A Study of Incremental Redundancy Hybrid ARQ over Markov Channel Models Derived from Experimental Data

Beatrice Tomasi*, Paolo Casari*, Leonardo Badia‡, Michele Zorzi*

*Department of Information Engineering, University of Padova — Via G. Gradenigo, 6/B, Padova, Italy
‡IMT Institute for Advanced Studies — Piazza S. Ponziano 6, Lucca, Italy
{tomasibe,casarip,zorzi}@dei.unipd.it   leonardo.badia@imtlucca.it

ABSTRACT

In this paper, we process channel Signal-to-Noise-Ratio time series gathered in the proximity of the Pianosa island, Italy, in Summer 2009. These traces are used to model the performance of capacity-achieving code ensembles as employed in an Incremental Redundancy (IR) Hybrid Automatic Repeat reQuest (HARQ) error control scheme. We apply a code-matched channel state quantization technique aimed at representing channel evolution over time with low quantization error; the evolution of the channel among the quantized states is then represented using a Markov model, over which we base the analytical evaluation of IR-HARQ performance.

Results confirm that IR-HARQ consistently improves link performance with respect to Type I HARQ. In addition, we observe that the different channel statistics due to different transmitter and receiver placements, as well as to the acoustic propagation conditions considered in our scenario, have an impact on HARQ performance. This impact is correctly captured by our Markov model, suggesting good adherence of the model to actual channel behaviors. The validation of the models (by simulating over different traces than those used to train the models) suggests that they are robust to moderate non-stationarity, making them good candidates to give a compact representation of the channel behavior, e.g., in network simulators.

Categories and Subject Descriptors

C.2.0 [**Communication/Networking and Information Technology**]: General—*Data communications*; I.6.4 [**Simulation and Modeling**]: Model Validation and Analysis

General Terms

Experimentation, Measurement, Performance, Theory, Verification

Keywords

Underwater channel measurements, Markov models, hybrid ARQ, analysis, simulation, validation

1. INTRODUCTION

Underwater acoustic networks have received considerable attention from the research community in recent times. On one hand, they represent an interesting scenario for several applications related to marine and oceanic monitoring and surveillance. On the other hand, being based on acoustic communications, they involve several challenges, foremost the propagation medium, which is both extremely hard to capture from the modeling standpoint and also quite difficult to utilize in an efficient manner.

Adopting a network perspective, i.e., translating the problem from a simple communication between a single transmitter-receiver pair to a set of interconnected devices, is not an easy task. To this end, proper medium access control (MAC) strategies must be defined, and further problems arise when dealing with routing and transport issues. While physical layer modeling, in spite of the aforementioned hurdles, has received significant attention, only recently has the research community begun to address these higher layer challenges [1, 2, 3, 4].

One particular problem which arises when dealing with applications for underwater networks is that of channel code design. The inherent unreliability of the underwater channel (even worse than that of the terrestrial radio channel) is often the cause of heavy loss of information and high error correlation. Additionally, acoustic communication is affected by long delays, which make the impact of errors and packet losses even more acute. Indeed, the presence of long delays plays against adopting simple strategies such as the repeated retransmission of corrupted data until success. Thus, underwater communication requires, possibly more than other traditionally lossy channels such as the terrestrial radio or the satellite channel, the employment of effective error control. The two basic techniques over time-varying channels are Forward Error Correction (FEC), which is based on the use of error-correction codes without any retransmission, and pure ARQ, which adopts retransmissions without coding. Both of them have pros and cons; we believe that the best strategy is to take a combined approach, so as to realize what is usually referred to as a hybrid ARQ (HARQ) scheme [5]. This solution, if properly designed, is able to combine the best of both techniques, namely to protect the data sent with error correcting codes already from the first transmission, yet allowing retransmissions if needed; this spares the need for introducing unnecessarily high amounts of redundancy in the codeword, which would be inefficient in underwater channels.

We believe that a deep understanding of HARQ techniques over the acoustic channel is necessary in order to properly model the performance of the applications of interest. In the literature, Markov chains are generally recognized as a useful means to represent wireless channels and analytically characterize techniques such as ARQ

/ HARQ operating on top of them [6, 7]. Many papers evaluate the performance of ARQ or HARQ over a Markov channel: for example, the authors of [8] analyze HARQ schemes representing a fading channel through a Markov chain, whereas in [9] the delay statistics of HARQ are evaluated when both the arrivals and the channel error process are Markov.

The approach considered in this paper is different from these related works, since we aim at identifying whether a Markov representation may be appropriate for the underwater acoustic channel, when HARQ is used on top of it, and in this case how the parameters of the Markov chain should be set.

Especially, a recognizable risk of Markov models for HARQ techniques is that they require a detailed model of the system. To take into account both coding and retransmissions, several quantities have to be tracked. Therefore, the resulting model is often considered to be intractable and cumbersome. In fact, in this paper we aim at disproving this misconception, by finding a Markov representation of practical value.

In more detail, our main contribution is the application, and validation with experimental trials, of a model for HARQ performance based on Markov models, which has been proposed in [10], with the aim of considerably simplifying the channel representation.

Such a general model finds suitable application in the underwater case, especially in view of the difficulties discussed above, since it enables a considerable simplification of the analysis in a precise and exact manner. More specifically, in this paper we will present both an analytical framework and supporting experimental evidence that show that the statistics of error correction techniques based on HARQ and run over an acoustic channel can be characterized by means of relatively simple Markov chains. In more detail, we will proceed as follows. We start from processing channel transmission data traces gathered during the SubNet'09 sea trials off the coast of the Pianosa island, Italy, and we consider these transmissions as punctured versions of the codewords of capacity-achieving codes [11]. By processing Signal-to-Noise-Ratio (SNR) time series,[1] we characterize the distribution of the receiver SNR to be approximately Gaussian and correspondingly derive the reliable region of the code. We then quantize channel states according to the methodology proposed in [5], i.e., by fixing a number of thresholds on the receive SNR levels, which are then used to estimate the distribution of the SNR for a subsequent transmission, conditioned on the present channel state. We finally evaluate the throughput performance of an Incremental Redundancy Hybrid ARQ (IR-HARQ) scheme by means of the estimated channel model, and compare its performance to that of Type I HARQ.

