Simultaneous and Hybrid beta-Encodings

C. Sinan Güntürk
Courant Institute of Mathematical Sciences
New York University
New York, NY 10012
Email: gunturk@courant.nyu.edu

***Abstract*—In this paper we will describe a constructive method to find β-encodings (β-representations) with special properties. These include simultaneous β-encodings with respect to several bases and hybrid $\Sigma\Delta/\beta$-encodings. The main motivation for the latter scheme is to have bit representations whose robustness to additive circuit noise is tunable.**

I. Introduction

Consider the standard binary (base-2) representation of a real number $x \in [0,1]$:

$$x = \sum_{n=1}^{\infty} b_n 2^{-n}, \quad b_n \in \{0,1\}. \tag{1}$$

While it is possible to think of each $b_n = b_n(x)$ as a separate function on $[0,1]$, typically these bits are extracted sequentially, using a recursive operation. One popular method is *successive approximation*. Let x_N be the approximation at the Nth step, which equals the N-bit truncation of the infinite series in (1)

$$x_N = \sum_{n=1}^{N} b_n 2^{-n}, \quad N = 1, 2, \dots, \tag{2}$$

with $x_0 = 0$. By setting

$$b_n = \begin{cases} 1, & \text{if } x - x_{n-1} \geq 2^{-n}, \\ 0, & \text{if } x - x_{n-1} < 2^{-n}, \end{cases} \tag{3}$$

the algorithm narrows down the range of x to a sub-interval whose length is halved after each iteration.

The above algorithm has to work around with an exponentially decreasing sequence of errors and thresholds. Instead, the following variation keeps all quantities macroscopic at all times:

Define $u_n := 2^n(x - x_n)$ for $n \geq 0$. It is easy to see that the sequence $(u_n)_0^\infty$ satisfies the recurrence relation

$$u_n = 2u_{n-1} - b_n, \quad n = 1, 2, \dots, \tag{4}$$

and the bits can simply be extracted via the formula

$$b_n = \lfloor 2u_{n-1} \rfloor = \begin{cases} 1, & u_{n-1} \geq 1/2, \\ 0, & u_{n-1} < 1/2. \end{cases} \tag{5}$$

The relation (4) is the *doubling map* in disguise; $u_n = T(u_{n-1})$, where $T : u \mapsto 2u \pmod 1$.

The base-2 representation essentially provides the most efficient encoding in the rate-distortion sense. Moreover, as seen above, there is a nice algorithmic circuit implementation for the purpose of analog-to-digital (A/D) conversion in which all quantities (the u_n and the b_n) are bounded and macroscopic (i.e., can be held as measurable electric charges). Hence, the case appears to be closed, at least in theory. However, in practice, the successive approximation algorithm (or the base-2 representation in general) hardly finds itself as the most popular choice of A/D conversion method. In real life, analog circuits are never perfect, suffering from arithmetic errors (e.g., through nonlinearity) as well as from quantizer errors (e.g., threshold offset), while being subject to thermal noise at the same time. All relations hold approximately, and therefore, all quantities are approximately equal to their theoretical values (and in this case, only for a finite number of iterations, given that dynamics of an expanding map has "sensitive dependence on initial conditions"). Note that this problem affects analog circuits mostly, as digital circuits are much more immune to small variations. Given that this is a central problem in A/D conversion (as well as in D/A conversion), many alternative bit representations of numbers (as well as of signals) have been devised in circuit engineering, such as $\Sigma\Delta$ modulation.

From a theoretical point of view, the imprecision problem associated to the successive approximation algorithm is not enough to discard the base-2 representation immediately. After all, we do not have to use this specific algorithm to extract the bits, and perhaps there are better (i.e., more resilient) algorithms for evaluating $b_n(x)$ for each x. However, the root of the real problem lies deeper in the fact that the bits in the base-2 representation are essentially uniquely determined, and are ultimately computed by a greedy method. Since $2^{-n} = 2^{-n-1} + 2^{-n-2} + \dots$, we have no choice other than to set b_n according to (3), except in the equality case, which happens if and only if x is a dyadic rational at scale n, i.e., $x = k2^{-n}$ where k is an odd integer. In this case, we have the option between choosing $b_n = 1$ followed by an infinite string of 0s, or $b_n = 0$ followed by an infinite string of 1s. It is clear that there is no way to recover from an erroneous bit computation: a 1 assignment for b_n when $x < x_{n-1} + 2^{-n}$ immediately means an "overshoot" from which there is no way to "back up" later. Similarly a 0 assignment for b_n when $x > x_{n-1} + 2^{-n}$ implies a "fall-behind" from which there is no way to "catch up" later.

