LLR-BASED SUCCESSIVE CANCELLATION LIST DECODING OF POLAR CODES

Alexios Balatsoukas-Stimming, Mani Bastani Parizi, Andreas Burg

École Polytechnique Fédérale de Lausanne, CH-1015 Ecublens, Switzerland

ABSTRACT

We present an LLR-based implementation of the successive cancellation list (SCL) decoder. To this end, we associate each decoding path with a *metric* which (i) is a monotone function of the path's likelihood and (ii) can be computed efficiently from the channel LLRs. The LLR-based formulation leads to a more efficient hardware implementation of the decoder compared to the known log-likelihood based implementation. Synthesis results for an SCL decoder with block-length of N = 1024 and list sizes of L = 2 and L = 4 confirm that the LLR-based decoder has considerable area and operating frequency advantages in the orders of 50% and 30%, respectively.

Index Terms— Successive Cancellation List Decoder, Polar Codes, Hardware Implementation

1. INTRODUCTION

Polar codes are a class of capacity-achieving error correcting codes with low-complexity encoding *and decoding* algorithms [1]. Specifically, the successive cancellation (SC) decoder has a structured nature which makes its hardware implementation attractive [2, 3, 4]. Moreover, the SC decoder can be implemented in a numerically accurate and stable way by representing the involved transition probabilities as log-likelihood ratios (LLRs).

Successive cancellation list (SCL) decoding was introduced in [5] to improve the finite block-length performance of polar codes. However, the original description of the SCL decoder is given in terms of the path likelihoods which is not suitable for practical implementation. Hence, the first hardware architecture for SCL decoding [6], uses log-likelihoods (LLs) to partially overcome the numerical stability issues. Unfortunately, the decoder still requires very large and irregular memory elements and processing elements that support large bit-widths, which induce a high cost in terms of hardware resources.

Contribution and Paper Outline: In this paper, we propose an LLR-based formulation of the SCL decoder and show that such a formulation can significantly improve the hardware architecture of [6] by solving all the aforementioned implementation problems. In Section 2, after a brief review of the SCL decoding algorithm, we introduce a *path-metric* which is iteratively updated as a function of the LLR of the bit being decoded given the past trajectory of the decoding path and its (tentative) value. We then prove that this metric is a monotone function of the path's likelihood which yields the LLR-based formulation of the SCL decoder. In Section 3, we provide a short review of the SCL decoder hardware architecture of [6] and compare the synthesis results for the LL- and LLR-based decoders.

Notation: Throughout this paper the boldface letters denote vectors. The elements of a vector \mathbf{x} are denoted as x_i . By \mathbf{x}^m we mean the sub-vector $[x_0, x_1, \ldots, x_m]^T$ if $m \ge 0$ and the null vector otherwise. If $\mathcal{I} = \{i_1, i_2, \ldots\}$ is a set of indices (such that $i_1 < i_2 < \ldots$), then $\mathbf{x}_{\mathcal{I}}$ denotes the sub-vector $[x_{i_1}, x_{i_2}, \ldots]^T$.

2. SC LIST DECODING OF POLAR CODES

A polar code of rate R < 1 and block length $N = 2^n$ is constructed by 'appropriately' choosing a subset $\mathcal{A} \subset \{0, 1, \ldots, N-1\}$ of cardinality $|\mathcal{A}| = NR$ called the *information indices*. The transmitter then constructs the vector $\mathbf{u} \in \{0, 1\}^N$ by putting the NR data bits on $\mathbf{u}_{\mathcal{A}}$ and fixing $\mathbf{u}_{\mathcal{F}}$, where $\mathcal{F} \triangleq \{0, 1, \ldots, N-1\} \setminus \mathcal{A}$, to known-to-receiver *frozen bits*. Subsequently, a codeword $\mathbf{x} = \mathbf{Gu}$ is computed and sent over the channel.¹.

