with $\bar{a} > 0$ discussed later)

Equation 8

$$\theta_t = \theta_{t-1} + \frac{\bar{a}}{r_{t-1}} \phi_{t-1}(y_t - y_t^*)$$

3) Solve* the below equation for u_t , noting that it is the only unknown:

$$\phi_t^T \theta_t = y_{t+1}^*$$

4) Update the algorithm state vector with the new information:

$$\phi_t = [y_t, \dots y_{t-n+1}, u_t, \dots, u_{t-m}, -y_t^*, \dots, -y_{t-l+1}^*]$$

5) Update r:

$$r_t = r_{t-1} + \phi_t^T \phi_t$$

*Note that solving for u_t will involve dividing by the corresponding entry in θ. If θ_0 or u_0 are chosen at random then division by zero is a probability zero event.

Basis of the Algorithm

While (Equation 6) and (Equation 7) motivate the search for a solution, they do not explain the authors' choice of state vector in the algorithm above. This choice becomes more obvious when we take the conditional expectation of (Equation 5) and substract $C(q^{-1})y_t^*$ from each side getting:

$$C(q^{-1})(\mathbb{E}[y_t|\mathcal{F}_{t-1}] - y_t^*) = q^{-1}G(q^{-1})y_t + q^{-1}f_0B(q^{-1})u_t - C(q^{-1})y_t^*$$

Where d has been replaced by 1 (which according to Appendix C causes $F(q^{-1}) = f_0$) and noting the unit delay operator is allowed to operated on the sigma-algebra. Looking closely we see that this is of the form:

$$C(q^{-1})(\mathbb{E}[y_t|\mathcal{F}_{t-1}] - y_t^*) = \phi_t^T\theta^* - y_t^*$$

Referring again to Appendix C for polynomial degrees.

Explain(4.6) in the paper.

Introduce z_t, e_t, v_t in a meaningful way.

The Proof

Main idea

The bulk of the proof consists of showing bounded mean square output tracking error. Upon obtaining this result the others (R.1 & R.2) follow nicely. The strategy here is to show that:

$$\lim_{\mathrm{n}\to\infty} \frac{1}{N}\sum_{t=0}^{N}\left(z(t)\right)^2 = 0 \; a.s.$$

(Recalling $z_{t-1} \triangleq \mathbb{E}[y_t|\mathcal{F}_{t-1}] - y_t^* = \mathbb{E}[y_t|\mathcal{F}_{t-1}] - \phi_{t-1}^T\theta_{t-1}$). The first term is the best possible estimate of the next output. Notice that it is not computable as the required information is not available to the system. The second term is the algorithm's estimate of the next output, given the current system parameters. This is done by defining the stochastic equivalent to a Lyapunov function.

The major steps of the proof are as follows:

1. Using the "Lyapunov function" mentioned above, along with the Neuveu Martingale Convergence Theorem show:

Equation 9

$$\lim_{N\to\infty}\left(\frac{N}{r(N)}\right)\frac{1}{N}\sum_{t=1}^{N}\left(z(t)\right)^2 = 0\ a.s$$

2. As a second step, show that

Equation 10

$$\lim_{N\to\infty}\inf\left(\frac{N}{r(N)}\right) > \frac{1}{K} > 0$$

3. Thus deduce that

Equation 11

$$\lim_{\mathrm{N}\to\infty}\frac{1}{N}\sum_{t=0}^{N}\left(z(t)\right)^2 = 0\ a.s.$$

4. The conclusions regarding the boundedness of the input and output sequences will be derived while progressing through steps 1 and 2.

Proof Details

Base Assumption and Initial Note

Recall assumptions:

A1: $d = r = s = 1$

A2: $Upper\ bounds\ for\ n, m, l\ are\ known.$

A3: $C(-) and\ B(-) have\ all\ zeros\ outside\ the\ closed\ unit\ circle.$

A4: $C(-) - \frac{\bar{a}}{2}\ is\ positive\ real.$

The following outline more or less mirrors the proof described in the original paper. This minimizes the amount back and forth for the reader. Some of the less obvious steps will be outlined and explained in a way that allow readers less familiar with measure-theoretic probability to follow easily. On that note, should the reader require a refresher, many important concepts of probability with Matingales are outlined in Appendix A.

Some of the lemmas are left unproved. This is done in the interest of remaining concise and because the goal of the report was to study adaptive control, not necessarily to delve in Martingale theory or obscure lemmas.