Due to this lack of robustness, the base-2 representation is not the preferred quantization method for A/D conversion. For similar reasons, it is also generally not the preferred method for D/A conversion. In practical settings, oversampled coarse quantization ($\Sigma\Delta$ modulation) is more popular [1][12], mostly due to its robustness achieved with the help of the redundant set of output codes that can represent each source value [8]. However, standard $\Sigma\Delta$ modulation is suboptimal as a quantization method (even though exponential accuracy in the bit rate can be achieved [7]).

A partial remedy comes with fractional base expansions, called β-expansions (β-representations) [2][4][6][11][13]. Fix $1 < \beta < 2$. It is well known that every x in $[0, 1]$ (in fact, in $[0, (\beta - 1)^{-1}]$) can be represented by an infinite series

$$x = \sum_{n=1}^{\infty} b_n \beta^{-n}, \tag{6}$$

for appropriate choice of the bit sequence (b_n). Clearly, the N-bit truncated approximations $\sum_{n=1}^{N} b_n \beta^{-n}$, which we will again denote by x_N, are only accurate to within $O(\beta^{-N})$, which is inferior to the accuracy of a base-2 representation. However, there is a crucial difference in that many distinct β-representations in the form (6) are now available. In fact, it is known that for any $\beta < 2$, almost all numbers (in the Lebesgue measure sense) have uncountably many distinct β-representations [14]. Among these, two of them are special [2][4]:

The *greedy selection* turns on the bits as soon as possible, i.e., sets

$$b_n = 1 \quad \text{iff} \quad x - x_{n-1} \geq \beta^{-n}. \tag{7}$$

On the other hand, the *lazy selection* waits until the last moment to turn on the bits, i.e., a bit is turned on only if the remaining bits cannot possibly catch-up later:

$$b_n = 1 \quad \text{iff} \quad x - x_{n-1} > \beta^{-n-1} + \beta^{-n-2} + \ldots, \tag{8}$$

and in general, these two selections result in different β-representations because for $\beta < 2$ we have

$$\beta^{-n-1} + \beta^{-n-2} + \cdots > \beta^{-n}.$$

This implies that if the greedy selection results in $b_n = 1$ followed by a sufficiently long stretch of 0s, then we have the option to set $b_n = 0$ as long as we follow it by a sufficiently long stretch of 1s. (For $\beta = 2$, "sufficiently" means "infinitely.") This redundancy makes β-representations an appealing alternative since it is possible to recover from (certain) incorrect bit computations. In fact, if these mistakes result from an unknown threshold offset in the quantizer, then it turns out that an intermediate *cautious selection* algorithm is completely robust provided a bound for the offset is known [4]. To make this statement more precise, consider the following algorithmic (dynamical) implementation:

Similar to before, define $u_n = \beta^n(x - x_n)$. Then we can extract from (6) the recursive relation

$$u_n = \beta u_{n-1} - b_n. \tag{9}$$

The greedy selection (7) can now be rewritten as

$$b_n = \begin{cases} 0, & u_{n-1} < 1/\beta, \\ 1, & u_{n-1} \geq 1/\beta. \end{cases} \tag{10}$$

Similarly, the lazy selection (8) is equivalent to

$$b_n = \begin{cases} 0, & u_{n-1} \leq 1/\beta(\beta{-}1), \\ 1, & u_{n-1} > 1/\beta(\beta{-}1), \end{cases} \tag{11}$$

which differs from the greedy selection in its threshold value and the specific choice made at the threshold value.

Now, the cautious selection starts with any two values a and b such that

$$1/\beta < a < b < 1/\beta(\beta{-}1), \tag{12}$$

and sets

$$b_n = \begin{cases} 0, & u_{n-1} \leq a, \\ 0 \text{ or } 1, & u_{n-1} \in (a, b), \\ 1, & u_{n-1} \geq b. \end{cases} \tag{13}$$

It can be easily seen that the precise values of a and b as well as the specific choice made in the interval (a, b) are unimportant in the sense that one always obtains a β-representation once it is guaranteed that the sequence (u_n) remains bounded.[1] This is the case even if a and b vary at each iteration while satisfying (12). Hence perfect encoding is possible with an imperfect (flaky) quantizer whose threshold value fluctuates in the interval $(1/\beta, 1/\beta(\beta - 1)$.

