NSF Highlights

Fundamental Limit of Error-Correction Achieved

Highlight ID: 16270

Error-correcting codes are all around us

Permission Not Granted
Credit: Hard Drive: Michael Connors, MorgueFile. Other Images: Unknown.

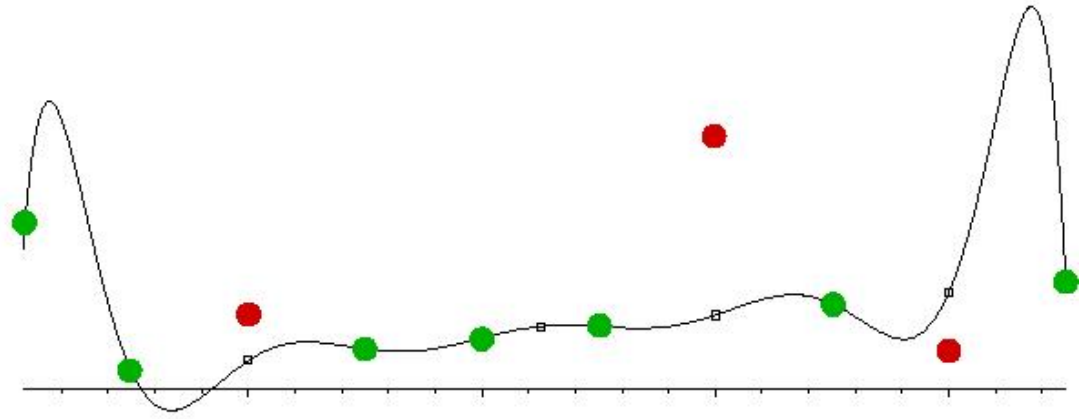

Univariate interpolation and Reed-Solomon decoding

Permission Granted
Credit: Venkatesan Guruswami, University of Washington

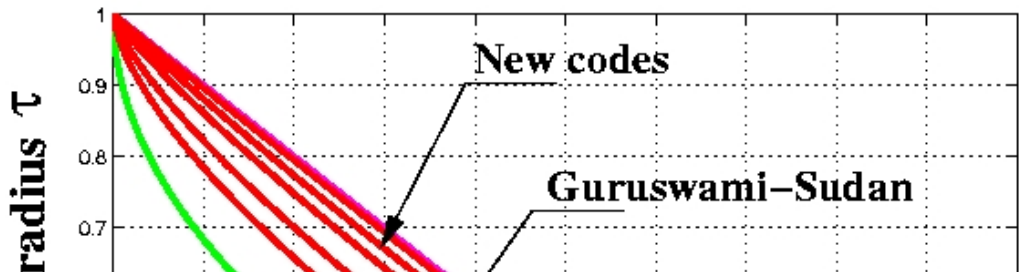

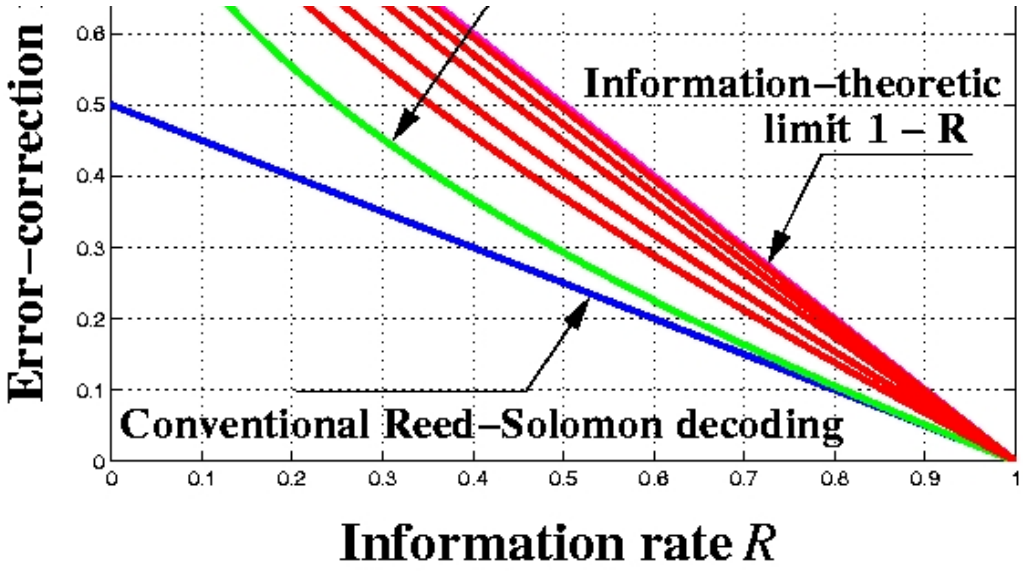

Trade-off between information rate and error-correction radius

Permission Granted
Credit: Venkatesan Guruswami, University of Washington

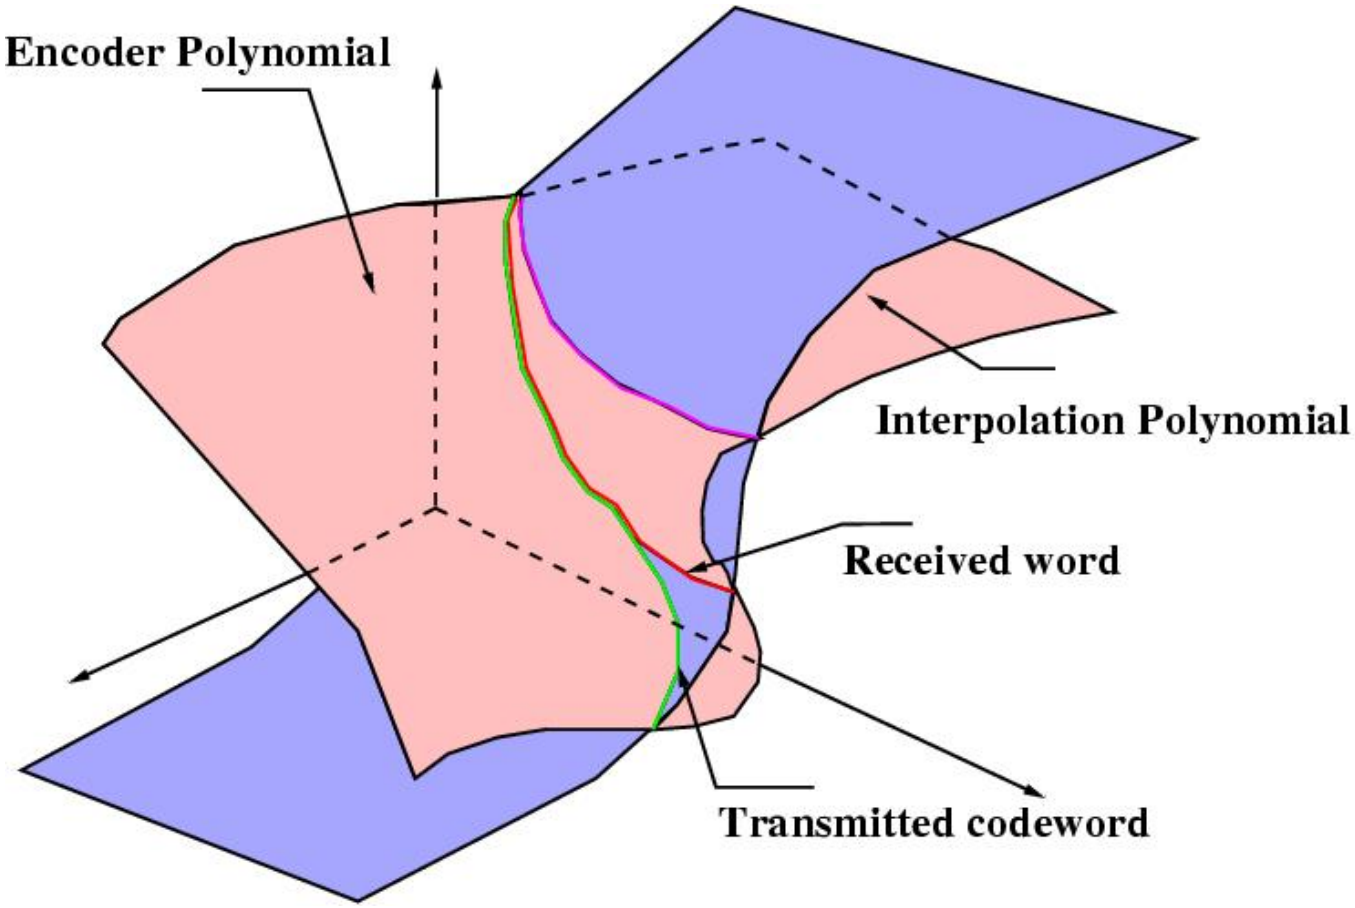

Conceptual sketch of the Parvaresh-Vardy coding scheme

Permission Granted
Credit: Venkatesan Guruswami, University of Washington

In the information age in which we now live, reliable transmission and storage of digital information is of paramount importance. What makes such reliable transmission and storage possible, despite the errors inherent to communication channels and storage media, are error-correcting codes, first conceived by Claude Shannon over 60 years ago. Recent advances made by two NSF-supported scientists --- Venkatesan Guruswami at the University of Washington and Alexander Vardy at the University of California San Diego, along with their graduate students --- have led to the discovery of error-correcting codes with the best possible trade-off between error-correction capability and redundancy. The newly discovered codes yield an improvement by a factor of two over conventional error-correction algorithms that are currently used in every CD player and every desktop PC, as well as a myriad other devices that directly impact our daily lives (Figure 1).