Proof

To obtain condition Equation 9 the authors, through a series of manipulations, define a stochastic Lyapunov function V.

Part 1: Obtaining a stochastic Lyapunov-type equation

First, $\tilde{\theta}_t \triangleq \theta_t - \theta^*$ is defined. $\tilde{\theta}_t$ is the error between the real algorithm parameters and the estimated algorithm parameters. As the point of the algorithm is to reduce this error over time, this seems as a good starting point for a Lyapunov function. I.e. it could have a role similar to energy in dynamical systems: if its value continually decreases, we can expect some converging behavior of the system.

Thus, the authors define a new stochastic process $V_t = \tilde{\theta}_t^{\ T} \tilde{\theta}_t$. Thus, recalling Equation 8 we get:

$$\tilde{\theta}_t = \tilde{\theta}_{t-1} + \frac{\bar{a}}{r_{t-1}} \phi_{t-1} e_t$$

(Recall that $e_t = y_t - y_t^*$)

$$V_t = \left[\tilde{\theta}_{t-1} + \frac{\bar{a}}{r_{t-1}} \phi_{t-1} e_t\right]^T \left[\tilde{\theta}_{t-1} + \frac{\bar{a}}{r_{t-1}} \phi_{t-1} e_t\right]$$

$$V_t = V_{t-1} + \frac{2\bar{a}}{r_{t-1}} \phi_{t-1}^T \phi_{t-1} e_t + \frac{\bar{a}^2}{(r_{t-1})^2} \phi_{t-1}^T \phi_{t-1} (e_t)^2$$

Were v_t is as defined previously. We thus get[1]

$$V_t = V_{t-1} + \frac{2\bar{a}}{r_{t-1}} \tilde{\theta}_{t-1}^T \phi_{t-1} (e_t - v_t) + \frac{2\bar{a}}{r_{t-1}} \tilde{\theta}_{t-1}^T \phi_{t-1} v_t + \frac{\bar{a}^2}{(r_{t-1})^2} \phi_{t-1}^T \phi_{t-1} [(e_t - v_t)^2 + 2v_t(e_t - v_t) + (v_t)^2]$$

The following variables are then defined

$$b_{t-1} = -\tilde{\theta}_{t-1}^T \phi_{t-1} = \theta^{*T} \phi_{t-1} - \theta_{t-1}^T \phi_{t-1}$$

$$z_{t-1} = e_t - v_t = (y_t - y_t^*) - (y_t - \mathbb{E}[y_t | \mathcal{F}_{t-d}])$$

b_t represents the difference in the actual output at the next instant (first term), and the next output predicted by the algorithm.

z_t represents the difference between the tracking error and the optimal prediction error. Ideally, the algorithm send this value to 0.

---

[1] Using the following substitutions

$$e_t = e_t - v_t + v_t$$

$$(e_t)^2 = (e_t - v_t)^2 + 2v_t(e_t - v_t) + (v_t)^2$$

Taking the expectation (see Appendix B2.1 for some pointers if needed) and then substituting b_t and z_t where appropriate:

Equation 12

$$\mathbb{E}[V_t|\mathcal{F}_{t-1}] = V_{t-1} - \frac{2\bar{a}}{r_{t-1}} b_{t-1} z_{t-1} + \frac{\bar{a}^2}{(r_{t-1})^2} \phi_{t-1}^T \phi_{t-1} (z_{t-1})^2 + \frac{\bar{a}^2}{(r_{t-1})^2} \phi_{t-1}^T \phi_{t-1} (\gamma)^2 \; a.s.$$

Note that for a function V_t to be a stochastic Lyapunov, it needs to be a supermartingale and non-negative. Therefore, some more modification should be made to Equation 12.

Additionally, note that Equation 12 is nearly in a form where the MCT could be used. This would be desirable as it would provide information regarding the summability of $\frac{2\bar{a}}{r_{t-1}} b_{t-1} z_{t-1}$. This motivates the next few steps:

Getting Inequality

Note: $r_{t-1} = r_{t-2} + \phi_{t-1}^T \phi_{t-1} \Rightarrow 1 = \frac{r_{t-2} + \phi_{t-1}^T \phi_{t-1}}{r_{t-1}} \Rightarrow \frac{\phi_{t-1}^T \phi_{t-1}}{r_{t-1}} \leq 1$

