Coarsely Quantized Redundant Representations of Signals

Sinan Güntürk (Courant Institute of Mathematical Sciences),
Thao Nguyen (City College, CUNY),
Alexander M. Powell (Vanderbilt University),
Özgür Yılmaz (University of British Columbia)

March 11 – March 16, 2006

1 Overview and Scientific Objective of the Workshop

Digital data is the driving force behind much of our modern technology. The Internet and cellular phones are ubiquitous examples of the need to handle information accurately, efficiently, and robustly. The fact that digital signals and data sets can be processed with great precision and speed places high demands on providing accurate conversion between the analog and digital worlds. However, the technology used in the analog to digital (A/D) conversion by necessity involves analog devices which have physical limitations that, at first sight, conflict with these accuracy demands. For example, typical printers are only capable of applying a very limited set of ink tones (which may be as small as a single black tone) to paper for rendering intermediate grayscale levels. Similarly, in audio applications, the inaccuracy of analog circuits in working with binary expansions is a problem that routinely needs to be addressed.

To cope with these problems, engineers have empirically developed special signal processing techniques leading to alternative signal and number representations that are quite different from the standard decimal or binary representations. Typical techniques take advantage of the fact that, although most analog devices used in A/D conversion fail to provide high precision in the amplitude domain, many of them are capable of sampling very densely in the time or the space domain over which the signals are defined. In these techniques, judiciously chosen dense arrangements of a limited set of discrete amplitude values are used to approximate the underlying analog signals. In the early days of printing, a similar task used to be carried out by an experienced halftoning artist. In today's digital world, there are several automated schemes that are implemented efficiently in consumer products, not only for digital halftoning, but also for audio encoding and decoding.

Although some of the very basic examples of such schemes were invented more than 30 years ago and the engineering practice has matured extensively in the meantime [1], a general mathematical theory of coarsely quantized representations did not experience a parallel development. In fact, many of the fundamental questions remained open until recently. In the past few years, a new interdisciplinary research activity was initiated by the breakthrough work of Daubechies and DeVore [2], and was followed by burgeoning interest among other mathematicians and engineers. This new theory has led to fruitful cross-pollination with applications in several disciplines of pure and applied mathematics, such as harmonic analysis, combinatorics, number theory, dynamical systems, approximation theory and information theory.

The BIRS workshop *Coarsely Quantized Redundant Representations of Signals* brought together expert mathematicians and theoretical engineers working on problems related or relevant to signal representations in coarsely grained environments. The strong interdisciplinary nature of the subject meant that mathematicians and engineers who are otherwise less likely to interact due to the distances in their respective fields of expertise had a great opportunity to meet, present their work and start future collaborations in the stimulating atmosphere of BIRS.

2 Background, Recent Developments and Open Problems

Coarse quantization is emerging as a part of a larger technological trend towards robust and redundant systems which utilize low cost components where the fundamental trade-off is between redundancy and hardware precision. Daubechies and Calderbank introduced the paradigm of “democracy of bits” to discuss this in a general setting [3]. This trend is further illustrated by the recent parallel advances in the sensors community, where dense microsensor networks using large numbers of highly correlated, but inexpensive, sensors are currently a huge area of focus. It is no coincidence that much of today’s cutting edge technology, such as Super Audio CD (SACD) systems, contains coarse quantization as a key component.

From an application point of view, the main focus of the workshop has been sigma-delta modulation in analog-to-digital encoding and error diffusion in digital halftoning:

Sigma-delta modulation is an oversampled analog-to-digital (A/D) conversion method for audio signals. A typical approach in obtaining digital representations of audio signals is to take the binary expansion of signal samples taken at the critical Nyquist rate of 44.1 KHz. However, this method suffers from high accuracy demands on the analog circuitry that is to be employed. Sigma-delta modulation, on the other hand, allocates as few as 1 bit for the discretization of signal samples taken at rates as high as 100 times faster than the critical rate. At the heart of this method lies an algorithm that recursively rounds (quantizes) each sample value to one of the few quantization levels in a way that is suited to achieve small global reconstruction error rather than small individual sample error.

The analogous problem in two dimensions is the digital halftoning process employed in image printing. While today’s printers can achieve high resolution in space (such as 1200 dpi), most often, only single-tone dots are available to be printed on the corresponding medium. In this case, only 0’s and 1’s can be used to represent efficiently any shade of the gray-scale. Error diffusion [4] can be thought of the exact analog of sigma-delta modulation in two dimensional signals and achieves this conversion effectively.