The philosophical reason for this robustness is the redundancy of the encoding. Our goal in this paper is to exploit the redundancy of β-representations further. More specifically, we will show constructively the existence of

- bit representations that yield multiple decodings (simultaneous β-encodings),
- β-representations with tunable asymptotic digit frequency (hybrid $\Sigma\Delta/\beta$-encodings).

We note that both of these objectives will be achieved in the setting of *signed digit expansions*, i.e., for $b_n \in \{-1, +1\}$ instead. This is mostly for technical convenience rather than a strict requirement since any $\{0, 1\}$ digit expansion can be mapped to a $\{-1, +1\}$ digit expansion and vice versa, once the range of x is suitably adjusted via $x \mapsto 2x - (\beta - 1)^{-1}$.

II. Simultaneous β-Encodings

Let us ask the following question: given (any) β_1 and β_2, is it possible to encode $x \mapsto (b_n)_1^\infty$ such that

$$x = \sum_{n=1}^{\infty} b_n \beta_1^{-n} = \sum_{n=1}^{\infty} b_n \beta_2^{-n} \ ? \tag{14}$$

If we were to seek $b_n \in \{0, 1\}$, this task would clearly be impossible to achieve in this form (i.e., without any adjustment in the reconstruction formula) since the map $t \mapsto \sum b_n t^n$ is strictly increasing. This obstacle is not present when $b_n \in \{-1, +1\}$, and the answer turns out to be positive:

Theorem 1: There exists $c > 0$ *such that for all* $1 < \beta_1, \beta_2 < 1+c$, *every* $x \in [-1, 1]$ *has a simultaneous* (β_1, β_2)*-representation, i.e., (14) holds.*

Note that redundancy of β-representations increases as $\beta \to 1$. It turns out that much more is true, in that we can pick several β_i and different x_i values:

[1] Note that since (b_n) is bounded, (9) implies that (u_n) either remains bounded (when $x_n \to x$) or blows up like β^n (when $x_n \not\to x$).

Theorem 2: Given any positive integer M, there exists $c = c_M > 0$ such that, for all

$$1 < \beta_1 < \beta_2 < \cdots < \beta_M < 1 + c,$$

there exists $\delta = \delta(\beta_1, ..., \beta_M) > 0$ such that, for all

$$(x_1, x_2, \ldots, x_M) \in [-\delta, \delta]^M,$$

there exists $(b_n)_1^\infty \in \{-1, +1\}^{\mathbb{N}}$ that simultaneously encodes all the x_i, i.e.,

$$x_i = \sum_{n=1}^{\infty} b_n \beta_i^{-n}; \quad i = 1, ..., M.$$

Remarks:

1. With regards to the comment following (14), the theorem continues to hold in the $0/1$ world, but for a translated copy of $[-\delta, \delta]^M$.

2. Not surprisingly, the result is for β_i's sufficiently close to 1. A necessary condition on $\beta_1, \ldots, \beta_M$ is that

$$\beta_1 \beta_2 \cdots \beta_M \leq 2, \tag{15}$$

so that c in the statement must satisfy $c \leq 2^{1/M} - 1$. There is a simple volume covering argument for this condition. Consider all N-bit truncated vectors

$$\mathbf{x}_N = \sum_{n=1}^{N} b_n \left(\beta_1^{-n}, \ldots, \beta_M^{-n}\right), \tag{16}$$

which provide us with at most 2^N distinct points in $[-\delta, \delta]^M$. Given the N-term error bound $\beta_i^{-N}(\beta_i - 1)^{-1}$ in the ith coordinate, we have that each sub-rectangle in $[-\delta, \delta]^M$ of dimensions $d_1 \times d_2 \times \cdots \times d_M$, where $d_i = 2\beta_i^{-N}(\beta_i - 1)^{-1}$, must contain such an N-term truncated vector. Hence by the pigeon-hole principle, we have

$$2^N \geq \frac{2^M \delta^M}{d_1 \cdots d_M} \geq \delta^M (\beta_1 - 1)^M (\beta_1 \cdots \beta_M)^N.$$

The result follows after taking the N'th root and letting $N \to \infty$.