Thus Equation 12 becomes:

$$\mathbb{E}[V_t|\mathcal{F}_{t-1}] \leq V_{t-1} - \frac{2\bar{a}}{r_{t-1}} b_{t-1} z_{t-1} + \frac{\bar{a}^2}{r_{t-1}} (z_{t-1})^2 + \frac{\bar{a}^2}{(r_{t-1})^2} \phi_{t-1}^T \phi_{t-1} (\gamma)^2 \; a.s.$$

And rewriting $\bar{a}^2 (z_{t-1})^2 = (\bar{a} + \rho)(z_{t-1})^2 - \rho (z_{t-1})^2$:

$$\mathbb{E}[V_t|\mathcal{F}_{t-1}] \leq V_{t-1} - \frac{2\bar{a}}{r_{t-1}} \left\{ b_{t-1} - \frac{\bar{a} - \rho}{2} z_{t-1} \right\} z_{t-1} - \frac{\rho \bar{a}}{r_{t-1}} (z_{t-1})^2 + \frac{\bar{a}^2}{(r_{t-1})^2} \phi_{t-1}^T \phi_{t-1} (\gamma)^2 \; a.s.$$

Using

$$h_{t-1} = b_{t-1} - \frac{\bar{a} - \rho}{2} z_{t-1}$$

One easily obtain:

$$h_{t-1} = \left[C(q^{-1}) - \frac{\bar{a} + \rho}{2} \right] z_{t-1}$$

Where the condition on ρ is such that $C(z) - \frac{\bar{a} - \rho}{2}$ is positive real (such a ρ guaranteed to exist by

Which expresses h_t as regression of z_t through a positive real transfer function. Note this will allow the use of the lemma A.4 (found in [1]) (Kalam-Popov-Yakubovich Lemma[2]). This ensures us that there exists

$$S_t = 2\bar{a} \sum_{j=1}^{t} h(j-1) z(j-1) + K, 0 < K < \infty \text{ such that } S_t \geq 0$$

[2] Sometimes referred to as positive real lemma

Then, the following variable is defined:

$$Z_t = V_t + \frac{S_t}{r(t-1)}$$

Substituting where appropriate:

$$\mathbb{E}[Z_t|\mathcal{F}_{t-1}] \leq V_{t-1} + \frac{S_t}{r_{t-1}} - \frac{2\bar{a}}{r_{t-1}} h_{t-1} z_{t-1} \ - \frac{\rho \bar{a} z_{t-1}^2}{r_{t-1}} + \frac{\bar{a}^2}{(r_{t-1})^2} \phi_{t-1}^T \phi_{t-1} (\gamma)^2 \ a.s.$$

Notice that the second and third terms combine:

$$\mathbb{E}[Z_t|\mathcal{F}_{t-1}] \leq V_{t-1} + \frac{S_{t-1}}{r_{t-1}} \ - \frac{\rho \bar{a} z_{t-1}^2}{r_{t-1}} + \frac{\bar{a}^2}{(r_{t-1})^2} \phi_{t-1}^T \phi_{t-1} (\gamma)^2 \ a.s.$$

Noting that $r_{t-2} \leq r_{t-1}$

$$\begin{aligned}\mathbb{E}[Z_t|\mathcal{F}_{t-1}] &\leq V_{t-1} + \frac{S_{t-1}}{r_{t-2}} - \frac{\rho \bar{a} z_{t-1}^2}{r_{t-1}} + \frac{\bar{a}^2}{(r_{t-1})^2} \phi_{t-1}^T \phi_{t-1} (\gamma)^2 \\ &= Z_{t-1} - \frac{\rho \bar{a} z_{t-1}^2}{r_{t-1}} + \frac{\bar{a}^2}{(r_{t-1})^2} \phi_{t-1}^T \phi_{t-1} (\gamma)^2 \ a.s.\end{aligned}$$

By lemma A.2 (found in [1]]):

$$\sum_{t=1}^{\infty} \frac{\phi_{t-1}^T \phi_{t-1}}{(r_{t-1})^2} < \infty$$

Thus, MCT can finally be used. Taking $\beta_{n-1} = \frac{\phi_{t-1}^T \phi_{t-1}}{(r_{t-1})^2}$, following result is obtained (taking in account that $\rho \bar{a} \neq 0$):

Equation 13

$$\sum_{t=1}^{\infty} \frac{z_{t-1}^2}{r_{t-1}} < \infty \ a.s.$$

Next, Kronecker's lemma (see Appendix B2.2) is used to obtain a long-range average conclusion based on Equation 13:

Note that in our case, the sequence $\sum_{t=1}^{\infty} \frac{z_{t-1}^2}{r_{t-1}}$ and r_n satisfy the lemma's requirements and we thus have

$$\lim_{n \to \infty} \frac{n}{n} \frac{1}{r_n} \sum_{t=1}^{n} \frac{r_t (z_t)^2}{r_t} = 0, a.s.$$

Hence the objective of the first part:

$$\lim_{N\to\infty}\left(\frac{N}{r(N)}\right)\frac{1}{N}\sum_{t=1}^{N}(z_t)^2 = 0\ a.s$$

Part 2. The previous result is used to prove the required claims.