Sigma-delta modulation has its roots in the 60’s when Inose and Yasuda developed the first unity bit encoding by negative feedback [5]. Similarly, the classical error diffusion algorithm of Floyd and Steinberg was invented in the 70’s [6]. In the late 80’s, interest in the information theory community was led by Gray who discovered some of the best known theoretical results, e.g., [7, 8]. It was, however, only in the late 90’s that an approximation theoretical framework was given to the problem by Daubechies and DeVore [2]. Since then, there has been a rapid development in the theoretical analysis of sigma-delta modulation with emerging connections to other mathematical fields [9, 10, 11].

At the core, the problem of signal quantization is linked with the underlying choice of signal representation. Overcomplete data expansions not only provide robustness with respect to noise and data loss, but most coarse quantization algorithms are, in fact, specifically built around exploiting redundancy. Frame theory [12] provides a mathematical framework for stably representing signals by overcomplete expansions. Work during the past two decades on sampling wavelet frames, Gabor frames, and Fourier frames has greatly advanced this field. More recently, finite frames have gained attention as a piece of this theory which is tailored for applications involving finite, but potentially high dimensional, data. For example, finite frame expansions have recently been proposed as joint source-channel codes for erasure channels [13], for multiple-antenna code design [14], and for modified quantum detection problems [15].

The problem of one-bit quantization is related to many other mathematical disciplines. Some of these would be regarded as classical now, but at the same time many connections are new or relatively unexplored.

- Combinatorics: combinatorial, geometric and linear discrepancy theory form the common grounds for understanding some of the fundamental problems in one-bit quantization, especially in understanding some of the universal lower bounds.

- Dynamical systems: most algorithms that yield efficient one-bit or low-bit representations are based on dynamical systems (sigma-delta modulation and error diffusion both fall into this category). Typically, underlying these dynamical systems are certain non-expanding piecewise affine maps on the Euclidean space. There is a strong connection between understanding the stability properties of these maps in high dimensions and the asymptotic performance of the associated algorithms.

- Functional and harmonic analysis: a natural analytical framework to study oversampled coarse quantization lies within the theory of frames. While the quantization problem in an orthonormal basis is trivial, there is no corresponding theory for redundant representations. Progress in this direction can also have far reaching consequences in compression and denoising.

3 Presentation Highlights

As noted earlier, one of the highlights of the workshop was the diversity of the scientific disciplines represented by the participants. Ingrid Daubechies and Ronald DeVore, two leading mathematicians in applied mathematics and approximation theory, were joined by experts in various disciplines, including John Benedetto and Yang Wang (harmonic analysis), Peter Casazza (functional analysis, Banach spaces, frame theory), Jan Allebach and Chai Wah Wu (halftoning, image processing), Tomasz Nowicki (dynamical systems), Benjamin Doerr (combinatorics, discrete mathematics), Vinay Vaishampayan, Helmut Boelcskei and Vivek Goyal (information theory, signal processing), Matt Yedlin (electrical engineering, circuit theory). This diversity helped create a productive and intellectually inspiring environment for the workshop.

Below is a list of talks presented at the workshop grouped with respect to subject.

3.1 Sigma-Delta Quantization and A/D Conversion

The talks in this category focused on sigma-delta quantization algorithms and their use as an A/D conversion method, emphasizing the mathematical as well as practical aspects.

The workshop started with an expository opening talk by *Ingrid Daubechies*, titled “Mathematical study of coarsely quantized representations of signals: what, why, how – an overview.” In this talk, Daubechies discussed the major practical challenges in A/D conversion, and she formulated the corresponding mathematical problems. Then, she gave a survey of the main results in this direction obtained since 1998, referring to work she has done jointly with Ron DeVore, as well as several results of Güntürk, Thao, Powell, Vaishampayan, and Yılmaz.