Our results are useful to assess the goodness of a Markov approximation to the evolution of the channel state when capacity-achieving codes are employed. In addition, our comparison of different links (i.e., between transmitters and receivers placed at different locations and depths) shows that the Markov model correctly captures different channel statistics, which translate into different performance of the underlying HARQ scheme, suggesting good adherence of our system model to the actual channel behavior. Such models may eventually be embedded into more complex network

---

[1]The statistics of the SNR are not the only channel property affecting communications performance: e.g., time-varying multipath patterns and the time spread of the resulting channel impulse response also have an impact on performance. However, in the presence of sufficiently powerful receiver-side signal processing, or of modulation schemes robust against multipath such as a frequency-hopping binary frequency shift keying, these phenomena may be compensated for or at least mitigated. Therefore, for simplicity we will concentrate on SNR statistics in this paper, neglecting other environment-induced propagation effects.

Figure 1: A scheme of the SubNet 2009 testbed deployment, East of the Pianosa Island.

simulators, in order to provide a compact and computationally efficient method of reproducing the statistics of links showing similar environmental characteristics (depth, distance, sound speed profile, etc.) to those studied in this paper.

2. EXPERIMENTAL SETUP

The SubNet 2009 sea trials were organized off the eastern shore of the Pianosa island, Italy ($42.585°$N, $10.1°$E), in order to serve as an experimental demonstration of JANUS [12], a signal format and transmission protocol designed for unsolicited broadcasting of information relevant to nautical vessels. The signal is composed of a preamble (either a hyperbolic frequency modulated sine wave or a sequence of 30 symbols taken from a predefined frequency-hopping pattern), employed as a probe for timing synchronization and for SNR estimation, and actual data transmission, using a 13-subcarrier Frequency-Hopping Binary Frequency Shift Keying (FH-BFSK) modulation with a predefined hopping pattern and a hopping rate of 1 hop per transmitted symbol. The data portion of the JANUS signal is further divided into a 144-bit header and an optional payload: for the purpose of our performance evaluation, we focus on transmissions bearing no payload.

The testbed used during SubNet 2009 consisted of one vertical array (VA) of hydrophones at different depths (of which we consider H1, H2 and H4, respectively located at a depth of 20, 40, and 80 m), and three Teledyne-Low Frequency acoustic modems [13], each placed on a tripod located on the sea floor at different depths (60 to 80 m) and distances from the VA. For reference, a scheme of the testbed and the sea trials location is depicted in Fig. 1. The three transmitters have been labeled T1 (1500 m from the VA, depth 60 m), T2 (2200 m from the VA, depth 70 m) and T3 (700 m from the VA, depth 80 m). Oceanographic parameters were sampled using a thermistor chain and an acoustic Doppler current profiler, both located close to the VA. In particular, the thermistor chain was designed for finer sampling in the mixed layer (i.e., at low depth) rather than in the lower layers, in order to better track temperature changes between 10 and 40 m of depth. Temperature is then used to infer sound speed profile (SSP) time series by means of such empirical equations as the Mackenzie formula [14], and by considering salinity measurements taken at the beginning of the campaign as approximately constant throughout the season.

Using this deployment, many experiments were conducted, lasting up to ten hours and involving several thousand JANUS transmissions, at different times of day and at different days along the

summer season. The trials we are focusing on in this paper took place between the end of May and the end of August 2009, and include more that 12000 transmissions, in different channel conditions, leading to different propagation effects such as multipath patterns [15]. In order to focus on a representative set of traces, we consider an experiment which took place on June 5. During this experiments, signals were transmitted once every 15 s for 9 hours over a downward-refractive channel [15].

3. CHANNEL MODEL

For the present discussion, we consider transmissions subject to an HARQ error control mechanism described as follows. Nodes transmit information frames, where each frame is actually derived from a long codeword of a low-rate code. Each frame is composed of multiple HARQ packets, or fragments, which for simplicity are assumed to be all of the same size. In addition, we assume that each fragment, if correctly received, is sufficient to recover the whole codeword. Every time a packet is sent, the receiver replies with a feedback ACK or NACK message, respectively indicating correct or incorrect packet reception. In the following, we will consider both Type I and Type II HARQ [16]. In Type I HARQ, only one coded fragment is sent per information frame, actually resulting in a FEC strategy to protect the frame against errors; in case a transmitter receives a NACK, it provides a retransmission of the same HARQ fragment. In Type II HARQ, instead, every information frame is associated to multiple HARQ fragments, and NACKs trigger the transmission of a new fragment each time: therefore, subsequent detection attempts are based on the availability of additional redundancy, and the corresponding scheme is usually referred to as Incremental Redundancy HARQ (IR-HARQ).

In what follows, we will match the transmission SNR traces gathered during the SubNet'09 campaign with the HARQ framework described above, by assuming that each packet transmission actually carries one HARQ fragment formed by resorting to "good" LDPC code ensembles such as those considered in [17]. The performance of these codes can be characterized from an information-theoretic point of view in terms of SNR thresholds: for a single transmission, a single threshold can be found that determines if decoding is immediately successful or not; for multiple transmissions, each bearing a different SNR in general, a *reliable* region can be defined as outlined in the following section [17].

The purpose of the following analysis is to determine reliable SNR regions for LDPC transmission over the links between the acoustic modems and the hydrophones. We remark that these depend on channel conditions, which in these cases are those experienced during the experiments listed at the end of the previous section. Starting from this analysis, we estimate a Markov model of the channel following the guidelines in [10] for the optimal quantization of the reliable SNR region and use the model for characterizing the performance of an IR-HARQ scheme based on the discussed LDPC code.