3. Theorem 2 does not give precise estimates on the relation between M and c (which determines the range of the β_i). In Figure 1, we plot all N-bit truncated vectors $\mathbf{x}_N$ defined in (16) for $N = 12$, $M = 2$, and two pairs of (β_1, β_2). The first pair $(1.4, 1.8)$ violates the necessary condition (15) and it can be seen that there is an open ball around the origin that appears to not contain any of these vectors. The second pair $(1.2, 1.3)$ does not violate (15), and the figure suggests that a neighborhood of the origin is indeed covered.

4. For related results in the context of Iterated Function Systems with overlaps, see [15].

III. Hybrid $\Sigma\Delta/\beta$-Encodings

The first order $\Sigma\Delta$ modulation is governed by the difference equation

$$u_n = u_{n-1} + x - b_n, \tag{17}$$

where

$$b_n = Q(u_{n-1}, x) \tag{18}$$

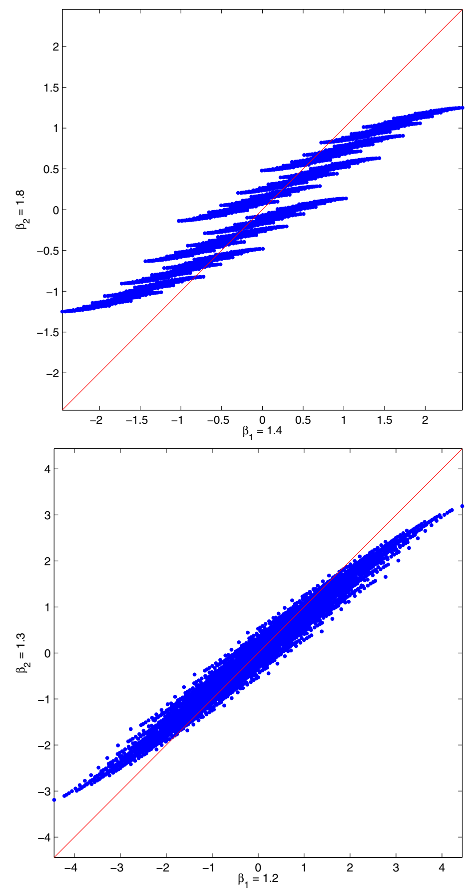

Fig. 1. Plot of all 12-bit truncated vectors $\mathbf{x}_{12}$ in two dimensions for two pairs. The first picture has $(\beta_1, \beta_2) = (1.4, 1.8)$ and the second picture has $(\beta_1, \beta_2) = (1.2, 1.3)$.

for some function Q which makes (u_n) bounded. Here, b_n takes on values in $\{-1, +1\}$ when $x \in [-1, 1]$. The standard example for Q is $Q(u, x) = \text{sign}(u + x)$, which amounts to greedily minimizing u_n given u_{n-1} and x, subject to (17).

The relation (17) and $|u_n| \leq C$ for all n imply that

$$\left| Nx - \sum_{n=1}^{N} b_n \right| \leq 2C; \tag{19}$$

hence, in particular, x is encoded as the average value of (b_n):

$$x = \lim_{N \to \infty} \frac{1}{N} \sum_{n=1}^{N} b_n. \tag{20}$$

Our next theorem says that every x can be assigned a bit sequence which is simultaneously a β-encoding as well as a $\Sigma\Delta$ encoding.

Theorem 3: There exists $c > 0$ such that for all $1 < \beta < 1 + c$, every $x \in [-1, 1]$ has a hybrid $\Sigma\Delta/\beta$ encoding (b_n) in $\{-1, +1\}$ in the sense that

$$x = \sum_{n=1}^{\infty} b_n \beta^{-n}$$

and

$$\left| x - \frac{1}{N} \sum_{1}^{N} b_n \right| \leq \frac{C}{N} \quad \textit{for all } N.$$

Remark: Note that the greedy selection that was mentioned earlier is generated by the map $T_\beta : u \mapsto \beta u \pmod 1$ via

$$u_n = T_\beta(u_{n-1}),$$

and

$$b_n = \lfloor \beta u_{n-1} \rfloor.$$

It is known that T_β is ergodic with respect to a T_β-invariant measure μ_β on $[0,1]$ which is absolutely continuous [13]. Therefore by the Birkhoff ergodic theorem, the digits (b_n) found via the greedy selection would generically have a fixed asymptotic frequency which is independent of x, i.e., for (Lebesgue) almost all x,

$$\frac{1}{N} \sum_{1}^{N} b_n \to c_\beta = \int_0^1 \lfloor \beta u \rfloor d\mu_\beta.$$

Our result implies that the bits in a general β-representation can be chosen in such a way that their asymptotic frequency can be tuned to the number that they represent, thereby providing an alternative description of the number. This property is further elaborated below.

