Beam Selection Based on Sequential Competition

Mostafa Khalili Marandi, Wolfgang Rave and Gerhard Fettweis
{mostafa.khalili_marandi, wolfgang.rave, gerhard.fettweis}@tu-dresden.de

***Abstract*—We present a novel M-ary sequential test for beam selection when knowledge about the SNR operating point is *not* available. The proposed sequential test adaptively changes the test length (the number of observations) according to the SNR operating point to achieve the desired performance. Moreover, to achieve the same performance in terms of captured signal power, the *sequential competition test* requires on average less observations (particularly in the lower SNR regime) in comparison to a perfectly tuned *fixed length* test assuming genie knowledge.**

***Index Terms*—Beam Selection, Sequential Test, Generalized Likelihood Ratio Test**

I. Introduction

Directional data transmission in practical multiple antenna systems can be simply achieved via electronically controllable beamforming networks like Butler Matrix [1]. Such a network provides a set of orthogonal beams (denoted as 'codebook') steered into different mainlobe directions.

A crucial question at the receiver side during beam training showed schematically in Fig 1, is how to efficiently select the best beam(s) capturing the highest amount of receive power from a given codebook.

State of the art beam-alignment techniques e.g. exhaustive or tree search with hierarchical codebooks, use a *fixed* length training sequence for evaluation of the best beam i.e. beam selection [2][3][4][5]. A *fixed* length test can only be optimally designed for one particular Signal to Noise Ratio (SNR). Below this point the decisions become unreliable and above it we take too many observations. However, in most practical applications, the knowledge about SNR is not or only roughly available and the problem needs to be solved over some non-negligible SNR range due to varying operation conditions. This fact motivated us to develop an adaptive *variable* length approach.

In this letter we present the underlying idea of beam selection based on sequential competition which was introduced in [6]. This algorithm learns on the fly the current statistics of the signals in terms of their amplitudes and noise variances based on the available observations from each beam[1]. In this way the decision on the best beam(s) can be made both as early as possible and adaptively according to the SNR operating point. This becomes particularly interesting when the number of beams from which we can choose, is large and/or the SNR value is low. This scenario is expected to occur frequently during initial acquisition or handover in communication systems using large antenna arrays.

[1]Note that this is a M-*ary* decision problem while M observation sequences are available and not just one observation sequence.

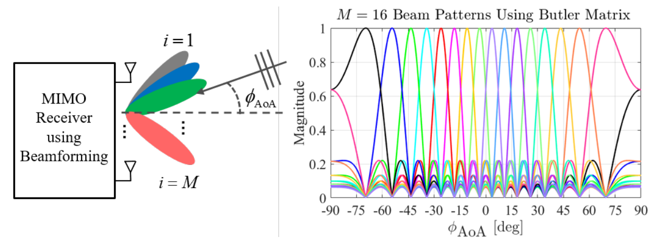

Fig. 1. Illustration of the beam selection problem for a single path channel: Which beam should be chosen to capture the highest amount of signal power?

II. Analog Beam Selection Problem

Consider the problem of finding the detector for selecting the best beam which captures the highest signal power, from a set of candidate beams each with different main lobe direction.

Although the channel model for mmWave systems contain features like multipath with multiple delays and angle of arrivals, we believe that the essential aspects of the detection problem under consideration that arises after correlation to a training sequence, are captured by a flat channel model. This leads to the following complex valued received samples which are observed separately[2] under each candidate beam as

$$r_m[n] = A_m s[n] + w_m[n] \ , \tag{1}$$

where $m \in \{1,\dots,M\}$ and $n \in \{1,\dots,N\}$ indicate the beam and sample indices. A_m is the combined effective channel and beamforming gain corresponding to beam m, that is treated as a deterministic unknown complex amplitude. $w_m[n]$ is complex zero mean white Gaussian noise (WGN) with unknown variance σ^2 under beam m. A pseudo-random sequence[3] with $s[n] \in \{\pm 1\}$, variance one and $P\{s[n] = +1\} = P\{s[n] = -1\} = 1/2$ is assumed for training so that $\mathrm{E}[s[n]s[n-k]] \simeq \delta[k]$ holds for its autocorrelation sequence.