3.1 Reliable SNR regions

The reliable region model for characterizing the performance of good code ensembles [11, 17] assumes that multiple fragments are transmitted sequentially, and that at the kth transmission decoding is based on all k fragments received so far, each bearing its own SNR value $s_1, \ldots, s_k$. The reliable region $\mathcal{R}(k)$ is defined as the subset of $\mathbb{R}^k$ containing the k-tuples of SNR values for which the decoding failure probability asymptotically vanishes as the codeword length increases. Any reliable region $\mathcal{R}(k)$ has the property that if a k-tuple of SNR values $(s_1, \ldots, s_k) \in \mathcal{R}(k)$ then the k-
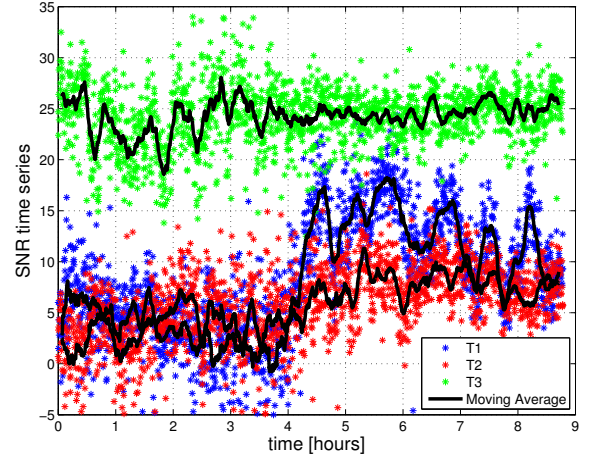
Figure 2: Measured time series of the SNR over the links from all transmitters to hydrophone H1. A moving average taken over 25 samples is superimposed to the SNR time series as a solid black line.

tuple $(s_1, \ldots, s'_k) \in \mathcal{R}(k)$ as well for every $s'_k > s_k$ [17]. This directly follows from the fact that if s_k suffices to enable correct decoding, any greater s'_k would suffice as well. In turn, this makes it possible to define the reliable region using a threshold model, where the minimum value the SNR of the kth transmission should have to ensure successful decoding depends on the sequence $\mathbf{s}^{(k-1)}$ of all previous $k-1$ SNR values. Such a threshold takes the form

$$\vartheta(\mathbf{s}^{(k-1)}) = \inf\{s_k : (s_1, \ldots, s_{k-1}, s_k) \in \mathcal{R}(k)\} \qquad (1)$$

which can be used to verify whether correct decoding occurred at the kth transmission, i.e., by checking if $s_k \geq \vartheta(\mathbf{s}^{(k-1)})$ or not.

If the joint probability density function (pdf) or, equivalently, the joint distribution of the SNRs is known, the probability distribution of a given SNR k-tuple can be derived; analogously, SNR regions in $\mathbb{R}^k$ can be mapped into probability regions using the cumulative distribution function (cdf) of the SNR [10]. The pdf of the SNR can be derived by fixing a link and analyzing the SNR time series over that link throughout the duration of an experiment. An example of an SNR time series is shown in Fig. 2, depicting the time evolution of the SNR over the links from all transmitters to hydrophone H1. A moving average of the time series taken over 25 samples is also shown as a solid black line. From the figure we see that the T3–H1 link (T3 is the closest to H1) experiences high SNR which is also quite stable over time, despite some events at the beginning of the experiment where the SNR drops 5 to 10 dB below its overall average value. These events are mainly due to environmental phenomena such as currents, causing temporary drops in the temperature of the upper water layers, which in turn affect how the acoustic energy propagates. We recall that H1 is placed at a depth of 20 m, and is therefore more vulnerable to changes in superficial layer propagation parameters than other hydrophones. Nevertheless, the average SNR level is quite high over the T3–H1 link, as the dominating effect here is the short distance (and the consequent low attenuation) between T3 and H1. A different behavior is observed over the T1–H1 and T2–H1 links, respectively the intermediate and longest distance links. In particular the larger distance causes the variance of the SNR to increase, and makes the links more sensitive to oceanographic phenomena, as the average value of the T1–H1 SNR oscillates from roughly 5 dB down to 0 dB or
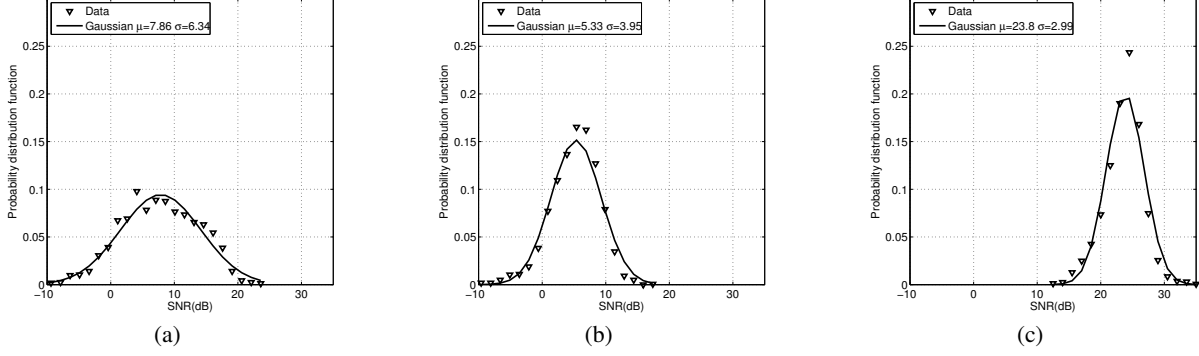
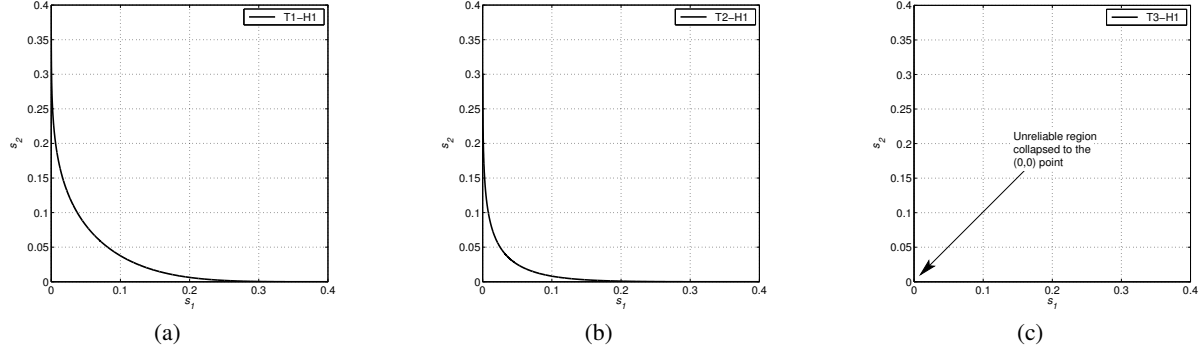Figure 3: Empirical pdf of the SNR and Gaussian fit over the T1–H1 link (a) T2–H1 link (b) and T3–H1 link (c).