IV. Applications

Consider $\Sigma\Delta$ modulation of a scalar x and its standard N-bit decoder defined by

$$x_N := \frac{1}{N} \sum_{n=1}^{N} b_n.$$

We have $|x - x_N| = O(N^{-1})$, which is highly suboptimal compared to the exponential accuracy of any β-representation. However, $\Sigma\Delta$ modulation has other virtues, such as its robustness against quantizer threshold offset error and variations in the initial condition [8]. It is also robust against certain perturbations of its recurrence relation (17). Let us consider the following noisy form of this relation

$$u_n = u_{n-1} + x - b_n + \xi_n, \tag{21}$$

where (ξ_n) is an i.i.d. sequence of random variables with zero mean. We now have

$$x - x_N = -\frac{1}{N} \sum_{n=1}^{N} \xi_n + \frac{1}{N}(u_N - u_0). \tag{22}$$

For practical reasons, (u_n) needs to be kept bounded. For simplicity, let us also assume that $|\xi_n| \leq \mu$ for some $\mu < 1$, and $|x| \leq 1 - \mu$. Under these assumptions, it is easy to verify by induction that the selection $b_n = \text{sign}(u_{n-1} + x)$ guarantees that $|u_n| \leq 1 + \mu$. The law of large numbers now implies that $x_N \to x$ almost surely, i.e., the standard decoder recovers x in the limit with probability 1.

We can also estimate the rate of convergence. It easily follows from (22) that

$$\mathbf{E}|x - x_N|^2 = O(N^{-1}). \tag{23}$$

By the central limit theorem, $(x - x_N)\sqrt{N}$ converges in distribution to a Gaussian of variance equal to the noise variance. Hence we can think of the error being typically of size $O(N^{-1/2})$. More precisely, the law of the iterated logarithm tells us that almost surely, we have

$$|x - x_N| \leq 2\sqrt{\frac{\log\log N}{N}} \tag{24}$$

for all but finitely many N.

Let us consider the same scenario in the setting of β-encoding. The noisy equation now reads

$$u_n = \beta u_{n-1} - b_n + \xi_n, \quad u_0 = x. \tag{25}$$

Let us assume that the b_n are chosen in $\{-1, +1\}$ (by any method) such that $|u_n| \leq C$ for all n. (Such stable methods exist. For example, if $\mu \leq 2 - \beta$, where μ is an absolute bound on the noise, then it is easy to verify that $b_n = \text{sign}(u_{n-1})$ implies $|u_n| \leq (1 - \mu)/(\beta - 1)$.)

The standard N-bit decoder for β-encoding is given by

$$x_N = \sum_{n=1}^{N} b_n \beta^{-n}.$$

The stability of (u_n) in (25) implies that

$$x - x_N = -\sum_{n=1}^{N} \xi_n \beta^{-n} + O(\beta^{-N}). \tag{26}$$

Note that in the limit $N \to \infty$, the error is a random series which converges almost surely to a random variable ξ given by

$$\xi = -\sum_{n=1}^{\infty} \xi_n \beta^{-n},$$

which has zero mean, but non-zero variance:

$$\sigma_\xi^2 = \sum_{n=1}^{\infty} \sigma_{\xi_n}^2 \beta^{-2n}.$$

Hence, we can conclude that with probability 1, the output of the decoder will be different from x.

This result is in large contrast with $\Sigma\Delta$ encoding in that it is impossible to have perfect β-encodings under additive circuit noise, even though we have superior N-term error decay in the non-noisy case.