Considering the vectors of correlated observations $\mathbf{y}_m = [s[1]r_m[1], \dots, s[N]r_m[N]]$, a fixed length test with length N decides for beam k over all candidate beams, based on minimum variance unbiased (MVU) signal power estimates, if

$$k = \underset{m\in\{1,\dots,M\}}{\operatorname{argmax}} \ \{|\mathbf{y}_m \mathbf{y}_m^{\mathrm{H}}|/N\} \ , \tag{2}$$

where $(.)^{\mathrm{H}}$ and $|.|$ denote conjugate transpose and absolute value operations, respectively.

[2]Depending on the number of RF chains with respect to number of antennas, the observations from different beams at each n can be made in parallel, serial or in a combined fashion.

[3]Assuming a set of sequences with good auto- and cross-correlation properties so that different users can be assigned to different sequences allows a generalization to multipath and multiuser scenarios.

Defining the maximum gain as $|A_{\max}| = \max\{|A_1|, \ldots, |A_M|\}$, the above detector results in the following normalized average loss of the signal magnitude

$$\bar{l} = 1 - \frac{1}{|A_{\max}|} \sum_{m=1}^{M} P\{k = m|\mathbf{y}_1, \ldots, \mathbf{y}_M\} |A_m| \ , \quad (3)$$

where $P\{k = m|\mathbf{y}_1, \ldots, \mathbf{y}_M\}$ denotes the probability of selecting beam m after observing the sequences $\mathbf{y}_1, \ldots, \mathbf{y}_M$.

The problem that arises in the design of the training sequence (or the detector) is to choose a specific length N to achieve a certain performance in terms of $\bar{l}$ in a range of scenarios that occur randomly in applications where exact knowledge about $\{|A_1|, \ldots, |A_M|\}$ with respect to σ^2 is not or only roughly available. Naively fixing the test length to some value N based on a certain assumed operating point can result in a strongly variable performance in practical scenarios. Additionally, if N is conservatively set to a high value based on the worst still acceptable operating point, a lot of time spent for detection of the best beam will be wasted, if the channel quality is actually better than expected. This is particularly important when the channel coherence time is limited and waisting time for training results in loss of throughput. This gives us the motivation to look for an adaptive variable length test that can achieve the same desirable performance over a range of operating points while keeping the test length (i.e. training time) as small as possible.

III. Generalized Likelihood Ratio Test: Detection of Signal Presence

The building block of our sequential test presented in the next section is the fixed length Generalized Likelihood Ratio Test (GLRT) to detect the presence or absence of the signal with unknown amplitude.

Consider the following composite detection problem

$$\begin{array}{ll} \mathcal{H}_0 & : y[n] = w[n] \\ \mathcal{H}_1 & : y[n] = A + w[n] \end{array} , \quad (4)$$

where $n = 1, \ldots, N$, A is a nonzero complex deterministic unknown amplitude and $w[n]$ is zero mean complex WGN with unknown variance σ^2.

Since the variance of the PDF under $\mathcal{H}_0$ is not known, a proper threshold to bound the probability of deciding $\mathcal{H}_1$ when $\mathcal{H}_0$ is true, typically denoted as the probability of false alarm P_{FA}, can not be found using a simple Neyman-Pearson (NP) approach [7]. Instead one can use the Maximum Likelihood (ML) estimates of the unknown parameters derived from the available observations and insert them into the likelihood functions under each hypothesis. The ML estimates of A and σ^2 under $\mathcal{H}_1$ are $\hat{A} = \mathbf{y}\mathbf{1}^T/N = \bar{\mathbf{y}}$ (the sample mean) and $\hat{\sigma}^2_{\mathcal{H}_1} = |(\mathbf{y} - \bar{\mathbf{y}})(\mathbf{y} - \bar{\mathbf{y}})^{\text{H}}|/N$, while under $\mathcal{H}_0$ the ML estimate of σ^2 is just $\hat{\sigma}^2_{\mathcal{H}_0} = |\mathbf{y}\mathbf{y}^{\text{H}}|/N$.