Figure 4: Reliable regions for the T1–H1 link (a) T2–H1 link (b) and T3–H1 link (c).

less and then increases again up to $15\,\mathrm{dB}$. While a more accurate model of SNR distribution would consider the changes in the SNR average value and variance over time, it would also complicate the analysis of the system below, without yielding significantly better insight. Therefore, in the following we will take a simpler approach and estimate the statistics of the SNR as if they were stationary over the whole duration of the experiments. However, we will also validate the model by training it over a subset of the SNR traces and comparing the results to simulations carried out over a different portion of the same experiment.

In this light, Figs. 3(a), 3(b) and 3(c) show the empirical pdf of the SNR over the links from all transmitters (T1 to T3) to receiver H1 (the hydrophone placed at a depth of $20\,\mathrm{m}$). Each figure also shows a fit performed using a Gaussian pdf of the form

$$f_{\Gamma}(\gamma) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(\gamma-\mu)^2}{2\sigma^2}\right), \qquad (2)$$

where the average value μ and variance σ^2 have been estimated to best fit the data in a least-squares sense; both μ and σ are reported in the legend. From (2), the probability that the SNR lies in any interval $[a, b]$ is then found straightforwardly as $\int_a^b f_{\Gamma}(\gamma)\,\mathrm{d}\gamma$. Starting from the SNR distribution, the reliable region of good LDPC code ensembles can be derived using the threshold model mentioned before and following the approach described in [17]. The reliable regions in the probability domain for links from T1, T2 and T3 to H1 are depicted in Figs. 4(a), 4(b) and 4(c), respectively, by considering two subsequent HARQ fragments. The reliable region lies to the upper right of the boundary curve shown as a solid black line; note that s_1 and s_2 represent the values taken by the cumulative distribution functions of the SNR of the first and second transmissions, respectively. Note that the reliable region is larger when the

SNR distribution has higher average and lower variance. In particular, due to the very high average SNR experienced by the T3–H1 link, the reliable region tends to occupy the whole space, and the unreliable region correspondingly collapses into the $(0, 0)$ point.

Figs. 4(a) to 4(c) refer to a specific experiment, and to a single receiver, but are representative of all other experiments and links between T1–T3 and the hydrophones: in fact, a generally good accordance between the SNR pdf and a proper Gaussian fit (albeit with different μ and σ for different experiments and links) was found to hold in all cases.

On top of the reliable region model, we construct a channel model following the guidelines for channel state quantization described in [10]. The procedure is briefly summarized in the following. Before channel quantization, a vector of channel states has the form $\mathbf{s}^{(k)}$, as defined above, where each element can take infinitely many values in $\mathbb{R}$. Quantizing channel states translates $\mathbf{s}^{(k)}$ into a k-tuple of discrete values, that evolve according to a Finite-State Markov Chain (FMSC). If N thresholds $\alpha_1, \ldots, \alpha_N$ divide $\mathbb{R}$ into $N + 1$ intervals $I_0, \ldots, I_N$, where $I_j = [\alpha_j, \alpha_{j+1}[$, and we define $\alpha_0 = 0$ and $\alpha_{N+1} = +\infty$, any real value becomes mapped into the discrete index j of the interval I_j it falls within. To formalize this mapping, define $d(s_k)$ as the function returning the interval of $\mathbb{R}$ where s_k is contained, i.e., $d(s_k) = j$ if $s_k \in I_j$. Now, by grouping the mappings for all elements of an SNR k-tuple $\mathbf{s}^{(k)}$ into a vector, we can write $\mathbf{d}^{(k)} = (d(s_1), \ldots, d(s_k))$, thereby establishing a map between every k-tuple $\mathbf{s}^{(k)}$ and an element of the set $\mathbb{Z}_{N+1}^k$, where $\mathbb{Z}_{N+1} = \{0, 1, \ldots, N\}$. In more detail, the vector $\mathbf{d}^{(k)}$ represents the fact that $\mathbf{s}^{(k)} \in I_{d(s_1)} \times I_{d(s_2)} \times \cdots \times I_{d(s_k)} = \mathcal{I}(\mathbf{d}^{(k)}) \subset \mathbb{R}^k$.

An FSMC channel model entails the assumption that the statistics of the SNR have the Markov property (a common means of

describing correlated SNR evolution over time [6]); the model can then be derived from the distribution of the SNR, which in this case is given by the Gaussian fitting discussed above. If the pdf of the SNR, denoted with γ, is called $f_\Gamma(\gamma)$, the state space of the channel is $\mathbb{Z}_{N+1}$, and the steady-state distribution is given by

$$\pi_i = \int_{\alpha_i}^{\alpha_{i+1}} f_\Gamma(\gamma)\,\mathrm{d}\gamma, \quad i = 0, \dots, N. \tag{3}$$

Analogously, the probability that a transition between state i and j occurs, t_{ij}, can be derived as follows

$$t_{ij} = \frac{\int_{\alpha_i}^{\alpha_{i+1}} f_\Gamma(\gamma_0) \int_{\alpha_j}^{\alpha_{j+1}} f_\Gamma(\gamma|\gamma_0)\,\mathrm{d}\gamma\,\mathrm{d}\gamma_0}{\int_{\alpha_i}^{\alpha_{i+1}} f_\Gamma(\gamma_0)\,\mathrm{d}\gamma_0}, \quad i, j \in \mathbb{Z}_{N+1} \tag{4}$$

where $f_\Gamma(\gamma|\gamma_0)$ is the conditional pdf of the SNR γ given the previous SNR value γ_0.

The channel transition probability matrix is then defined as $\mathbf{T} = (t_{ij})$ for $i, j \in \mathbb{Z}_{N+1}$. We remark that given the limited size of the available SNR data sets, estimating a close-form fit of conditional distributions may yield little significance. We therefore resort to direct estimation of channel transition probabilities from the data, by taking the relative frequencies of SNR transitions between any two intervals I_i, I_j, where $i, j \in \mathbb{Z}_{N+1}$ represent the indices of the starting and ending intervals respectively.