A hybrid $\Sigma\Delta/\beta$-encoding, on the other hand, would be able to bring the best of both worlds. It can operate with two different linear decoders on N bits,

$$D_{\Sigma\Delta}(b_1, \ldots, b_N) = \frac{1}{N} \sum_{n=1}^{N} b_n, \quad \text{and}$$

$$D_\beta(b_1, \ldots, b_N) = \sum_{n=1}^{N} b_n \beta^{-n},$$

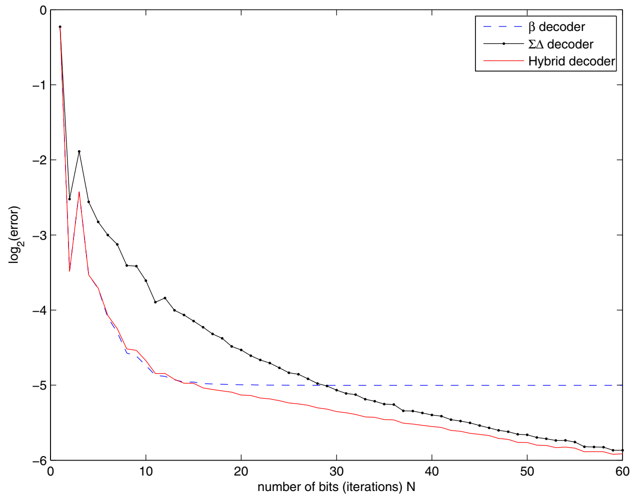

Fig. 2. Error decay rate comparison of the affine decoder (27) with the standard $\Sigma\Delta$ and β-decoders in a hybrid $\Sigma\Delta/\beta$-representation. The accuracy of the β-decoder saturates around the noise level. Here, $\beta = 1.3$, and ξ_n is uniformly distributed in $[-\mu, \mu]$, where $\mu = 0.15$.

and therefore with any affine combination of the two:

$$D = \alpha D_{\Sigma\Delta} + (1-\alpha) D_\beta, \quad \alpha \in \mathbb{R}. \tag{27}$$

The value of the free parameter α can now be optimized based on the noise level bound μ and N. For $N \leq |\log \mu| / \log \beta$, it is natural to pick D_β to take advantage of the exponential decay of error. Once N hits this critical value, then the noise term dominates in (26), and the accuracy of the β-representation saturates. When $N \geq 1/\mu^2$ (which would generally happen later), the accuracy of the $\Sigma\Delta$ representation will reach this saturated level, and therefore switching to $D_{\Sigma\Delta}$ becomes advantageous as the error will continue to decrease (though at a slower rate). In the intermediate range, a careful optimization in fact can do slightly better than both decoders. This is shown in Figure 2.

V. Discussion of the Method

The underlying method for these results is the construction of appropriate difference equations of the form

$$u_n = \sum_{k=1}^{L} h_k u_{n-k} - b_n, \tag{28}$$

or, of the form

$$u_n = \sum_{k=1}^{L} h_k u_{n-k} + x - b_n, \tag{29}$$

with the following properties:

(a) The roots of the characteristic polynomial $P(z) = z^L - \sum_{k=1}^{L} h_k z^{L-k}$ contain $1, \beta_1, \beta_2, \ldots, \beta_M$. In addition, $h_L \neq 0$, and $h_{L-1} = \cdots = h_{L-M} = 0$.

(b) There exist quantization rules for $b_n \in \{-1, +1\}$ such that (u_n) remains bounded; in fact $|u_n| \leq 1$.

Let us see how these two properties are used for the simultaneous (β_1, β_2)-encoding of a pair of numbers (x_1, x_2). Here, $M = 2$. Consider (28) for $n \geq 1$, with initial conditions $u_{-L+1}, \ldots, u_{-1}, u_0$ to be set later. Define $h_0 = -1$, so that property (a) above implies

$$\sum_{k=0}^{L} h_k \beta_{1,2}^{-k} = 0. \tag{30}$$

Take $\beta \in \{\beta_1, \beta_2\}$. Using the boundedness of (u_n) we find that

$$\begin{aligned}
\sum_{n=1}^{\infty} b_n \beta^{-n} &= \sum_{n=1}^{\infty} \sum_{k=0}^{L} h_k u_{n-k} \beta^{-n} \\
&= \sum_{k=0}^{L} h_k \beta^{-k} \sum_{n=-k+1}^{\infty} u_n \beta^{-n} \\
&= \sum_{k=0}^{L} h_k \beta^{-k} \sum_{n=-k+1}^{1} u_n \beta^{-n},
\end{aligned} \tag{31}$$

where we have used (30) in the last step. We now set $u_{-L+3} = \cdots = u_0 = 0$. (Note that by property (a) we necessarily have $L \geq 3$.) This, along with property (a) that $h_{L-1} = h_{L-2} = 0$, reduces the first sum in (31) to the case $k = L$ and the second sum to the two terms $n = -L+1, -L+2$, where the sum over k corresponding to $n = 1$ vanishes again due to (30). Hence, the simultaneous (β_1, β_2)-representation objective boils down to solving the system of equations

$$\begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = h_L \begin{bmatrix} \beta_1^{-1} & \beta_1^{-2} \\ \beta_2^{-1} & \beta_2^{-2} \end{bmatrix} \begin{bmatrix} u_{-L+1} \\ u_{-L+2} \end{bmatrix}. \tag{32}$$