By replacing the unknown parameters with their ML estimates in the Gaussian PDFs, the GLR can be calculated as

$$L_{\text{G}}(\mathbf{y}) = \frac{p(\mathbf{y}; \hat{A}, \hat{\sigma}^2_{\mathcal{H}_1}, \mathcal{H}_1)}{p(\mathbf{y}; \hat{\sigma}^2_{\mathcal{H}_0}, \mathcal{H}_0)} = \left(\frac{\hat{\sigma}^2_{\mathcal{H}_0}}{\hat{\sigma}^2_{\mathcal{H}_1}}\right)^{\frac{N}{2}} . \quad (5)$$

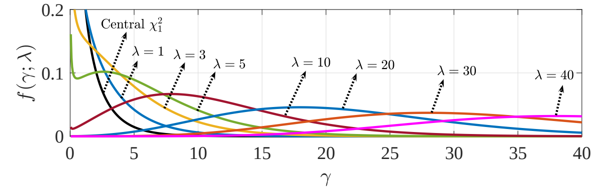

Fig. 2. Illustration of χ_1^2 distributions with different λ values.

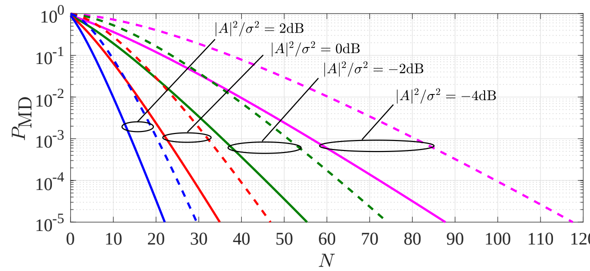

Fig. 3. Decay of P_{MD} with increasing test length N for different SNR values where $P_{\text{FA}} = 0.1$ for solid lines and $P_{\text{FA}} = 0.01$ for dashed lines.

Eq. (5) indicates that deciding for $\mathcal{H}_1$ makes sense when the fit of the data to the signal amplitude estimate $\hat{A} = \bar{\mathbf{y}}$ produces a smaller error, as measured by $\hat{\sigma}^2_{\mathcal{H}_1}$, compared to the fit to the no signal hypothesis reflected by the estimate $\hat{\sigma}^2_{\mathcal{H}_0}$.

The remaining task is to find a proper decision threshold in order to bound P_{FA}. Let us define the modified GLR statistic $\gamma \equiv 2 \ln L_{\text{G}}(\mathbf{y})$. A non-trivial result that can be found in [8] states that for large N the random variable γ follows either a central or a non-central χ^2-distribution with one degree of freedom so that

$$\gamma = N \ln \left(\frac{\hat{\sigma}^2_{\mathcal{H}_0}}{\hat{\sigma}^2_{\mathcal{H}_1}}\right) \sim \begin{cases} \chi_1^2, & \text{under } \mathcal{H}_0 \\ \chi_1^2(\lambda), & \text{under } \mathcal{H}_1 \end{cases} . \quad (6)$$

The ratio $\lambda = N(|A|^2/\sigma^2)$ that plays the role of the deflection coefficient is denoted in statistics as the non-centrality parameter of the $\chi_1^2(\lambda)$-PDF (see Fig. 2). Since now the PDF of γ under $\mathcal{H}_0$ is completely known, using the NP design criterion, we can ensure that P_{FA} will not surpass a predefined value by finding a proper threshold γ_{th}. Noting that a χ_1^2 r.v. γ is related to a standard normal r.v. $x \sim \mathcal{N}(0, 1)$ as $\gamma = x^2$, it follows that $P_{\text{FA}} = \Pr\{\gamma > \gamma_{\text{th}}; \mathcal{H}_0\}$ can be expressed as a sum of Q-functions so that $P_{\text{FA}} = \Pr\{x > \sqrt{\gamma_{\text{th}}}\} + \Pr\{x < -\sqrt{\gamma_{\text{th}}}\} = 2Q(\sqrt{\gamma_{\text{th}}})$. This specifies γ_{th} in terms of P_{FA} as

$$\gamma_{\text{th}} = \left[Q^{-1}\left(\frac{P_{\text{FA}}}{2}\right)\right]^2 . \quad (7)$$