The basic idea of error-correction coding is relatively simple. Suppose that a sender (Alice) wishes to communicate to a receiver (Bob) a message consisting of k

symbols in such a way that Bob can always recover the message perfectly, even if e symbol errors occur during its transmission. To do so, Alice first encodes the message in a string of n > k symbols, using an error-correcting code that is capable of correcting any e transmission errors. The ratio R = k/n is called the "information rate" of the code, while the ratio t = e/n is known as its "error-correction radius." Obviously, we would like both R and t to be as large as possible, transmitting information at a high rate while, at the same time, correcting many errors. However, these are conflicting goals: to correct more errors, one has to add more redundant symbols to the message. Ever since the dawn of coding theory in the 1940s, researchers have wondered what is the best possible trade-off between these parameters R and t. Even more importantly, could such trade-off be achieved efficiently --- that is, with polynomial-time encoding and decoding algorithms? The recent results of Vardy (with his student Farzad Parvaresh) and Guruswami (with his student Atri Rudra) provide a fascinating answer to both questions. They show that it is possible to achieve, with constructive codes and polynomial-time decoding algorithms, the ultimate information-theoretic limit t = 1-R. This means that the redundancy in the encoding can be as close as one desires to the proportion of errors we wish to correct, which is the best one can hope for.

To describe the remarkable journey that culminated in these results, we need to go back to the work of Irving Reed and Gustave Solomon in 1960. Reed and Solomon advised Alice to think of the k symbols she wishes to convey as the coefficients of a polynomial, and send to Bob the values of this polynomial at some n different points. They showed that Bob can always recover Alice's polynomial perfectly by interpolating through the n values he receives, even if (n-k)/2 of them are in error (Figure 2). The Reed-Solomon codes thus establish the following trade-off between information rate and error-correction radius: t = (1-R)/2. Almost half a century after their invention, Reed-Solomon codes are still ubiquitous today in applications ranging from magnetic recording to satellite communications and fiber-optic networks.

For several decades, this classical trade-off t = (1-R)/2 was considered the best one could hope for, until Madhu Sudan from MIT surprised the scientific community in 1996 by showing that Reed-Solomon codes can correct more errors than previously thought possible. Sudan's algorithm does not always enable Bob to deduce Alice's message uniquely, but rather produces a small list of possible messages one of which is guaranteed to be correct. In practice, it turns out that error patterns that may cause such list-ambiguity are extremely rare, and list-decoding works just as well as conventional unique decoding. Somewhat counter-intuitively, Sudan's algorithm used interpolation in the domain of bivariate polynomials to recover the univariate polynomial that encodes Alice's message. Later, Guruswami and Sudan extended this idea by allowing multiplicities at the interpolation points, an elegant and powerful feature that has been vital to subsequent developments. This enabled them to establish the trade-off t = 1- sqrt{R} between rate and decoding radius (see Figure 3). Soon afterwards, Ralf Koetter from the University of Illinois and Vardy discovered a way to assign multiplicities in the Guruswami-Sudan algorithm so as to take into account the probabilistic measurements provided by a communication channel, thereby making this algorithm suitable for use in a wide variety of applications. These results have won broad acclaim, including the Nevanlinna Prize, the ACM Doctoral Dissertation Award, and two IEEE Information Theory Society Best Paper Awards. Further details on this NSF-supported work can be found at http://nsf.gov/discoveries/disc_summ.jsp?cntn_id=100256.

As years since the 1999 publication of the Guruswami-Sudan paper went by, efforts to improve the algorithm further did not meet with much success. There loomed the possibility that correcting a larger number of errors would cause too much ambiguity, making it impossible to accomplish decoding with a small list. It became clear that decoding beyond the Guruswami-Sudan radius, if at all possible, would require radically new methods. Such a breakthrough method was discovered by Parvaresh and Vardy [1], who showed in 2005 that even more errors could be corrected. To achieve this result, they ventured beyond the bivariate interpolation methods of Sudan and used polynomials in M variables, where M > 2 is a design parameter. They furthermore devised a clever new variant of Reed-Solomon codes based on the idea of including more information in the encoding of every message. Specifically, to encode a message polynomial f(X), Parvaresh and Vardy evaluate (at some n points, as before) both f(X) and a carefully chosen related polynomial g(X); the ingenious way in which g(X) is picked based upon f(X) forms the crux of the coding scheme. The general concept is sketched-out in Figure 4 for the case of trivariate interpolation (M = 3). The resulting scheme does enable correcting more errors, but unfortunately it also doubles the redundancy of the code, thereby reducing its information rate. Thus, although Parvaresh and Vardy succeeded in beating the Guruswami-Sudan radius for low rates, it appeared unlikely that their approach could be extended to high rates, which is often the regime of interest in communication settings.