Useful lemmas:

Going back to the main proof development:

From lemma A.5 (which follows easily once A.1 (see B2.3 for some clarifications) is obtained), if we reconsider our system as having input y_t and w_t, there exists an N' such that:

Equation 14

$$\frac{1}{N}\sum_{t=1}^{N}(u_t)^2 \leq \frac{K_1}{N}\sum_{t=0}^{N}(y_{t+1})^2 + K_2, for\ N > N'\ a.s.$$

This is one of our objectives: the inputs are mean squared bounded.

Recalling that r(N) is an affine combination of $(y_{t-1})^2, \dots, (y_{t-n})^2, (u_{t-1})^2, \dots, (u_{t-m})^2$ (affine because the y*'s produce a constant as they are computed offline).

Thus Equation 15 follows pretty easily:

Equation 15

$$\frac{r(N)}{N} \leq \frac{K_3}{N}\sum_{t=1}^{N}(y_{t+1})^2 + K_4, for\ N > N' a.s.$$

Keeping in mind that $y_i = z_{i-1} + y_i^* + v_i$ and using a trick similar to what was used in Appendix B2.3 to bound$(y_i)^2$ the following is obtained (M_2 arises because the y* are actually computable):

$$\frac{1}{N}\sum_{i=1}^{N}(y_{i+1})^2 \leq \frac{3}{N}\sum_{i=1}^{N}(z_i)^2 + M_2 + \frac{3}{N}\sum_{i=1}^{N}(v_{i+1})^2$$

By:

$$\lim_{N\to\infty}\sup\frac{1}{N}\sum_{i=1}^{N}(w_i)^2 < \infty\ a.s.$$

We must have:

$$\lim_{N\to\infty} \sup \frac{1}{N}\sum_{i=1}^{N}(v_i)^2 < \infty \; a.s.$$

Thus:

Equation 16

$$\frac{1}{N}\sum_{i=1}^{N}(y_{i+1})^2 \leq \frac{3}{N}\sum_{i=1}^{N}(z_i)^2 + M_3 \; a.s. for\; N \geq N''$$

Then substituting Equation 16 in Equation 15:

Equation 17

$$\frac{r(N)}{N} \leq \frac{C_1}{N}\sum_{i=1}^{N}(z_i)^2 + C_2, for\; N > \max{(N', N'')}$$

N', N'' defined in Equation 15 and Equation 16 respectively. Equation 14and Equation 17 provide bounds for $\frac{r(N)}{N}$.

Proof by Contradiction

Now, assume a sample unbounded input sequence $\frac{1}{N}\sum_{i=1}^{N}(y_{i+1})^2$. If the algorithm is correct, this should not be allowed and we should have a contradiction.

According to our assumption and by the definition of $r(N)$:

$$\lim_{N\to\infty} \sup \frac{r(N)}{N} = \infty$$

Thus, from Equation 17

$$\lim_{N\to\infty} \frac{1}{N}\sum_{i=1}^{N}(z_i)^2 = \infty$$

Now define:

$$\bar{z}(N) = \frac{1}{N}\sum_{i=1}^{N}(z_i)^2$$

Starting from Equation 17:

$$\left\{\frac{r(N)}{N} \leq \frac{C_1}{N}\sum_{i=1}^{N}(z_i)^2 + C_2\right\} * \frac{1}{\bar{z}(N)} \; for\; N > \bar{N}$$

$$\left\{\frac{r(N)}{\bar{z}(N)} \leq C_1 + \frac{C_2}{\bar{Z}(N)}\right\} for\ N > \bar{N}$$

$$\left(\frac{r(N)}{N}\right)^{-1} \bar{z}(N) \geq \frac{\bar{z}(N)}{C_1 \bar{z}(N) + C_2} for\ N > \bar{N}$$