Similarly, the second type of error, the probability of misdetection $P_{\text{MD}} = 1 - \Pr\{\gamma > \gamma_{\text{th}}; \mathcal{H}_1\}$ can be stated in closed form as

$$P_{\text{MD}} = Q(\sqrt{\lambda} - \sqrt{\gamma_{\text{th}}}) - Q(\sqrt{\lambda} + \sqrt{\gamma_{\text{th}}}). \quad (8)$$

It is worth noting that, for large enough N, the test

$$\gamma \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \gamma_{\text{th}} \quad (9)$$

is a Uniformly Most Powerful (UMP) test. This means the performance of this detector in terms of P_{MD} achieves the bound given by the clairvoyant NP detector that assumes perfect knowledge of the non-centrality parameter λ under $\mathcal{H}_1$.

As shown in Fig. 3, the probability of misdetection P_{MD} decreases exponentially as N increases due to the Q-function. In addition, the rate of decrease depends strongly on SNR per observation $|A|^2/\sigma^2$. The sequential competition test that we introduce next exploits this strong dependence of P_{MD} on $|A|^2/\sigma^2$ and N (while P_{FA} is fixed), and turns it into an advantage.

IV. Sequential Competition Test

Consider again the initial M-*ary* decision problem stated in Section II where separate observation sequences corresponding to different beams are available and the aim is to detect the beam that captures the highest signal magnitude. This time, instead of comparing the estimates of the values $\{A_1, \ldots, A_M\}$ directly against each other as in Eq. (2), let us rather compare them separately to the absence of a signal. This means that under each beam $m \in \{1, \ldots, M\}$ we formulate the same *binary* hypothesis test stated in the previous section as

$$\begin{array}{ll} \mathcal{H}_{m,0}: & y_m[n] = w_m[n] \\ \mathcal{H}_{m,1}: & y_m[n] = A_m + w_m[n] \end{array}, \tag{10}$$

however, this time n can grow until a decision criterion is fulfilled. Obviously, $\mathcal{H}_{m,0}$ is the wrong hypothesis under each beam, assuming that some signal is observable but with different strength. On the other hand, $\mathcal{H}_{m,0}$ acts as a virtual common reference in the set of M parallel binary tests, into which the M-*ary* test is decomposed.

Let us denote the probability of selecting the presence of the signal in the binary test of beam m after n observations as $P_{\mathcal{H}_{m,1}}(n)$. Comparing the binary tests of beam m and beam m', it follows from Eq. (8) that for $|A_m| > |A_{m'}|$ and a common decision threshold γ_{th} that $P_{\mathcal{H}_{m,1}}(n) > P_{\mathcal{H}_{m',1}}(n)$. This is simply a consequence of the fact that the accumulated deflection coefficient $\lambda_m = n|A_m|^2/\sigma^2$ will grow more quickly than $\lambda_{m'} = n|A_{m'}|^2/\sigma^2$. Therefore, if we use the modified GLR statistic introduced in the last section as a decision metric, the beam that observes the stronger signal will on average cross the threshold earlier. This fact leads to the following sequential competition test applied to stochastic paths $\gamma_m[n]$ for $m = 1, \ldots, M$ stated with pseudo code in Algorithm 1.

At each step n all stochastic paths $\gamma_m(n)$ corresponding to beams $m = 1, \ldots, M$ are compared to the fixed common threshold γ_{th}. The test terminates as soon as one of the paths surpasses the threshold while the index of this path indicates the selected beam. Otherwise, we continue by taking the next observation into account. The interpretation is that we let the beams compete to distinguish themselves from pure zero mean WGN with the same unknown variance, and the one which does it faster is the winning beam in the competition. The test length n is now a random variable with average $\bar{n}$.