The receiver observes a noisy version of \mathbf{x} denoted as \mathbf{y} and has to decode \mathbf{x} or, equivalently, \mathbf{u} . Since the sub-vector $\mathbf{u}_{\mathcal{F}}$ is already known, the decoder's task reduces to estimating $\mathbf{u}_{\mathcal{A}}$. To this end, Arıkan proposes the SC decoding procedure summarized in Algorithm 1. In line 5 of Algorithm 1, $W_n^{(i)}(\mathbf{y}, \mathbf{u}^{i-1}|u_i)$ represents the likelihood of u_i given the channel output \mathbf{y} and \mathbf{u}^{i-1} considering $u_{i+1}, u_{i+2}, \ldots, u_{N-1}$ as unknown bits.

Algorithm 1: Successive Cancellation Decoding [1].			
1 for $i = 0, 1,, N - 1$ do			
2 if $i \notin \mathcal{A}$ then			
3 $\hat{u}_i \leftarrow u_i;$	// Frozen Bits		
4 else			
$\mathbf{s} \left[\begin{array}{c} \hat{u}_i \leftarrow \arg \max_{u_i \in \{0,1\}} W_n^{(i)}(\mathbf{y}, \hat{\mathbf{u}}^{i-1} u_i) ; \end{array} \right]$			
6 return $\hat{\mathbf{u}}_{\mathcal{A}}$;			

SC decoding is sub-optimal since at each step $i \in A$ the decoder ignores the information it possesses about the future frozen bits $\{u_j : j > i, j \in F\}$. In return for this sub-optimality, the likelihoods $W_n^{(i)}(\mathbf{y}, \hat{\mathbf{u}}^{i-1}|u_i)$ can be computed efficiently and the decoding complexity scales like $O(N \log N)$. Furthermore, Arikan shows that, as long as the code rate is below the channel capacity, there exists $A \subset \{0, 1, \ldots, N-1\}$ for which the block-error probability of this scheme vanishes as N increases [1, 7].

2.1. SC List Decoding

Unfortunately, the sub-optimality of SC decoding is still significant in small-to-moderate block-lengths used in practice. A successive cancellation list (SCL) decoder has been proposed in [5] to partially compensate for this sub-optimally.

In short, SCL decoding is carried out as follows: at each decoding step $i \in A$, instead of fixing a decision on u_i (line 5 of Algorithm 1), two *decoding paths* corresponding to either possible value of u_i are created and decoding is continued in two parallel decoding threads. In order to avoid the exponential growth of the number of decoding paths, at each step only the L most likely paths are retained. Finally, the decoder will end up with a list of L candidates for

¹The structure of **G** permits this computation to be done using $O(N \log N)$ binary additions.

 $\mathbf{u}_{\mathcal{A}}$ out of which the most likely one is declared as the final estimate, $\hat{\mathbf{u}}_{\mathcal{A}}$. This procedure is described in Algorithm 2.

Algorithm 2: SC List Decoding [5].			
1 for $i = 0, 1,, N - 1$ do			
2	if $i \not\in \mathcal{A}$ then		
3	$ \hat{u}_i[\ell] \leftarrow u_i, \forall \ell \in \{0, 1, \dots, L-1\}; $		
4	else		
5	if Less than L paths are active then		
6	Duplicate all the paths and continue with both		
	possible values of u_i ;		
7	else		
8	Sort the likelihoods $\{W_n^{(i)}(\mathbf{y}, \hat{\mathbf{u}}^{i-1}[\ell] u_i) : \forall \ell \in$		
	$\{0, 1, \dots, L-1\}, \forall u_i \in \{0, 1\}\};$		
9	Continue along the L most likely paths;		
10 $\ell^* \leftarrow$ The index of the most-likely path;			
11 return $\hat{\mathbf{u}}_{\mathcal{A}}[\ell^*];$			

Simulation results show that, with relatively small list sizes, the SCL decoder's performance is very close to the optimal ML decoder. More importantly, with a clever choice of the data structures, SCL decoding can be done in $O(L N \log N)$ time complexity [5].

2.2. LLR-based Path Metric Computation

Algorithms 1 and 2 are both valid high-level descriptions of SC and SCL decoding, respectively. However, implementing the decoders using the likelihoods directly is risky as likelihoods become very small numbers and the decoder will be prone to underflow errors.