Paralleling the original proof again:

since $\lim_{N\to\infty} \sup(\bar{z}(N)) = \infty$, there is some subsequence $\{N_K\}$ such that

$$\lim_{K\to\infty} \inf \left[\frac{r(N_K)}{N_K}\right]^{-1} \frac{1}{N_k} \sum_{i=1}^{N_k} (z_i)^2 \geq 1/C_1$$

Which contradicts:

$$\lim_{n\to\infty} \left(\frac{N}{r(N)}\right) \frac{1}{N} \sum_{t=1}^{N} (z_t)^2 = 0\ a.s$$

Thus, by contradiction $\frac{1}{N}\sum_{i=1}^{N}(y_{i+1})^2$ is bounded. This proves our second objective: mean square bounded inputs.

From the above boundedness condition, and the nonnegative nature of $\frac{r(N)}{N}$:

$$\lim_{N\to\infty} \frac{\sup r(N)}{N} < \infty$$

It is thus rather intuitive that

$$\lim_{N\to\infty} inf \frac{N}{r(N)} > \frac{1}{K} > 0$$

Following the original proof, $z_{i-1}\ is\ \mathcal{F}_{i-1}$ measurable. We thus, finally, have the following:

$$\mathbb{E}\{(y_i - y_i^*)^2 | \mathcal{F}_{i-1}\} = \gamma^2 + (z_{i-1})^2\ a.s.$$

$$\boxed{\lim_{N\to\infty} \frac{1}{N} \sum_{i=1}^{N} \mathbb{E}\{(y_i - y_i^*)^2 | \mathcal{F}_{i-1}\} = \gamma^2\ a.s.}$$

And we are done.

Appendix A: Preliminary Knowledge

Martingale Theory and Measure-Theoretic Probability

A key principle used in proving the proposed properties of the algorithm is the Martingale Convergence Theorem. In a nutshell, the Martingale convergence theorem is somewhat analogous to the monotonic convergence theorem in the sense that it provides condition for which a Martingale (a certain kind of sequence of random variables) converges.

Some preliminary background in measure-theoretic probability is required. It will be assumed that the reader understands what is meant by σ-algebra and $\mathcal{F}_1/\mathcal{F}_2$-measurable in the context of the measurable spaces $(\Omega_1, \mathcal{F}_1)$ and $(\Omega_2, \mathcal{F}_2)$

In the special case where a function (most commonly a random variable) is said to be $\mathcal{F}_n$-measurable, the second measurable space is understood to be $(\mathbb{R}, \mathcal{B})$. Where $\mathcal{B}$ is the Borel σ-algebra.

Definition: σ algebra generated by a collection of maps on Ω

Suppose the following collection of maps: $\{Y_\gamma : \gamma \in \Gamma\}$ such that $Y_\gamma : \Omega \to \mathbb{R}\ \forall \gamma$. Then, the σ algebra generated by that collection is as follows:

$$\mathcal{Y} \coloneqq \sigma(\{\omega \in \Omega : Y_\gamma(\omega) \in B\} : \gamma \in \Gamma, B \in \mathcal{B})$$

i.e. it is the smallest σ algebra such that all Y_γ are $\mathcal{Y}$-measurable.

Some Intuition

Say we run some experiment. Some information about ω is obtained by observing the result of some random variable Y. The algebra generated is the collection of events F, for which, for each w, we can decide whether or not F has occurred based on the Y observations

Definition: Filtration

A sequence $\{\mathcal{F}_n : n = 0,1,2, \dots\}$ is called a filtration of the measurable space $(\Omega, \mathcal{F})$ if the following holds:

- $\mathcal{F}_n$ is a sub-algebra of $\mathcal{F}\ \forall n$
- $\mathcal{F}_n \subset \mathcal{F}_{n+1} \forall\, n$

Some Intuition:

In the discrete space case, sigma algebras can be seen as partitions of the whole space into atoms. As sigma algebras include all their unions/complements, the second condition is equivalent to “refining” the space as n grows.

Adapted

A stochastic process $\{X_n : n = 0,1, \dots\}$is said to be adapted to a filtration $(\mathcal{F}_n)$ if the following holds:

- X_n is $\mathcal{F}_n$-measurable $\forall n$

We are now ready to introduce the concept of Martingale. Informally, a Martingale is a mathematical model of a fair wager.