The next talk in this category was presented by *Thao Nguyen* and entitled “The Tiling Phenomenon of Sigma-Delta Modulators.” As of now, sigma-delta quantization is the only known and efficient algorithm for the coarse quantization of redundant expansions. While this method has proved successful in practice, its mathematical analysis is difficult. This is due to the fact that sigma-delta modulation is based on a nonlinear dynamical system. In general, one does not know how to derive explicitly the expression of the output of such a system. It was discovered however that the quantizers of an important generic family of sigma-delta modulators possess an outstanding property called “tiling”, which provides substantial algebraic information about their outputs. Nguyen gave a retrospective on this research from its origin to current investigations. He covered several issues including the mathematical origin of tiling and the consequence of tiling for the rigorous prediction of the quantization error.

Matt Yedlin focused on the practical aspects of sigma-delta algorithms in his talk “Industrial Applications of Sigma-Delta Converters.” Yedlin emphasized the major role sigma-delta based A/D conversion plays in modern technological applications including digital audio receivers, sampling music synthesizers, biomedical data acquisition (EEG), seismometers and wireless communication systems. After a historical account of the subject, which included the first patent filed about sigma-delta converters, Yedlin focused on the current industrial applications of sigma-delta A/D converters over a wide range of frequency and resolution settings. Yedlin then profiled a number of currently available commercial products, ranging from the 19 Hz 24 bit AD (Analog Devices) 7783 which is used for pressure and temperature sensing to a high-speed sigma-delta converter, AD 7725 with an input bandwidth of 350kHz, used for radar and sonar data acquisition. Finally, Yedlin illustrated some of the practical issues with a hardware demonstration of the AD7725 sigma-delta converter.

Bin Han, in his talk titled “Time Average MSE Analysis for the First Order Sigma-Delta Modulator,” presented improved estimates for the mean square error (MSE) of first-order sigma-delta modulation. The goal of Han's research is to close the gap between the numerically obtained (faster) error decay rates (as a function of the redundancy) and the rigorously obtained (slower) error decay rates. Building upon several results of Daubechies, DeVore and Güntürk, Han proved improved estimates for the MSE under certain additional assumptions on the input signals.

3.2 Frame expansions and quantization

A natural analytical framework to study oversampled coarse quantization lies within the theory of frames. While the quantization problem in an orthonormal basis is trivial, the corresponding problem for redundant representations is highly challenging and mostly unresolved. The talks in this category highlighted the recent progress toward a comprehensive theory of quantization for redundant frame expansions.

The first talk was presented by *John Benedetto* on "Sigma-Delta quantization and finite frames". Benedetto focused on first-order sigma-delta schemes in the setting of finite frames, and presented several error estimates that prove that sigma-delta quantizers outperform pulse code modulation (PCM) schemes when the frame is sufficiently redundant. His presentation emphasized the techniques for obtaining refined error estimates based on analytic number theory tools (e.g., [10]). The theory presented in Benedetto's talk and its extension to higher order schemes is a collaboration with Alex Powell and Özgür Yılmaz. The analytic number theoretic portion also includes Aram Tangboondouangjit in the collaboration.

Mark Lammers's talk "Alternate Dual Frames and Sigma-Delta Quantization" focused on a joint work with Alex Powell and Özgür Yılmaz. Lammers investigated the use of alternate dual frames in sigma-delta quantization of finite frame expansions, and proved that reconstruction with alternate dual frames can substantially reduce quantization error in many settings, including pointwise and MSE error estimates for higher order sigma-delta schemes. In particular, Lammers showed that using alternate dual frames is an effective way of dealing with the "boundary term" problem, and allows for k-th order sigma-delta schemes to achieve pointwise error of order N^{-k} when the canonical dual frame may not yield the same result.

Next, *Peter Casazza*, a leading expert in frame design, gave a stimulating talk in which he identified several frame design problems and conjectures for sigma-delta quantization. This talk provoked many questions and comments from the participants, helping the participants organize and isolate various critical issues and relevant problems.

In his talk "Sigma-Delta Quantization and the Traveling Salesman Problem," *Yang Wang* considered sigma-delta quantization in the finite frame expansion setting with the maximal and mean square errors as measures of quantization distortion. Wang showed that this problem is related to the classical Traveling Salesman Problem (TSP) in the Euclidean space. Using the fact that sigma-delta error bounds depend on ordering of frame elements and incorporating some *a priori* bounds for the Euclidean TSP, Wang showed that, in general, sigma-delta modulation is superior to PCM in the quantization error performance. Wang also gave a recursive algorithm for finding an ordering of the frame elements that leads to good maximal error *and* mean square error at the same time.

