Searching d-defective sets with queries of size k

D. K. Vu*

Consider a set X of n elements. We wish to identify a particular subset Y containing at most d unknown elements. To this end, we perform a series of experiments with the following property: when testing a subset $A \subseteq X$, we receive a *positive* result if and only if A contains at least one of these d unknown elements. In practice, we often have the additional constraint that $|A| \leq k$, and we desire to minimize the total number of queries while yet determining Y exactly. This can be done adaptively, meaning that the answer of to a query influences which queries are made in the course of the search, or non-adaptively, where all questions are determined in advance. In the non-adaptive case, a successful family of such queries is often referred to as a *(d-)separating family.*

This question was first posed by A. Rényi in 1961. For the case of $d = 1$ G. O. H. Katona solved the adaptive case and provided upper and lower estimates for the non-adaptive case in 1966. In 2013, É. Hosszu, J. Tapolcai and G. Wiener simplified the proof remarkably. Using some of their ideas, we obtain similar results for general d. While the adaptive case is very similar, we also provide new (and to our knowledge the first non-trivial) upper and lower bounds in the non-adaptive case. We do so by examining the relationship between the girth of hypergraphs and separability.

In this talk the focus will be on the cases of $d = 2, 3$ for illustrative purposes.

This is joint work with F. S. Benevides, M. Delcourt, D. Gerbner, C. Palmer and B. Sinaimeri.

*Department of Mathematical Sciences, University of Memphis, Memphis, TN-38152, USA, dominik.vu@memphis.edu