Martingale

An adapted stochastic process $\{X_n, \mathcal{F}_n : n = 0,1, \dots\}$ is a called a Martingale if

- $(\mathcal{F}_n)$ is a filtration, and $X_n$is adapted to it
- X_n is integrable $\forall n$
- $\mathbb{E}[X_{n+1}|\mathcal{F}_n] = X_n, \forall n$

Similarly, one can define sub and super-martingales as follows (respectively favourable and unfavourable games):

Sub:

- $\mathbb{E}[X_{n+1}|\mathcal{F}\_n] > X_n, \forall n$

Super:

- $\mathbb{E}[X_{n+1}|\mathcal{F}\_n] < X_n, \forall n$

One can think of the filtration as the amount of knowledge about the process available at instant n. If X_n is a betting game, a fair game would leave the player, in expectation, with X_n fortune after the n+1 game. Hence we have the terms unfavourable game for a supermartingale and favourable for a submartingale.

Statement of (a) Martingale Convergence Theorem

Let $\{X_n, \mathcal{F}_n : n = 0,1, \dots\}$ be a sub martingale. Additionally, suppose $\mathbb{E}[|X_n|]$ is bounded (i.e. absolutely integrable). Then there exists a finite integrable random variable such that:

$$\lim_{n \to \infty} X_n = X_\infty \;\; a.e.$$

Requires Proof.

Corollary 1

A positive supermartingale converges to a finite r.v.

Requires Proof.

Some Properties

(P1) A random variable X is $\sigma(Y)$-measurable if and only if it is a function of Y.

Requires Proof.

(P2) If X is $\mathcal{G}$ measurable,

$$\mathbb{E}[X|\mathcal{G}] = X \; a.e.$$

Requires Proof.

(P3) If $\mathcal{F}_1 \subset \mathcal{F}_2$ then

$$\mathbb{E}[\mathbb{E}[X|\mathcal{F}_1]|\mathcal{F}_2] = \mathbb{E}[\mathbb{E}[X|\mathcal{F}_2]|\mathcal{F}_1] = \mathbb{E}[X|\mathcal{F}_1]$$

Requires Proof.

Appendix B: Details on the paper's Lemma and Derivation

B1: Statement Without Proof: Neveu's Martingale Convergence Theorem (MCT)(Lemma A.3)

Let $\{T_n\}, \{\alpha_n\}, \{\beta_n\}$ be sequences of non-negative random variables adapted to and increasing sequence of σ-algebras $\mathcal{F}_n$ such that

$$\mathbb{E}[T_n|\mathcal{F}_{n-1}] \leq T_{n-1} - \alpha_{n-1} + \beta_{n-1}$$

If $\sum_1^\infty \beta_n < \infty \; a.s.$, then$T_n$ converges almost surely to a finite random variable T and $\sum_1^\infty \alpha_n < \infty \; a.s.$

Proof Sketch:

Corollary 1 (above) implies that a non-negative supermartingale converges with probability one.

$$\mathbb{E}[T_n|\mathcal{F}_{n-1}] \leq T_{n-1} + \beta_{n-1}$$

With the above requirement on $\beta's$ can be seen as a supermartingale with some diminishing terms (due to their sum converging and their non-negativity) added to it. It is thus expected that they should not affect convergence. Hence,

$T_n \rightarrow T \; a.s.$ with T a non-negative random variable (requires proof)

Let's now define:

$$Z_n = T_n + \sum_{l=1}^{n-1} \alpha_l$$

Thus:

$$\mathbb{E}[Z_n|\mathcal{F}_{n-1}] = T_{n-1} + \beta_{n-1} + \sum_{l=1}^{n-1} \alpha_l$$

Notice that sine $\alpha_n$is non-negative:

$$\sum_{l=1}^{n-2} \alpha_l \leq \sum_{l=1}^{n-1} \alpha_l$$

Hence:

$$\mathbb{E}[Z_n|\mathcal{F}_{n-1}] \leq T_{n-1} + \beta_{n-1} + \sum_{l=1}^{n-1} \alpha_l - \alpha_{n-1} = Z_{n-1} + \beta_{n-1}$$

Again using the convergence argument presented above, we have $Z_n \to Z$ a.s., Z non-negative. Which forces

$$\sum_{1}^{\infty} \alpha_n < \infty \; a.s.$$

Which concludes our proof sketch.