Algorithm 1 Sequential Competition Test

input: $M, P_{\mathrm{FA}}, N_{\max}, n = 1, \gamma_m[1] = 0$
$\gamma_{\mathrm{th}} = [Q^{-1}(P_{\mathrm{FA}}/2)]^2$
while $\max_m(\gamma_m[n]) < \gamma_{\mathrm{th}} \wedge n \leqslant N_{\max}$ **do**
  $n = n + 1$
  for $m = 1, \ldots, M$ **do**
    $\mathbf{y}_m = [y_m[1], \ldots, y_m[n]]$
    $\bar{\mathbf{y}}_m = \sum_{i=0}^{n} y_m[i]/n$
    $\hat{\sigma}^2_{\mathcal{H}_{m,0}} = |\mathbf{y}_m \mathbf{y}_m^{\mathrm{H}}|/n$
    $\hat{\sigma}^2_{\mathcal{H}_{m,1}} = |(\mathbf{y}_m - \bar{\mathbf{y}}_m)(\mathbf{y}_m - \bar{\mathbf{y}}_m)^{\mathrm{H}}|/n$
    $\gamma_m[n] = n \ln(\hat{\sigma}^2_{\mathcal{H}_{m,0}}/\hat{\sigma}^2_{\mathcal{H}_{m,1}})$
  end for
end while
output: $\operatorname{argmax}_m(\gamma_m[n])$

The common threshold γ_{th} for all virtual binary tests is selected using Eq. (7) without knowing the values $\{|A_1|, \ldots, |A_M|\}$ and σ^2. The higher we put the threshold (i.e. choosing a smaller allowed P_{FA} under each virtual binary test) the longer it takes on average for the sequential competition test to terminate and make a decision. This however, leads to a more accurate discrimination of the strongest beam among all competitors and therefore better performance in terms of $\bar{l}$.

In the following toy example depicted in Fig. 4 with $M = 2$ and ratio $r = |A_1|/|A_2| = 0.5$, we observe that the sequential competition test shows an essentially invariant adaptive performance in terms of $\bar{l}$ for varying value of the differential SNR between the candidate amplitudes. This is in strong contrast to the fixed length test that uses length of $N_{\mathrm{fix}} = 25$. The reason is that the average test length $\bar{n}$ (over multiple realizations) of the sequential competition test adapts itself to $(|A_1| - |A_2|)^2/\sigma^2$ as shown in the bottom plot of Fig 4.

In addition, as depicted in Fig. 5, the sequential test reacts adaptively as well to the varying ratio $r = |A_1|/|A_2|$ by changing the average number of observations. For r close to one, where the possible average loss $\bar{l}$ is negligible but the detection of the stronger beam is more difficult, the performance of the sequential test in terms of $\bar{l}$ gets closer to the upper-bound $(1 - r)/2$ given by the random selection i.e. coin tossing (top plot), while requiring only few observations on average (bottom plot), which makes sense.

Observe further that, the sequential competition test requires on average even less observations $\bar{n}$ in the lower SNR regime to achieve a certain value of $\bar{l}$ compared to the required number $N_{\mathrm{fix}}^{\mathrm{genie}}$ of a fixed length test tuned with genie knowledge that would achieve the same $\bar{l}$ (see Fig. 4, bottom). This can be understood intuitively, because according to the fluctuations of the competing stochastic processes around their typical behavior, the test exactly terminates when a reliable discrimination becomes possible, so that the test terminates earlier *on average* (which was the original motivation of *Wald* to develop his test [9]). This property is of particular interest, because the sequential test reduces training time at exactly those smaller

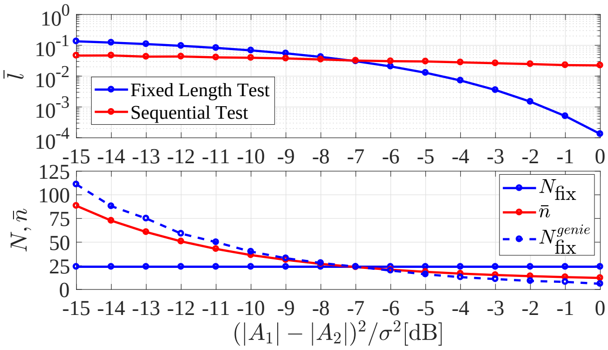

Fig. 4. Comparing the achieved $\bar{l}$ between a fixed length test with length N_{fix} and the sequential test with equal γ_{th} based on $P_{\text{FA}} = 10^{-3}$.