The next presentation in this category was by *Vern Paulsen* titled "Frame Paths and Sigma-Delta Quantization," a joint work with Bernhard Bodmann. Building on the ideas of Benedetto-Powell-Yılmaz, Paulsen studied the performance of finite frames for the encoding of vectors by applying sigma-delta quantization algorithms to the sequence of frame coefficients. Paulsen's focus was on frame paths which allow the user to systematically increase the redundancy of the frame to compensate for the errors created by quantization. Paulsen showed that the earlier bounds for the worst-case quantization error can be improved. Moreover, he gave lower bounds on the worst-case error. In addition, Paulsen introduced some new frame paths, which can in turn be used to construct finite frames with arbitrarily high redundancy that have some interesting features.

Bernhard Bodmann gave the next talk "Zero-terminated frame paths and second order sigma-delta quantization" on a joint work with Vern Paulsen, which extended the ideas presented in Paulsen's talk to higher order sigma-delta schemes. In particular, Bodmann studied the performance of finite frames for the encoding of vectors by applying standard second order sigma-delta quantization to the frame coefficients, where the frames under consideration are obtained from regular sampling of a frame path in a Hilbert space. In order to achieve error bounds that are comparable to the results for second-order sigma-delta quantization of oversampled bandlimited functions, frame paths that terminate in the zero vector were constructed.

Shidong Li talked about "A Dual Frame Formula and Pseudoframes with Applications." In his talk, Li obtained a dual frame formula with which optimal duals can be obtained through a parametric sequence of functions. Next, Li considered *pseudoframes for subspaces* (PFFS), an extension of the notion of frames. PFFS are collections similar to a frame for a subspace X in H. Unlike the elements of a frame for X, however, the elements of a pseudoframe for X need not be contained in X. Li discussed characterizations, constructions, and several key properties of PFFS, and gave examples that illustrate the benefits of extending the concept of frames. Using the theory of PFFS, Li constructed compactly supported and/or fast decaying dual Gabor functions as well as arbitrary compactly supported (pseudoframe) biorthogonal duals of a B-spline Riesz sequence. Other highlights of Li's presentation include the existence of tight pseudo-duals of frames of translates, the illustration of how PFFS can be incorporated in the construction of biorthogonal wavelets so as to obtain results that are more favorable, how PFFS can be exploited in a quite general optimal noise suppression problem, and finally, examples where pseudo-dual or dual frame formula can be used to reduce perturbation and/or quantization noise in a general redundant frame or pseudoframe representation.

In her talk "Fusion Frames: Redundant Representations under Distributed Processing Requirements," *Gitta Kutyniok* considered a different generalization of the concept of a frame for a Hilbert space. The focus of her collaboration with Peter Casazza and Shidong Li was on applications under distributed processing requirements such as sensor networks or sensorineural systems. Kutyniok noted that for such applications frames can be used locally, but the global structure cannot be handled by using classical frame theory. She then introduced the notion of a fusion frame, which is a sequence of subspaces satisfying a frame-like property. This property was used to ensure that local collections of frame elements linked together by using a fusion frame yield a global frame structure. In this sense the new theory of fusion frames provides the link between distributed and centralized structures. Kutyniok showed that the theory of fusion frames has in fact many similarities with the classical frame theory for sequences of vectors. Moreover, this theory can be used to ease the construction of frames by considering local structures. An application to noise reduction under distributed processing requirements was also discussed.

Helmut Boelcskei's talk was on "Noise shaping and predictive quantizers of order $N > 1$ for arbitrary frame expansions." After briefly reviewing the classical paper by Tewksbury and Hallock on "Oversampled, linear predictive and noise-shaping coders of order $N > 1$," Boelcskei discussed his work with Hlawatsch that extends of the ideas of Tewksbury and Hallock to the case of oversampled filter banks. This work was then used to outline how noise shaping and linear predictive coders can be applied to arbitrary frame expansions. In the last part of his talk, Boelcskei described such an extension and provided a proof of the reconstruction MSE decaying as r^{-2N+1} where r denotes the frame redundancy and N is the order of the noise shaping filter.