It is worth noting that the transition probabilities strongly depend on the number of states used to quantize the channel, and therefore on the number of thresholds used for delimiting SNR intervals (or, equivalently, SNR probability intervals). In this paper, we will employ only two thresholds, resulting in a total of three channel states. This choice results in a very simple channel model, but still provides sufficient quantization accuracy, as shown by the numerical results.

4. MODELS FOR HARQ SCHEMES

The previous section focused on describing how a FSMC model of the channel can be derived from experimental data. We now focus on how to employ this FSMC to model an HARQ error control process based on the good LDPC code ensembles the FSMC has been matched to. Recall that $\mathbf{T}$ denotes the transition probability matrix of the FSMC. The FSMC has in general $N + 1$ states $0, 1, \dots, N$, if N thresholds are chosen to quantize the channel behavior. With no loss of generality, assume that 0 is the best state, i.e., the one associated to the highest values of the SNR, whereas N is the worst state. Let us define a map $g(j)$ which associates each state index j, $j = 0, \dots, N$, to an "error level" a packet would incur if transmitted while the channel is in state j. The error level is a non-decreasing function of the state index, and is employed to describe the usefulness of the packet being transmitted for the decoding of the LDPC codeword at the receiver as follows. Recall from the beginning of Section 3 that every information frame is encoded and divided in HARQ fragments, to be transmitted sequentially, and that a single correct HARQ fragment is always sufficient to successfully decode the whole information frame. However, corrupted fragments may still be used at the receiver, according to the type of HARQ scheme, as will be explained in detail later. To model this HARQ feature, we assume that a successful decoding takes place after reception of HARQ fragment k only if the overall error level (which is defined as the error level of fragment k for Type I HARQ, and as the sum of the error levels of all HARQ fragments received so far for IR-HARQ) is lower than or equal to a certain threshold θ_k [5].

With the above in mind, and for a fixed SNR statistics over the link, the performance of the HARQ scheme only depends on the round-trip time of the channel, m, and on the maximum number of retransmissions allowed before an information frame is discarded, denoted by F [5]. Recall from Section 2 that our data set contains transmissions performed once every $15\,\mathrm{s}$, which is much larger than the average propagation delay, given the distances between the transmit and receive hardware. Hence, we have to fix $m = 1$ for our data set, which corresponds to assuming that a slotted approach is taken, whereby each slot is long enough to accommodate the maximum round-trip time (that between T2 and the VA) and the time required for acoustic reverberation to fade out. Taking higher values for m would correspond to assuming that the links span a distance of more than $22.5\,\mathrm{km}$, which would make the measured SNR statistics meaningless. Finally, note that setting $m = 1$ means that transmitted frames are actually sent only after receiving the ACK/NACK feedback related to the previous message, so that at every time instant there is at most one message in flight over the channel (i.e., the ARQ scheme is Stop-and-Wait). Due to the long propagation delay, we also assume that each corrupted packet can be retransmitted only once, i.e., $F = 1$, and if the retransmission also fails the packet is discarded.

In general, at each time we need to keep track of the number of retransmissions already made and of the correspondingly accumulated error level, as well as of the channel state. The resulting model can be derived in general form following the approach described in [5]. In the following subsections we will instead focus on a simpler model for the Type I and Type II HARQ schemes in the specific case $F = 1, m = 1, N = 2$.

4.1 Type II (Incremental Redundancy) HARQ

We start with the description of the Type II HARQ (IR-HARQ) scheme. We employ two thresholds to describe the channel behavior, resulting in three channel states, namely 0 (error-free state), 1 (some recoverable errors), 2 (worst state with unrecoverable errors). In general, each state j is associated to a different error level incurred by a fragment transmitted when the channel is in that state through the map $g(j)$ (see Section 4). According to the above description, we simply set $g(j) = j$.

If the transmission of the first fragment occurs when the channel is in state 0 (error-free state), the frame is correctly received and no retransmission is needed. If instead the first fragment is in error, a new fragment is transmitted, and the frame can still be recovered if the pair of states for the two attempts is $(1, 0)$, $(1, 1)$ or $(2, 0)$, whereas we assume that the situations $(2, 1)$, $(1, 2)$, and $(2, 2)$ correspond to too many errors, and do not lead to successful frame recovery. This corresponds to the set of thresholds $\theta_0 = 0$ and $\theta_1 = 2$.

Hence, we represent the system state with a pair (S, r), where S represents the previous channel state (required to track the evolution of the Markov channel) and r denotes the number of transmissions already made for the current frame, i.e., $r = 0$ for the first fragment and $r = 1$ for a retransmission. In this simplified model, there is no need for an explicit variable tracking the error level of the packet, as it is identical to S. Note also that there are only *five* possible states (S, r), although S and r can take three and two values, respectively. Indeed, the combination $S = 0, r = 1$ is invalid as it would correspond to the retransmission of a correctly received frame.

Define now σ_{Sr} as the steady-state probability of being in state (S, r). If t_{ij} denotes the probability of making a transition from

channel state i to state j, the balance equations between such steady-state probabilities can be found as follows

$$\sigma_{00} = \sum_{i=0}^{2} \sigma_{i0}t_{i0} + \sum_{i=1}^{2} \sigma_{i1}t_{i0} \qquad (5)$$

$$\sigma_{S0} = \sum_{i=1}^{2} \sigma_{i1}t_{iS} \quad \text{for } S = 1, 2 \qquad (6)$$

$$\sigma_{S1} = \sum_{i=0}^{2} \sigma_{i0}t_{iS} \quad \text{for } S = 1, 2 \qquad (7)$$

These equations can be put into a system which can be solved by imposing the additional condition that the sum of all σ_{Sr} equals 1. After solving the system, one can directly compute the throughput, Θ (average number of successful frames per slot), the average number of retransmissions per correctly decoded information frame, N_{fr}, and the probability that a frame is discarded, P_{fd} (i.e., the fraction of frames that are not correctly received), as

$$\Theta = \sum_{i=0}^{2} \sigma_{i0}t_{i0} + \sum_{i=1}^{2} \sigma_{i1}t_{i0} + \sigma_{11}t_{11} = \sigma_{00} + \sigma_{11}t_{11} \,, \quad (8)$$

$$N_{fr} = \frac{\sigma_{11}t_{10} + \sigma_{11}t_{11} + \sigma_{21}t_{20}}{\sum_{i=0}^{2} \sigma_{i0}t_{i0} + \sum_{i=1}^{2} \sigma_{i1}t_{i0} + \sigma_{11}t_{11}} \,, \qquad (9)$$

$$P_{fd} = \frac{\sum_{i=0}^{2} \sigma_{i0}(t_{i1}t_{12} + t_{i2}(1 - t_{20}))}{\sum_{i=0}^{2} \sigma_{i0}} \,. \qquad (10)$$