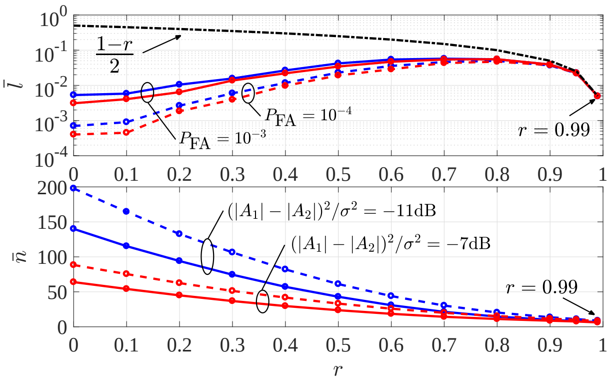

Fig. 5. Performance of the sequential test over ratio r. The upper-bound (beam selection by coin tossing) on $\bar{l}$ is depicted as black dot dashed line.

values of $(|A_1| - |A_2|)^2/\sigma^2$ where many observations are needed. On the other hand, when $(|A_1| - |A_2|)^2/\sigma^2$ becomes large the detection problem becomes easy and we do not need many observations in the first place.

V. Numerical Evaluation

We numerically studied the performance of the *sequential competition test* in the reference channel model described in Eq. (1) with a uniform linear array with 16 antenna elements using the codebook of a Butler matrix with 16 beams. The AoA was distributed uniformly in $[-90^\circ, 90^\circ]$ over the simulation runs while SNR was defined as $|A_{\text{max}}|^2/\sigma^2$ indicating the maximum available SNR of the best beam. The quantities $\bar{l}$ and $\bar{n}$ were estimated at each SNR point based on 10^4 simulation runs for SNR values in the interval $[-8, 8]$ dB. For comparison we consider the fixed length detector based on MVU power estimates stated in Eq. 2, with $N_{\text{fix}} \in [10, 25, 50, 75]$. As shown in Fig. 6, the sequential competition test with the same γ_{th} for all beams based on $P_{\text{FA}} = 10^{-3}$ keeps $\bar{l}$ in an interval of $[0.07, 0.14]$ considered to be sufficient in practice, while adaptively decreasing the average test length $\bar{n}$ as $|A_{\text{max}}|^2/\sigma^2$ grows larger. In contrast, fixed length tests show strongly variable performance. For instance, the fixed length test with $N_{\text{fix}} = 25$ results in $\bar{l}$ in an interval of $[0.005, 0.55]$.

In case of finite transmission time (e.g. limited channel coherence time), the shorter we spend time for training, the more time is left for data transmission. On the other hand, the more time we spend for training, the loss of SNR due to inaccurate training will reduce. This means for $N_{\max}$ possible channel uses within finite transmission time, we can evaluate the ratio between effective data rates by sequential and fixed length tests via

$$\frac{R_{\text{eff,seq}}}{R_{\text{eff,fix}}} = \frac{\mathbb{E}\left[(1 - \frac{n}{N_{\max}}) \log(1 + (1 - l_{\text{seq}})^2 \frac{|A_{\max}|^2}{\sigma^2})\right]}{\mathbb{E}\left[(1 - \frac{N_{\text{fix}}}{N_{\max}}) \log(1 + (1 - l_{\text{fix}})^2 \frac{|A_{\max}|^2}{\sigma^2})\right]}, \quad (11)$$

where l_{seq} and l_{fix} are normalized signal magnitude loss under each individual simulation run and $\mathbb{E}$ denotes the expectation over all simulation runs. As depicted in Fig. 6 bottom, the sequential competition test fulfills the training-transmission trade-off better compared to fixed length tests by providing higher effective data rate at almost all SNR operating points.