"PCM Quantization Errors and the White Noise Hypothesis" by *David Jimenez* was the last talk in this category. The White Noise Hypothesis (WNH), introduced by Bennett approximately a half century ago, assumes that in PCM quantization scheme the errors in individual channels behave like white noise, i.e., they can be modeled as independent and identically distributed random variables. The WNH has been the key in estimating the mean square quantization error in the case of PCM. In this joint work with Yang Wang, Jimenez took a close look at the validity of the WNH. He proved that in a redundant system the errors from individual channels can never be independent, implying that, strictly speaking, WNH is not valid. Jimenez also presented numerical experiments which confirm this theoretical result whenever the quantization is coarse, i.e., the quantizer step size is large. On the other hand, Jimenez showed that with fine quantization, the WNH is essentially valid, in that the errors from individual channels become asymptotically *pairwise* independent and each uniformly distributed in $[-\Delta/2, \Delta/2)$, where Δ denotes the stepsize of the quantization.

3.3 Digital Halftoning

While today's printers can achieve high resolution in space (such as 1200 dpi), most often, only single tone dots are available to be printed on the corresponding medium. In this case, only 0's and 1's can be used to represent efficiently any shade of the gray-scale. Error diffusion can be thought of the exact analog of sigma-delta modulation in two dimensional signals and achieves this conversion effectively. One of the goals of this workshop was to provide a platform where researchers in the fields of audio quantization and digital halftoning could meet. The talks listed in this category are on digital halftoning.

The first talk in this category was "Digital halftoning - a model-based perspective" by *Jan Allebach*. Allebach started his talk with an overview of the basic principles of digital halftoning algorithms, which are used in virtually all printing devices and many display systems as well, in order provide a visually pleasing rendering of continuous-tone images that can be stably reproduced by the output device. He then described the two dominant desktop printing technologies, inkjet and electrophotography, and discussed how the physical characteristics of these technologies impact the preferred choice of halftoning algorithm. Next, he discussed hardware and software architectures used in desktop printing systems, and showed how these architectures impact the computational approaches that can be used for digital halftoning. After this extensive review, Allebach described how model-based approaches can yield high quality halftoning algorithms that perform robustly on the intended platform, subject to the constraints on computational architecture and resources. He first discussed the direct binary search (DBS) algorithm that minimizes a cost function based on visual quality, and represents a gold standard of sorts for dispersed-dot halftoning. Allebach noted that the DBS algorithm is heuristic in nature, however, it is still possible to show that it simultaneously minimizes mean-squared and maximum error. Furthermore, DBS is too computationally demanding to be used in desktop printing applications; but it serves as the basis for the design of other algorithms that are implementable with real systems. Then Allebach focused on two such solutions: tone dependent error diffusion, which achieves nearly the same quality as DBS, and can be found in inkjet products, and the hybrid screen which is suitable for electrophotographic engines. The hybrid screen creates stochastic, dispersed-dot textures in highlights that gradually transform to clustered-dot, periodic textures in mid-tones. Allebach concluded his talk with a discussion of measures for halftone image quality where he proposed a new approach to assessing image quality based on the directional sequency spectrum.

Tomasz Nowicki's talk was titled "On the existence of a minimal attractor in Convex Dynamics," joint work with Charles Tresser. The talk was on the convex dynamics concerning a family of the piecewise affine maps (translations) which were inspired by some applications in printing and other analog-to-digital (or continuous-to-discrete) coding of sequences of signals which employ error diffusion. Here, the pieces are Voronoi regions of the corners (sets of points for which such a corner is the closest one) of some polytope(s) and the translation vectors are the vectors from the respective corners to arbitrary points of the polytope. The fundamental theorem Nowicki presented was on the boundedness of the dynamics. He showed that there exists a unique minimal bounded invariant set which contains the polytope itself and all the corners of the Voronoi regions.

In his talk "Discrepancy and Digital Halftoning," *Benjamin Doerr* first gave a brief introduction to the field of discrepancy theory. Doerr explained various notions of discrepancy, such as combinatorial, geometric, and linear discrepancy, and tied these notions with hypergraph colouring problems. He then showed how Asano et al [16] modeled the digital halftoning problem as a linear discrepancy problem. Doerr proposed a novel solution to this problem based on a technique he introduced, which he called "dependent randomized rounding". In his solution, Doerr is able to obtain probabilistic error bounds as well as deterministic ones.