Eq. (8) is derived by summing the probabilities of all transitions that correspond to a successful frame in a slot, which include all cases in which a transmission occurs in channel state 0, plus the case in which a retransmission in state 1 follows an erroneous transmission that was itself in state 1. (Note the simplification allowed by the balance equation (5).) Eq. (9) is again obtained by enumerating all events that correspond to the successful delivery of a frame in a slot: the sum of the probabilities of these events is the denominator of (9), whereas the numerator is the sum of the probabilities of only those that correspond to a success after retransmission. (Note that the denominator in this case is the throughput Θ.) Finally, Eq. (10) is derived as follows. When the first fragment of a frame is transmitted, we must have $r = 0$, so that the only three possible states are $(S, 0), S = 0, 1, 2$. Given that the channel state is S, the probability that the frame is discarded is the probability that the two transmission attempts have an error level that is not sufficient for decoding, which is equal to $t_{S1}t_{12} + t_{S2}(1 - t_{20})$. The final result is obtained by averaging this probability over the *normalized* distribution of the channel state, $\sigma_{S0}/\sum_{i=0}^{2} \sigma_{i0}$.

4.2 Type I HARQ

As a term of comparison, we also consider a Type I hybrid HARQ where only one HARQ fragment is actually transmitted per information frame, so that encoding provides only some protection against channel errors over a single packet transmission. However, no more than one HARQ fragment is considered at every decoding attempt: therefore, further retransmissions can help the decoding process only by providing more chances to incur a sufficiently high SNR value. This entails the definition of a single threshold which is equal to the noise threshold of the LDPC code ensemble, and leads therefore to two cases: if the SNR is above the threshold, the frame is correctly decoded; otherwise, it is retransmitted once, and if again the SNR threshold is not met, it is discarded. For a fair comparison, we keep the maximum number of retransmissions as $F = 1$ in this case as well.

This situation can be modeled using the same set of balance equations as before. However, in Eqs. (8)–(10) the effect of incremental redundancy must be removed, i.e., two subsequent transmissions of the same frame that both experience channel state 1 no longer yield a correct decoding, and therefore any term that relates to this event is to be counted as a failure rather than a success. In this case we have

$$\Theta = \sigma_{00} \,, \qquad (11)$$

$$N_{fr} = \frac{\sigma_{11}t_{10} + \sigma_{21}t_{20}}{\Theta} \,, \qquad (12)$$

$$P_{fd} = \frac{\sum_{i=0}^{2} \sigma_{i0}(t_{i1}(1 - t_{10}) + t_{i2}(1 - t_{20}))}{\sum_{i=0}^{2} \sigma_{i0}} \,. \qquad (13)$$

5. RESULTS

We now show a comparison of throughput, probability of frame discarding and average number of HARQ fragments per correctly decoded frame as derived by the analysis discussed above. The analysis is compared to simulation results obtained by reproducing the evolution of the Type I and II HARQ over the SNR traces employed to derive the Markov models.

All results are plotted as a function of the average transmit power: for each point on the curve, SNR traces are offset in order to simulate a different transmit power; in turn, this changes the average value of the distribution of the SNR, affecting the size of the reliable code region, and the performance of the decoding process. All other parameters are set as discussed above.

5.1 Training over a complete SNR trace

Figs. 5, 6 and 7 show the throughput Θ, average number of retransmissions N_{fr} and probability of frame discarding P_{fd}, respectively. The T1–H1 and T3–H1 links are considered. For this evaluation, the Markov models have been trained over the whole SNR traces of the experiment: this is meant as a general sanity check for the use of such models in the context of underwater networks. Later in this section we will also validate the results by training the model over a portion of the dataset in order to obtain analytical results, while performing simulations over a different portion of the traces of the same experiment.

We start by considering throughput in Fig. 5. The curves show two expected behaviors, namely i) that the performance of error control schemes improves with higher average SNR (which increases the length of sojourns in favorable states of the Markov model); ii) that Type I HARQ is consistently outperformed by Type II HARQ which is in line with common wisdom (e.g., see [9]).

In addition, we observe that the analysis fits simulation results quite accurately; this suggests that Markov models are in fact a good choice to achieve a high-level representation of the channel behavior as they can correctly reproduce the statistics of the SNR process.

The SNR distribution significantly impacts the behavior of the curves: in fact, the slope of the transitional portion of the curve (when throughput increases from 0 to 1) is steeper for more stable links, and milder for links exhibiting greater variance. For example, compare HARQ performance with the behavior of SNR over time as seen in Fig. 2: note that while T3 experiences high, fairly stable SNR, T1 is subject to an initially lower SNR that however increases after roughly 4 hours since the beginning of the experiment. Nevertheless, due to the greater dispersion of data over a larger interval of SNR values throughout the experiment, T1's SNR distribution has higher variance than T3's. Therefore, the throughput increase over the T3–H1 link is steeper, and T3's throughput eventually tops
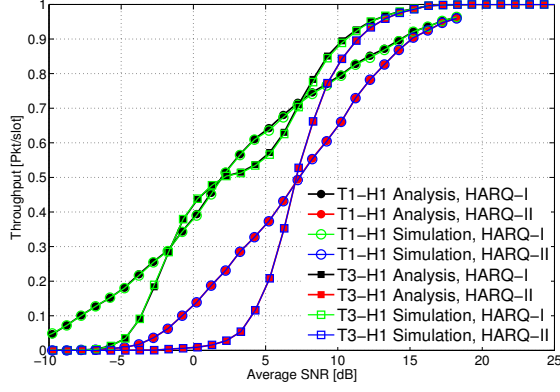
Figure 5: Throughput Θ as a function of the average SNR for links T1–H1 and T3–H1.