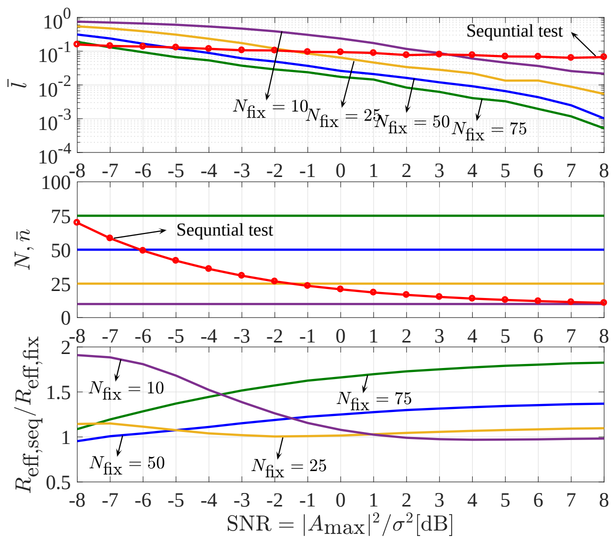

Fig. 6. Achieved $\bar{l}$, $\bar{n}$ and ratio between effective rates using sequential competition test and fixed length tests with $N_{\max} = 150$.

VI. Concluding Remarks

We proposed a novel M-*ary* sequential hypothesis test based on GLR statistics to solve the composite beam selection problem. The proposed *sequential competition test* shows adaptivity w.r.t. the SNR operating point. Furthermore, it speeds up beam selection on average even when compared to an optimally (i.e. having genie knowledge) tuned fixed length test at lower SNR where it matters the most or most time will be spent for training in a possible application. Almost the same quantitative behaviors can be observed in a more complicated multipath mmWave channel which requires additional correlation steps. Therefore, these benefits can be of interest in systems with large number of candidate beams (e.g. mmWave Massive MIMO systems) as well as under conditions where the training time is limited due to small channel coherence time. Our future work will focus on multi-user beam selection using sequential competition test.

REFERENCES

[1] J. Butler, R. Lowe *"Beam-forming Matrix Simplifies Design of Electronically Scanned Antennas,"* Electronic Design, vol. 9, pp. 170-173, April 12, 1961

[2] C.N. Barati, S. Amir Hosseini, S. Rangan, P. Liu, T. Korakis, S. S Panwar and T. S. Rappaport *"Directional Cell Discovery in millimeter Wave Cellular Networks,"* IEEE Trans. on Wireless Commun., vol. 14, no. 12, pp. 6664-6678, Dec. 2015

[3] M. Hussain and N. Michelusi *"Energy-Efficient Interactive Beam-Alignment for Millimeter-Wave Networks,"* IEEE Trans. on Wireless Commun., 2018

[4] J. Zhang, Y. Huang, Q. Shi, J. Wang, L. Yang *"Codebook Design for Beam Alignment in Millimeter Wave Communication Systems,"* IEEE Trans. on Commun., vol. 65, no. 11, pp. 4980-4995, 2017

[5] X. Song, S. Haghighatshoar and G. Caire *"A scalable and statistically robust beam alignment technique for mm-wave systems,"* IEEE Trans. on Wireless Communications, vol. 17, no. 7, pp. 4792-4805, May 2018

[6] M. Khalili Marandi, W. Rave, G. Fettweis *"An Adaptive Sequential Competition Test for Beam Selection in Massive MIMO Systems,"* 10th IEEE Sensor Array and Multichannel Signal Processing Workshop (SAM), pp. 104-108, Jul. 2018

[7] J. Neyman and E. S. Pearso *"On the Problem of the Most Efficient Tests of Statistical Hypotheses,"* Phil. Trans. of the Royal Society A, vol. 231, pp. 289-337, 1933

[8] S. M. Kay *"Fundamentals of statistical signal processing, Vol. II Detection Theory,"* Prentice Hall 1998

[9] A. Wald *"Sequential tests of statistical hypotheses,"* Ann. Math. Statistics, vol. 16, pp. 117-186, 1945