Chai Wah Wu presented his work during his talk titled "Some techniques for speeding up digital halftoning algorithms." In the first part of his talk, Wu described some techniques to speed up two high-quality digital halftoning algorithms: Direct Binary Search (DBS) and error diffusion. Noting that DBS operates by considering pairs of halftone pixels to swap and does so when the swap reduces the overall error, Wu showed that one can reduce the number of operations needed to evaluate and update a potential swap of pixels in each step of DBS by extracting some redundancy in the computation. He also considered some heuristics for selecting which pairs of pixels to evaluate that reduce the total number of pairs evaluations. In the second part of the talk, Wu considered the use of lookup tables to speed up the error diffusion algorithm. In particular, he presented an implementation of the Floyd-Steinberg error diffusion algorithm where the computation requires only 1 addition and 3 table lookup operations per pixel. Based on this, Wu concluded that it may be possible to implement error diffusion-like algorithms at speeds similar to screening algorithms.

3.4 Sparse expansions and quantization

The role of sparse expansions in signal processing and information theory has been understood recently, with the leading works of Candès, Donoho, Gilbert, Strauss and Tropp. The talks in this category are on the theory of sparse expansions and the related quantization issues.

In his talk "Instance-optimal estimates in Compressed Sensing," *Ronald DeVore* first gave an extensive review of the recently developing theory of compressed sensing. In DeVore's words, "Discrete Compressed Sensing samples a discrete signal $x \in \mathbb{R}^N$ by n linear measurements, each an inner product of x with a vector from $\mathbb{R}^N$. If n is the number of measurements allocated to the sensor then the whole process can be represented by an $n \times N$ matrix Φ. The vector $y = \Phi(x)$ represents the n samples we have about x. A decoder Δ is a mapping from $\mathbb{R}^n \to \mathbb{R}^N$. The vector $\Delta(\Phi(x))$ is the approximation we have to x from the information y." In his talk, DeVore discussed how well such an encoding-decoding scheme can perform given the pair n and N by comparing the error $\|x - \Delta(\Phi(x))\|_{\ell_p}$ to the best k-term approximation error $\sigma_k(x)_{\ell_p}$.

Next, *Vivek Goyal* talked on "Sparsity and Quantization: When Undersampling is Oversampling." Goyal noted that in a variety of applications, sparse solutions to underdetermined linear systems of equations are preferred, and for at least thirty years, practitioners have used 1-norm minimizations to promote sparsity. Recent research has turned this ad hoc procedure into the basis for certain deterministic and probabilistic guarantees, fueling further work in sparsity-based modeling. In his talk, Goyal related sparsity-based modeling to his earlier work on quantized overcomplete expansions, showing that large improvements are obtained by exploiting boundedness of quantization error. Goyal also discussed universality of such source coding techniques.

Holger Rauhut considered the problem of reconstructing a sparse multivariate trigonometric polynomial from few sample values in his talk "Random Sampling of Sparse Trigonometric Polynomials". Rauhut investigated the problem in a probabilistic framework in which he modeled the sampling points as being random and independently distributed according to one of the following two probability models: (a) continuous uniform distribution on the unit cube, (b) discrete uniform distribution on a finite set of equidistant sampling nodes. The latter corresponds to the problem of recovering a sparse vector from few samples of its discrete Fourier transform. Rauhut studied two recovery methods: Basis pursuit (BP), which was studied recently by Candès, Romberg, and Tao in the context of the discrete Fourier transform, and orthogonal matching pursuit (OMP), a greedy algorithm, studied recently by Tropp et al. for the recovery problem in the context of Gaussian and Bernoulli measurements. Rauhut then presented a result that applies simultaneously to both of the probability models (a) and (b) and to both BP and OMP. His theorem states that if the number of samples N is large enough compared to the sparsity (but possibly much smaller than the overall dimension of the underlying space of trigonometric polynomials) then with high probability the polynomial can be recovered both by BP and OMP. Rauhut noted that this includes a previous result of Candès, Romberg and Tao as a special case (with a different proof). Finally, Rauhut provided numerical experiments which confirm the proposed methods.

The last talk in this category was "The diffusion framework: a computational approach to data analysis and signal processing on data sets" by *Yosi Keller*. Keller noted that the diffusion framework is a computational approach to high dimensional data analysis and processing. Based on spectral graph theory, Keller defined diffusion processes on data sets, then noted that these agglomerate local transitions reflect the infinitesimal geometries of high-dimensional datasets, and finally used these observations to obtain meaningful global embeddings. The eigenfunctions of the corresponding diffusion operator (Graph Laplacian) provide a natural embedding of the sets into a Euclidean space. Keller then showed that the eigenfunctions of the Laplacian form manifold adaptive bases, which pave the way for the extension of signal processing concepts and algorithms from n-dimensional Euclidean space $\mathbb{R}^n$ to general data sets. As an example, Keller applied this approach to collaborative filtering and improving the resolution of coarsely quantized signals.