B2: Following Along the Proof:

B2.1

To properly condition Equation 12 on $\mathcal{F}_{t-1}\sigma\{x_0, w_1, w_2, \dots, w_{t-1}\}$, there are a few points to note:

- By definition, b_{t-1} is a function of $[y_{t-1}, \dots, y_{t-n}, u_{t-1}, \dots, u_{t-m}, \dots, -y^*_{t-1}, \dots - y^*_{t-l}]$, and by (P1) it is thus measurable with respect to an algebra included in $\mathcal{F}_{t-1}$, and so measurable w.r.t. $\mathcal{F}_{t-1}$.
- By a similar argument (this time using equation4.6 from the paper), e_t-v_t is $\mathcal{F}_{t-1}$ measurable as well.
- By property (P2) we then have , $\mathbb{E}[z_{t-1}|\mathcal{F}_{t-1}] = z_{t-1} \; a.e$ and $\mathbb{E}[V_{t-1}|\mathcal{F}_{t-1}] = V_{t-1} \; a.e.$
- $\mathbb{E}[v_t|\mathcal{F}_{t-1}] = \mathbb{E}[y_t|\mathcal{F}_{t-1}] - \mathbb{E}[\mathbb{E}[y_t|\mathcal{F}_{t-1}]|\mathcal{F}_{t-1}] = 0$ by linearity of expectations and (P3) on the second term.
- $\mathbb{E}[v_t(e_t - v_t)|\mathcal{F}_{t-1}] = (e_t - v_t)\mathbb{E}[v_t|\mathcal{F}_{t-1}] = 0$, by (P2)
- $\mathbb{E}[(v_t)^2|\mathcal{F}_{t-1}] = \gamma^2 \; a.s.$ (see first part of the paper)

B2.2 Kronecker's Lemma

If $(x_n)_{n=1}^{\infty}$ $is\ a\ convergent\ infinte\ sequence\ of\ real\ numbers$. For some $0 \leq r_0 \leq r_1 \leq \dots$ and $r_n \to \infty$ then we have the following:

$$\lim_{n\to\infty} \frac{1}{r_n} \sum_{k=1}^{n} r_k x_k = 0$$

Requires proof.

For an accessible proof consult Wikipedia article on Kronecker's Lemma.

B2.3 Clarifications on Lemma A.1

Following the paper and using simple back substitution it is straightforward to obtain

$$h_t = CA^t x_0 + Dz_t + \sum_{j=1}^{t} CA^{j-1} B z_{t-j}$$

To obtain the next inequality we use the following trick:

Let $h_t = \alpha + \beta + \gamma$, where $\alpha, \beta, \gamma \in \mathbb{R}^{S\,X\,1}$, then

$$\|h_t\|^2 = \|\alpha\|^2 + \|\beta\|^2 + \|\gamma\|^2 + 2\langle\alpha,\beta\rangle + 2\langle\alpha,\gamma\rangle + 2\langle\beta,\gamma\rangle$$

Now, obviously we have $\|\alpha - \beta\|^2 + \|\gamma\|^2 \geq 0$

This, by symmetry and norm distributivity, gives:

$$\|\alpha\|^2 + \|\beta\|^2 + \|\gamma\|^2 \geq 2\langle\alpha,\beta\rangle$$

One similarly obtains:

$$\|\alpha\|^2 + \|\beta\|^2 + \|\gamma\|^2 \geq 2\langle\alpha,\gamma\rangle$$

$$\|\alpha\|^2 + \|\beta\|^2 + \|\gamma\|^2 \geq 2\langle\beta,\gamma\rangle$$

Thus, replacing each mixed term in the initial $(h_t)^2$ equation, we obtain:

$$\|h_t\|^2 \leq 3\left(\|cA^t x_0\|^2 + \|Dz_t\|^2 + \left\|\sum_{j=1}^{t} CA^{j-1}Bz_{t-j}\right\|^2 \right)$$

One extra step:

$$\left\|\sum_{j=1}^{t} CA^{j-1}Bz_{t-j}\right\|^2 \leq \left[\sum_{j=1}^{t} \|C\|\left\|A^{j-1}\right\|\|B\|\left\|z_{t-j}\right\|\right]^2$$

By repeated Triangle Inequality and Cauchy-Schwartz Inequality. Thus

$$\|h_t\|^2 \leq 3\left(\|cA^t x_0\|^2 + \|Dz_t\|^2 + \left[\sum_{j=1}^{t} \|C\|\left\|A^{j-1}\right\|\|B\|\left\|z_{t-j}\right\|\right]^2 \right)$$

The rest of the proof falls pretty easily as in the original paper. We do note the following:

$\lambda^{j-1} = \lambda^{-1}\lambda^{\frac{j}{2}}\lambda^{\frac{j}{2}}$ instead of $\lambda^{j} = \lambda^{\frac{j}{2}}\lambda^{\frac{j}{2}}$. This ultimately has no incidence as $\lambda^{-1}$is a constant and gets absorbed in K_3.

