The Baker Prize committee of the Information Theory Society, composed of Max Costa, Andrea Goldsmith, Gerhard Kramer, Muriel Medard (chair) and Gregory Wornell, has considered several papers for possible Baker Award nomination from the Information Theory Society. After considering many influential papers in the appropriate time range for the prize, the committee decided to consider the papers that were premiated. These papers were considered in terms of innovation, technical quality and impact. The committee had three conference calls, as well as extensive online discussions.

While the committee had varied views, the final recommendation, particularly for its impact, is the paper

Viveck R. Cadambe, Syed A. Jafar, “Interference Alignment and the Degrees of Freedom of the K user Interference Channel”, IEEE Transactions on Information Theory, Aug. 2008, Vol. 54, No. 8,pages 3425 - 3441.

The basic premise of the ground-breaking result reported in this paper is that, in a network of K >= 2 distributed wireless nodes, each node can transmit information such that it has interference-free communication half of the time, with the other half of the time dominated by interference from other nodes. Thus, no matter how many nodes are in the wireless network, a given user gets \half the cake," achieving data rates up to half what would be possible with no interfering nodes in the shared spectrum. This surprising result is in stark contrast to the conventional wisdom on which many current wireless network designs are based, where spectrum is chopped up and allocated to each node in an orthogonal manner so that for K nodes, each node can only communicate using a fraction 1/K of the total bandwidth, thereby getting only 1/K of the cake. Interference alignment means that communication in wireless networks are not limited by the number of nodes sharing the spectrum, a truly revolutionary notion.

This paper builds on the prior work by Maddah-Ali, Motahari and Khandani (ISIT 2006) and by the same authors on a specialized communication scenario known as the X channel, where the first explicit interference alignment schemes were presented. This paper is the first to introduce the full-edged idea of interference alignment in a general context: interference networks with an arbitrary number of users, with interference alignment presented as a broad principle. The alignment of vector subspaces needed for everyone to get "half the cake" is at first sight an over-determined problem that affords no solution. It is intimately connected to the existence of common invariant subspaces of an increasing number of linear transformations: a problem that is well known in linear algebra to have no solution. Cadambe and Jafar successfully circumvent this hurdle by an ingenious variation: instead of asking for exact common invariant subspaces, the paper looks for "almost invariant" subspaces, i.e., vector subspaces that are mapped almost entirely to themselves by all channel matrices. This leads to the alignment scheme in this paper over multiple channel symbols where the almost common invariant subspace formulation becomes increasingly a completely invariant subspace solution. Despite the technical sophistication underlying this derivation, the results of the paper are nicely summarized by well-chosen illustrative examples, making the paper broadly accessible.

The paper has ignited the imagination of the information theory community, leading to many follow-up works with extensions such as ergodic, distributed, and opportunistic interference alignment as well as application of the interference alignment idea to a wide range of interference networks. Given the many open information theory problems related to channels with interference, this result is one of the most significant in wireless information theory over the last decade. The paper has had deep and immediate impact in both theory and practice. In fact, few other results in information theory over the last decade have moved so quickly from theory to implementation, MIMO being a possible other example of such success. There was some discussion about the novelty of the approach, since the basic technique had already been developed and analyzed for smaller dimensions.

Another strong contender, which the committee strongly encourages future committees to consider, is the paper

E. Arikan , "Channel Polarization: A Method for Constructing Capacity-Achieving Codes for Symmetric Binary-Input Memoryless Channels", IEEE Transactions on Information Theory, Jul. 2009

Arikan's paper is linked directly to Shannon's work: it is based on the most elemental ideas of information theory (for instance, it has almost no references to prior work, except the author's earlier research in this direction). It is a direct answer to Shannon's challenge of constructing codes that achieve channel capacity. It is very carefully written and takes time to develop the ideas. Reading this paper is a thrilling experience because it takes many unexpected turns, and new substantial ideas appear not once (as in Guruswami/Rudra's work), not twice, but many times. Its culmination, a martingale proof of the capacity achieving property in Theorem 1, is an altogether stunning accomplishment. However, the paper does not stop there: it goes on to analyze the complexity of code construction and decoding, to quantify the rate of polarization, and to discuss a link to, and difference with, another recursive code construction, the Reed-Muller codes

This is a paper that at the same time closes a long- standing problem and opens a new line of research: it has closed Shannon’s problem of finding explicit capacity-achieving codes, and opened avenues for the improvement and applications of the polar codes. A large number of papers have already been written in this area: it has been shown that the recursive construction procedure can have many forms, a precise analysis has been performed for the error probability of decoding, intricate links to Reed-Muller codes and their decoding algorithms have been uncovered. Polar codes were used in a construction that saturates the rate region of some multiple-access channels, to achieve capacity of secrecy systems, to achieve the

theoretical compression rate of lossy source coding. However, it goes in a similar direction as concatenated codes: in both cases, one can provably achieve capacity, the codes are non-uniform and hence non-explicit in the sense that for every length a search needs to be performed to find the right "parameters" (the puncturing pattern for polar codes, the inner code for concatenated codes), and the proximity to channel capacity as a function of the length is not that great (better for concatenated codes and their new graph based versions of Barg and Zemor than for polar codes).

Respectfully submitted,

Muriel.