4 Outcome of the Meeting and Scientific Progress Made

The organization of this meeting was motivated by the increasing interest in the mathematical and electrical engineering community on coarse quantization strategies for redundant representations of signals. The exclusively positive feedback received from the participants emphasized the scientific appropriateness of the meeting as well as the exciting nature of doing interdisciplinary research. The superb location and facilities of the BIRS was also a big bonus. New collaborations have started (e.g., a joint mathematics-circuits initiative between UBC, Princeton, NYU, and Georgia Tech; joint work by Doerr, Güntürk and Yılmaz) and further meeting ideas were developed (e.g., a summer school for mathematicians and circuit engineers). Compared to some of the previous similar meetings and focused research activities on this topic over the past years, this workshop has been one of the largest and most extensive. It is expected that the trend will continue and more meetings will be organized in the near future. (One of the invitees of the meeting, who unfortunately could not attend, is currently planning the organization of an ICASSP special session for 2007 on the same subject.)

References

[1] J. Candy and G. Temes, Eds., *Oversampling Delta-Sigma Data Converters: Theory, Design, and Simulation*, Wiley-IEEE Press, 1992.

[2] I. Daubechies and R. DeVore, Reconstructing a bandlimited function from very coarsely quantized data: A family of stable sigma-delta modulators of arbitrary order, *Annals of Mathematics*, **158 (2)** (2003), 679–710.

[3] A.R. Calderbank and I. Daubechies, The pros and cons of democracy, *IEEE Transactions on Information Theory*, **48 (6)** (2002) 1721–1725.

[4] R. Ulichney, *Digital Halftoning*, MIT Press, 1987.

[5] H. Inose and Y. Yasuda, A unity bit coding method by negative feedback, *Proceedings of the IEEE*, **51 (11)** (1963) 1524–1535.

[6] R. Floyd and L. Steinberg, An adaptive algorithm for spatial greyscale, *Proc. SID*, **17 (2)** (1976), 75–77.

[7] R.M. Gray, W. Chou and P.W. Wong, Quantization noise in single-loop sigma-delta modulation with sinusoidal inputs, *IEEE Transactions on Communications*, **37 (9)** (1989), 956–968.

[8] R.M. Gray, Quantization noise spectra, *IEEE Transactions on Information Theory*, **36 (6)** (1990), 1220–1244.

[9] C.S. Güntürk and N.T. Thao, Ergodic dynamics in sigma-delta quantization: Tiling invariant sets and spectral analysis of error, *Advances in Applied Mathematics*, **34** (2005), 523–560.

[10] C.S. Güntürk, Number theoretical error estimates in a quantization scheme for bandlimited signals. In *Unusual Applications of Number Theory*, DIMACS: Series in Discrete Mathematics and Theoretical Computer Science, **64**, 81–94, American Mathematical Society, 2004.

[11] J. Benedetto, A. Powell and Ö. Yılmaz, Sigma-delta quantization and finite frames, *IEEE Transactions on Information Theory*, **52 (5)** (2006), 1990–2005.

[12] I. Daubechies, *Ten Lectures on Wavelets*, SIAM, Philadelphia, 1992.

[13] V. Goyal, J. Kovačević, and J. Kelner, Quantized frame expansions with erasures, *Appl. Comput. Harmon. Anal.*, **10** (2001), 203–233.

[14] B. Hochwald, T. Marzetta, T. Richardson, W. Sweldens, and R. Urbanke, Systematic design of unitary space-time constellations, *IEEE Transactions on Information Theory*,**46 (6)** (2000) 1962–1973.

[15] Y. Eldar and G. Forney, Optimal tight frames and quantum measurement, *IEEE Transactions on Information Theory*, **48 (3)** (2002) 599–610.

[16] T. Asano, N. Katoh, K. Obokata, and T. Tokuyama, Matrix rounding under the Lp-discrepancy measure and its application to digital halftoning, *SIAM Journal on Computing* **32 (6)** (2003), 1423–1435.