Appendix C: Simulation Results

The below simulations show the result of the algorithm attempting make the arbitrary system:

$$x_{t+1} = 0.75x_t + 0.9u_t + 0.5w_t$$

$$y_t = 2x_t + w_t$$

$$w_t \; i.i.d. \; N(0{,}0.25^2)$$

track two reference sequences. Notice the "learning" period at the beginning.

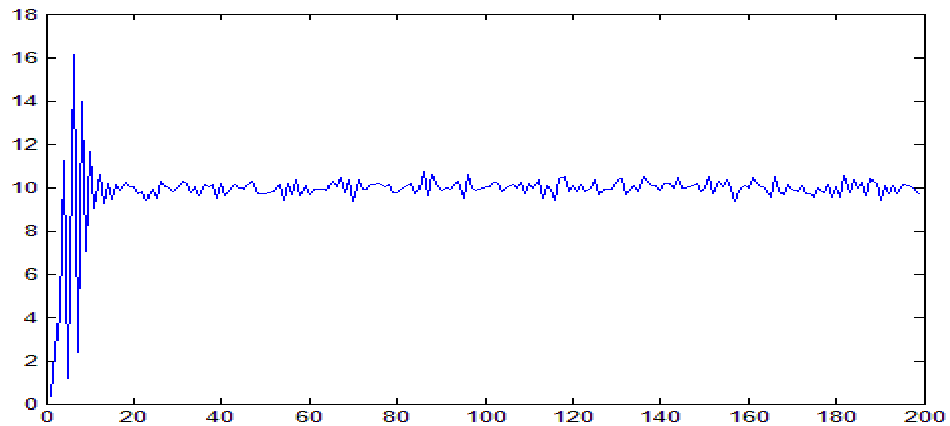

Figure 1: Tracking a fixed reference signal of 10 for 200 time steps.

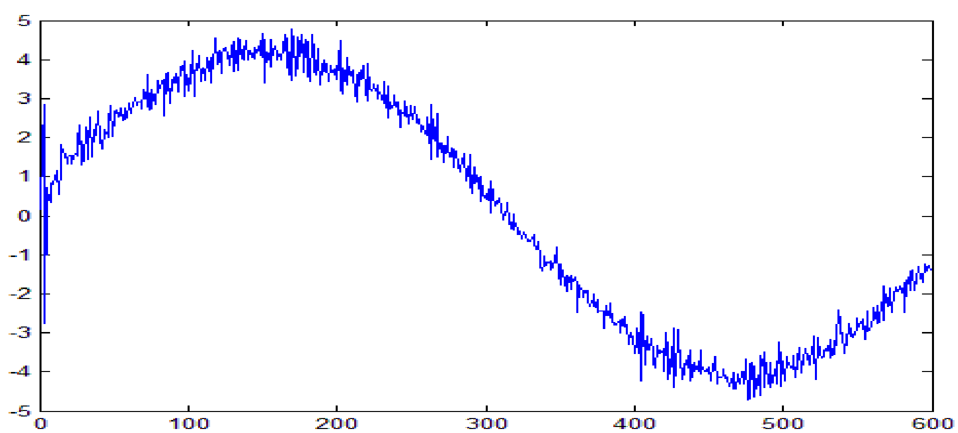

Figure 2: Tracking a sinusoidal reference.

References

[1] Goodwin G.C., Ramadge P. J., Caines P.E., “Discrete Time Stochastic Adaptive Control” SIAM J. Control and Optimization, Vol 19, No. 6, November 1981

[2] Notes on Elementary Martingale Theory by professor John B. Walsh (UBC), retrieved from http://www.math.ubc.ca/~walsh/walsh.html

[3] Wikipedia, Online Encyclopedia, http://en.wikipedia.org/wiki/Kronecker%27s_lemma

[4] Williams D. (1991) Probability with Martingales, Cambridge: Cambridge University Press

[5] A Brief Primer ON Conditional Expectation, retrieved from professor Ben Green’s website: http://www.dpmms.cam.ac.uk/~bjg23/ergodic/conditional-expect.pdf