Yet, in a recent paper [2] published in January 2008, Guruswami and Rudra managed to do just that. By orchestrating an algebraic miracle of sorts, they ensured that the second polynomial g(X) in the Parvaresh-Vardy scheme becomes essentially identical to the first: when chosen according to the Guruswami-Rudra method, the encoding of g(X) is just a cyclic shift of the encoding of f(X). There is then no need to explicitly send the extra encoding at all, and hence there is no loss in rate. The upshot of all this is a truly remarkable result: Guruswami and Rudra showed that the ultimate error-correction radius t=1-R can be reached and, moreover, it can be reached constructively, with polynomial-time encoding and decoding. This achieves the information-theoretic limit on the best possible trade-off between rate and decoding radius, for all possible rates! This furthermore corrects twice as many errors as the conventional decoding algorithms for Reed-Solomon codes (see Figure 3).

Although very recent, the coding schemes invented by Parvaresh and Vardy [1] and Guruswami and Rudra [2] have already found striking applications in diverse areas, ranging from compressed sensing to random number generation. Their work was recognized by prestigious Best Paper Awards at the IEEE Symposium on Foundations of Computer Science (FOCS'05) and at the IEEE Conference on Computational Complexity (CCC'07). It was also featured in a recent perspective article in Science [3].

It remains to be seen whether, some years from now, we will all be using the new decoding algorithms whenever we play a CD or access a computer hard disk. Numerous challenges must be overcome in order to reduce to practice the theoretical promise of the results discussed in the foregoing paragraphs. What is already clear, however, is that Vardy, Guruswami, and their graduate students have achieved an elusive milestone that has been sought by researchers in coding theory ever since the birth of this field 60 years ago.

REFERENCES:

[1] F. Parvaresh and A. Vardy, Correcting errors beyond the Guruswami-Sudan radius in polynomial time, Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science (FOCS), pp. 285-294, Pittsburgh, PA., October 2005.

[2] V. Guruswami and A. Rudra, Explicit codes achieving list decoding capacity: Error-correction with optimal redundancy, IEEE Transactions on Information Theory, vol. 54, no.1, pp. 135-150, January 2008.

[3] B. Chazelle, Perspectives in computer science: Coding and computing join forces, Science Magazine, vol. 317, no.5845, pp. 1691-1692, September 2007.

Primary Strategic Outcome Goal:

- Disciplinary/Interdisciplinary Research (Anything not covered by one of the 12 categories below.)
- CAREER

Secondary Strategic Outcome Goals:

How does this highlight address the strategic outcome goal(s) as described in the NSF Strategic Plan 2006-2011?:
This is a breakthrough, as well as the near-final word, on the central problem in coding theory. It has profound theoretical significance and overwhelming practical importance.

Does this highlight represent transformative research? *If so, please explain why.*
The National Science Board has defined transformative research as "Research that has the capacity to revolutionize existing fields, create new subfields, cause paradigm shifts, support discovery, and lead to radically new technologies." National Science Board: Enhancing Support of Transformative Research at the National Science Foundation

No

Does this highlight represent Broadening Participation? *If so, please explain why.*
The concept of broadening participation includes: individuals from underrepresented groups, certain types of institutions of higher education, geographic areas (e.g. EPSCoR states), and organizations whose memberships are composed of institutions or individuals underrepresented in STEM or whose primary focus is on broadening participation in science and engineering. It is important to note that underrepresented groups vary within scientific fields.

No

Are there any existing or potential societal benefits, including benefits to the U.S. economy, of this research of which you are aware? *If so, please describe in*

the space below.
It is important for NSF to be able to provide examples of NSF-supported research that have or may have societal benefits.

Yes

This work has the potential to make all forms of electronic data storage smaller, faster, and more reliable. It also has the potential to make all forms of electronic communication faster and more reliable.

CSE/CCF 2008

Program Officer: Sirin Tekinay & Richard Beigel

NSF Award Numbers:

0343672
Award Title: CAREER: Error-Correcting Codes --- List Decoding and Related Algorithmic Challenges
PI Name: Venkatesan Guruswami

Institution Name: University of Washington

PE Code: 2860

0514890
Award Title: Collaborative Research: Next Generation Decoders for Reed-Solomon Codes
PI Name: Alexander Vardy

Institution Name: University of California-San Diego

PE Code: 7351

NSF Contract Numbers:

NSF Investments: None Applicable

Entered on 02/19/2008 by Richard Beigel