T1's. While not shown here due to lack of space, other links (e.g., from T1 and T3 to H4) show generally similar trends. In particular, links from closer transmitters consistently exhibit better performance, and in addition more stable links show steeper transitions in the throughput curves.

Consider now Figs. 6 and 7 showing the number of retransmissions and the probability of frame discarding, respectively, for the T1–H1 and T3–H1 links. The general trends, in terms of the relative performance of the two links, reflect that observed for throughput. However, Fig. 6 shows a further difference between the Type I and Type II HARQ policies: as long as the average SNR is low, the number of retransmissions incurred by Type I HARQ is 0.5, whereas that of Type II HARQ is 1. In fact, N_{fr} is conditioned on having correctly decoded an information frame: in unfavorable channel conditions this event is so rare that Type I HARQ experiences a success during the first fragment transmission or during the first retransmission with equal probability. With Type II HARQ, instead, the transmission of a second fragment always improves the decoding performance, making it more likely to successfully decode after the first retransmission, thereby shifting N_{fr} to 1. Fig. 7 supports this interpretation by showing that Type II HARQ indeed yields a lower probability of frame discarding.

5.2 Model validation

In all figures discussed above, the simulation results match the analysis quite accurately. This is a consequence of having trained the Markov models over the full SNR traces, and of having performed simulation runs over the same traces. In this subsection, we validate the accuracy of the models by training over a different portion of data set than used in the simulation. In any event, all portions belong to the same experiment. We consider first a mildly non-stationary case, whereby the model is trained over the portion from one half to three quarters of the SNR time series obtained at H1, whereas simulations are run on the last quarter of the same time series. We observe from Fig. 2 that, e.g., the T3–H1 link is very stationary during this portion of the experiment, while T1 undergoes broader changes as the average SNR oscillates and tends to decrease toward the end of the experiment. This case is covered by Fig. 8, showing throughput, number of retransmissions and probability of frame discarding. The limited amount of variations in the SNR traces from the portion used in the analysis and that used in simulations reflects into a very good agreement between simulation and analysis for both the T3–H1 link and the T1–H1 link. The only slight disagreement comes from T1–H1's N_{fr} metric (Fig. 8(b)) for
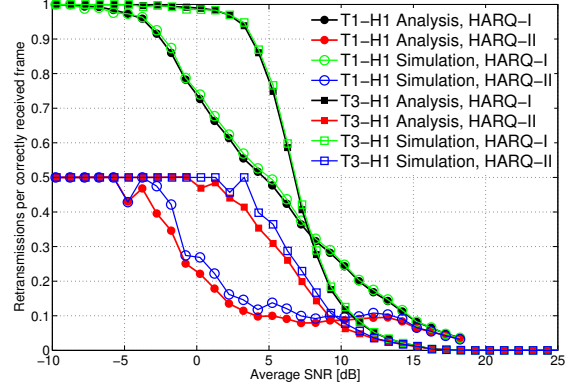
Figure 6: Number of retransmissions per correctly received frame, N_{fr}, as a function of the average SNR for the links T1–H1 and T3–H1.
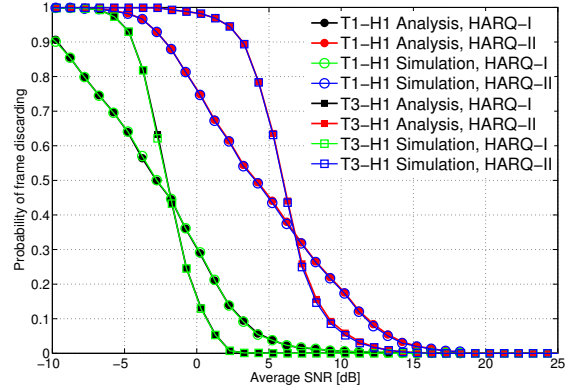

Figure 7: Probability of discarding a data frame, P_{fd}, as a function of the average SNR for links T1–H1 and T3–H1.

an average SNR of 1 to 4 dB, and is due to an insufficient number of correctly received packets, leading to lower statistical significance of the expectations performed to calculate the value of the metric.

For comparison, we also considered a non-stationary case, whereby the model is trained over the first half, whereas simulations are run on the second half of the SNR time series of H1. In this case, both the T3–H1 and the T1–H1 links experience variations, which are limited for T3–H1, but of greater entity for T1–H1, whose average SNR increases by about 10 dB in the second half of the experiment, with respect to the first half. This case is covered by Fig. 9, where we observe that the model is acceptably accurate for the T3–H1 link, and less accurate for T1–H1. In any event, despite some disagreement, the accuracy is still acceptable for the model to work at least as an approximation of the link behavior. This may still make it suitable, e.g., to be implemented in a network simulator.