Since $h_L \neq 0$, and $\beta_1 \neq \beta_2$, this system of equations can be solved in the unknown initial conditions (u_{-L+1}, u_{-L+2}) for any pair (x_1, x_2) as long as the system remains stable with these initial conditions. Property (b) says that it suffices to pick the initial conditions in $[-1, 1]$, therefore the image of $[-1, 1]^2$ under the map defined by the RHS of (32) is precisely the set of pairs (x_1, x_2) for which this method succeeds.

Remarks: The method is easily extended to an arbitrary value of M. The properties (a) and (b) can be achieved for all M, by a modification of the method used in [7].

The solution for the hybrid $\Sigma\Delta/\beta$-encodings follow a similar line, but uses the additional zero at $z = 1$, as well as the modified difference equation (29). This method can also be extended to generalize Theorem 3 to an arbitrary order $\Sigma\Delta$ and arbitrary number of distinct β_i's. Details will be provided in a separate publication.

Acknowledgment

The work of the author has been supported in part by the following: National Science Foundation Grant CCF-0515187, Alfred P. Sloan Research Fellowship, NYU Goddard Fellowship.

REFERENCES

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

[2] K. Dajani and C. Kraaikamp, *From greedy to lazy expansions and their driving dynamics*, Expo. Math., 20, (2002), no. 4, pp. 315–327.

[3] I. Daubechies, R. DeVore, *Reconstructing a Bandlimited Function From Very Coarsely Quantized Data: A Family of Stable Sigma-Delta Modulators of Arbitrary Order*, Ann. of Math., vol. 158, no. 2, pp. 679–710, Sept. 2003.

[4] I. Daubechies, R. DeVore, C. S. Güntürk, and V. Vaishampayan, *A/D conversion with imperfect quantizers*, IEEE Transactions on Information Theory, vol. 52, no. 3, pp. 874–885, March 2006.

[5] I. Daubechies, C. S. Güntürk, Y. Wang, and Ö. Yılmaz, *The Golden Ratio Encoder*, in preparation.

[6] I. Daubechies, Ö. Yılmaz, *Robust and practical analog-to-digital conversion with exponential precision*, IEEE Trans. Inform. Theory 52 (2006), no. 8, 3533–3545.

[7] C. S. Güntürk, *One-Bit Sigma-Delta Quantization with Exponential Accuracy*, Comm. Pure Appl. Math., vol. 56, pp. 1608–1630, no. 11, 2003.

[8] C. S. Güntürk, J. C. Lagarias, and V. A. Vaishampayan, *On the Robustness of Single-Loop Sigma-Delta Modulation*, IEEE Trans. Inform. Th., vol. 47, No. 5, pp. 1735-1744, July 2001.

[9] C. S. Güntürk, *Interpolation by Power Series with* ± 1 *Coefficients: Simultaneous Beta-Expansions*, in preparation.

[10] H. Inose and Y. Yasuda, *A Unity Bit Coding Method by Negative Feedback*, Proc. IEEE, vol. 51, pp. 1524-1535, Nov. 1963.

[11] A. N. Karanicolas, H.-S. Lee and K. L. Bacrania, *A 15-b 1-Msample/s Digitally Self-Calibrated Pipeline ADC*, IEEE Journal of Solid-State Circuits, vol. 28, No. 12, pp. 1207-1215, Dec. 1993.

[12] S. R. Norsworthy, R. Schreier, and G. C. Temes, Eds., *Delta-Sigma Data Converters: Theory, Design and Simulation*, IEEE Press, 1996.

[13] W. Parry, *On* β*-expansions of real numbers*, Acta. Math. Acad. Sci. Hungar., vol. 15, pp. 95-105, 1964.

[14] N. Sidorov, *Almost every number has a continuum of* β*-expansions*, Amer. Math. Monthly 110 (2003), no. 9, 838–842.

[15] N. Sidorov, *Combinatorics of linear iterated function systems with overlaps*, Nonlinearity 20 (2007), no. 5, 1299–1312.