6. CONCLUSIONS

In this paper we applied, and validated with experimental data, a model for HARQ performance based on Markov chains. The study is based on the processing of SNR traces extracted from transmissions performed during a campaign of sea trials in 2009. Starting from these traces, we have extracted channel statistics and used them to derive reliable SNR regions for good code ensembles. We have then estimated a Markov model that tracks the evolution of the
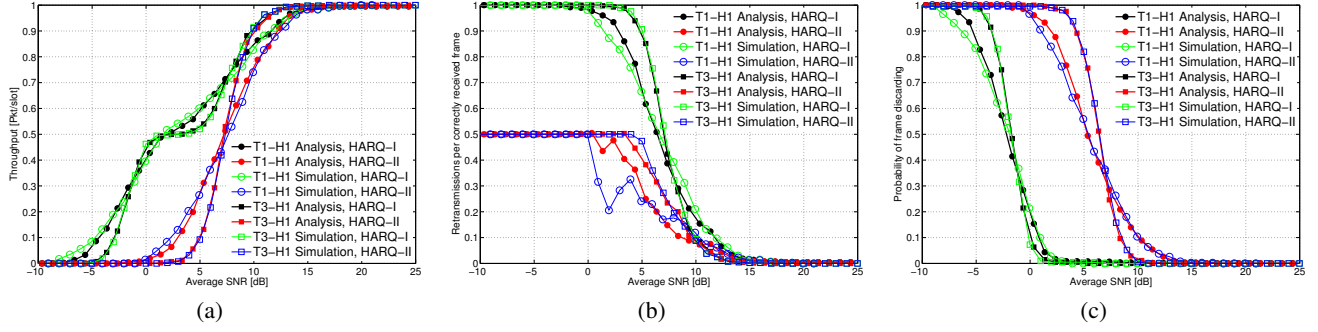
Figure 8: Model validation, mildly non-stationary case: model trained over the *third quarter*, simulations run over the *fourth quarter* of the SNR time series. (a) Throughput, Θ (b) Average number of retransmissions, N_{fr} (c) Probability of frame dropping, P_{fd}.
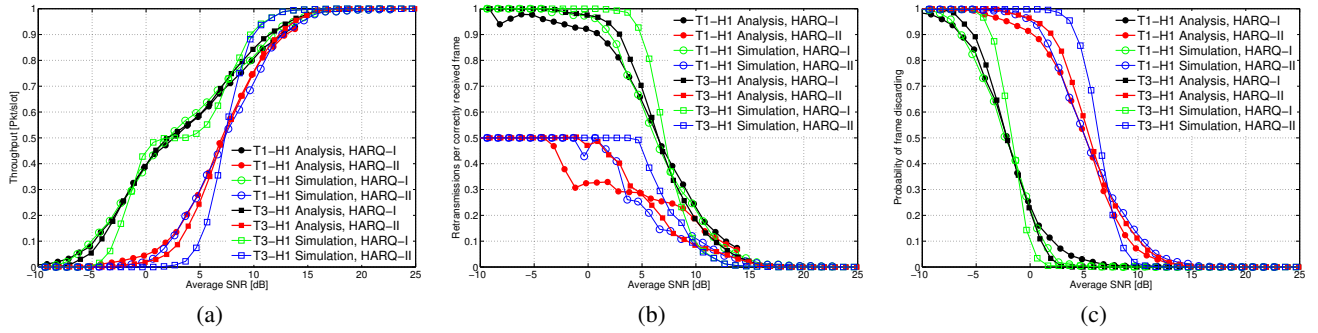
Figure 9: Model validation, non-stationary case: model trained over the *first half*, simulations run over the *second half* of the SNR time series. (a) Throughput, Θ (b) Average number of retransmissions, N_{fr} (c) Probability of frame dropping, P_{fd}.

channel state in time, and finally employed the model to estimate the performance of incremental redundancy (Type II) HARQ, as compared to Type I HARQ. The model has been shown to significantly simplify channel representation, yet to satisfactorily adhere to the actual channel behavior, thereby allowing to characterize the performance of HARQ policies with low complexity, even in the presence of variations of moderate entity in the channel statistics with respect to those used to train the model.

Acknowledgment

7. REFERENCES

[1] B. Peleato and M. Stojanovic, "Distance aware collision avoidance protocol for ad hoc underwater acoustic sensor networks," *IEEE Commun. Lett.*, vol. 11, no. 12, pp. 1025–1027, Dec. 2007.

[2] K. B. Kredo II and P. Mohapatra, "A hybrid medium access control protocol for underwater wireless networks," in *Proc. ACM WUWNet*, Montreal, Quebec, Canada, Sep. 2007.

[3] M. Zorzi, P. Casari, N. Baldo, and A. F. Harris III, "Energy-efficient routing schemes for underwater acoustic networks," *IEEE J. Sel. Areas Commun.*, vol. 26, no. 9, pp. 1754–1766, Dec. 2008.

[4] D. Pompili, T. Melodia, and I. F. Akyildiz, "Routing algorithms for delay-insensitive and delay-sensitive applications in underwater sensor networks," in *Proc. ACM Mobicom*, 2006.

[5] L. Badia, M. Levorato, and M. Zorzi, "Markov analysis of selective repeat type II hybrid ARQ using block codes," *IEEE Trans. Commun.*, vol. 56, no. 9, pp. 1434–1441, Sep. 2008.

[6] Q. Zhang and S. A. Kassam, "Finiste-state Markov model for Rayleigh fading channels," *IEEE J. Sel. Areas Commun.*, vol. 17, no. 5, pp. 867–880, May 1999.

[7] M. Rossi, L. Badia, and M. Zorzi, "SR ARQ delay statistics on N-state Markov channels with non-instantaneous feedback," *IEEE Trans. Wireless Commun.*, vol. 5, no. 6, pp. 1526–1536, Jun. 2006.

[8] Q. Zhang and S. A. Kassam, "Hybrid ARQ with selective combining for fading channel," *IEEE J. Sel. Areas Commun.*, vol. 17, no. 5, pp. 867–880, May 1999.

[9] J.-B. Seo, Y.-S. Choi, S.-Q. Lee, N.-H. Park, and H.-W. Lee, "Performance analysis of a Type-II Hybrid-ARQ in a TDMA system with correlated arrival over a non-stationary channel," in *Proc. ISWCS*, 2005, pp. 59–63.

[10] L. Badia, M. Levorato, and M. Zorzi, "A channel representation method for the study of hybrid retransmission-based error control," *IEEE Trans. Commun.*, vol. 57, no. 7, pp. 1959–1971, Jul. 2009.

[11] E. Soijanin, N. Varnica, and P. Whiting, "Punctured vs. rateless codes for hybrid ARQ," in *Proc. of IEEE ITW*, Punta del Este, Uruguay, 2006, pp. 155–159.

[12] K. McCoy, "JANUS: from primitive signal to orthodox networks," in *Proc. of IACM UAM*, Nafplion, Greece, Jun. 2009.

[13] "Teledyne benthos undersea systems," www.benthos.com.

[14] K. V. Mackenzie, "Nine-term equation for the sound speed in the oceans," *J. Acoust. Soc. Am.*, vol. 70, no. 3, pp. 807–812, 1981.

[15] B. Tomasi *et al.*, "Experimental study of the space-time properties of acoustic channels for underwater communications," in *Proc. IEEE/OES Oceans*, Sydney, Australia, May 2010.

[16] S. Lin and D. J. Costello, *Error Control Coding: Fundamentals and Applications*. Prentice-Hall, 1983.

[17] R. Liu, P. Spasojevic, and E. Soljanin, "Reliable channel regions for good binary codes transmitted over parallel channels," *IEEE Trans. Inf. Theory*, vol. 52, no. 4, pp. 1405–1424, Apr. 2006.