A practical SC decoder, therefore, receives the channel output in the form of LLRs $(\ln \frac{W(y_i|0)}{W(y_i|1)}, i \in \{0, 1, \dots, N-1\}$ where $W(y|x), x \in \{0, 1\}$ is the channel transition probability) and computes the *decision LLRs*, $L_n^{(i)} \triangleq \ln \frac{W_n^{(i)}(y, \hat{u}^{i-1}|0)}{W_n^{(i)}(y, \hat{u}^{i-1}|1)}$, which are sufficient statistics for decisions in line 5 of Algorithm 1. Furthermore, the computations are numerically stable and the decoding involves $O(N \log N)$ arithmetic operations in total [1, Section VIII].

For the SCL decoder, however, the LLRs seem to be insufficient for choosing the L most likely paths in line 8 of Algorithm 2. In [5] the decoder, therefore, computes a scaled version of the pair of likelihoods $W_n^{(i)}(\mathbf{y}, \hat{\mathbf{u}}^{i-1}[\ell]|u_i), u_i \in \{0, 1\}$ assuming the channel output is provided in the form of *pairs* of likelihoods $(W(y_i|x_i), x_i \in \{0, 1\}, i \in \{0, 1, \ldots, N-1\})$. In order to avoid the underflows, all likelihoods are scaled by a common factor at each intermediate step of the computations. This normalization step is circumvented in the hardware implementation of [6] by performing the computations in the log-likelihood domain.

Luckily, it turns out that the decoding paths can also be ordered according to their likelihoods using only the decision LLRs and the past trajectory of each path as we shall demonstrate in the following.

Theorem 1. For each path $\ell \in \{0, 1, ..., L - 1\}$ and each step $i \in \{0, 1, ..., N - 1\}$ let the path-metric be defined as:

$$\mathsf{PM}_{\ell}^{(i)} \triangleq \sum_{j=0}^{i} \ln \left(1 + e^{-(1-2\hat{u}_{j}[\ell]) \cdot \mathsf{L}_{n}^{(j)}[\ell]} \right), \tag{1}$$

where $\mathsf{L}_{n}^{(i)}[\ell] = \ln \frac{W_{n}^{(i)}(\mathbf{y}, \hat{\mathbf{u}}^{i-1}[\ell]|0)}{W_{n}^{(i)}(\mathbf{y}, \hat{\mathbf{u}}^{i-1}[\ell]|1)}$. If all the information bits are uniformly distributed in $\{0, 1\}$, for

If all the information bits are uniformly distributed in $\{0, 1\}$, for any pair of paths $\ell, \ell' \in \{0, 1, \dots, L-1\}$,

$$W_n^{(i)}(\mathbf{y}, \hat{\mathbf{u}}^{i-1}[\ell] | \hat{u}_i[\ell]) < W_n^{(i)}(\mathbf{y}, \hat{\mathbf{u}}^{i-1}[\ell'] | \hat{u}_i[\ell'])$$

if and only if

$$\mathsf{PM}_{\ell}^{(i)} > \mathsf{PM}_{\ell'}^{(i)}$$

In view of Theorem 1, one can implement the SCL decoder using L parallel low-complexity and stable LLR-based SC decoders as the underlying building blocks and, in addition, keep track of Lpath-metrics. The metrics can be updated iteratively as the decoder proceeds according to

$$\mathsf{PM}_{\ell}^{(i)} = \mathsf{PM}_{\ell}^{(i-1)} + \ln\left(1 + e^{-(1-2\hat{u}_i[\ell])\mathsf{L}_n^{(i)}[\ell]}\right).$$
(2)

Any comparison of the likelihoods of the paths can be done equivalently using the values of the path-metrics.

Recall that the SC decoder's decisions (in line 5 of an LLRbased implementation of Algorithm 1) would be $\hat{u}_i = \delta(\mathsf{L}_n^{(i)})$ where $\delta(x) = \frac{1}{2}(1 - \operatorname{sign}(x))$. Moreover, (2) is well-approximated as

$$\mathsf{PM}_{\ell}^{(i)} \approx \begin{cases} \mathsf{PM}_{\ell}^{(i-1)} & \text{if } \hat{u}_{i}[\ell] = \delta(\mathsf{L}_{n}^{(i)}[\ell]), \\ \mathsf{PM}_{\ell}^{(i-1)} + |\mathsf{L}_{n}^{(i)}[\ell]|, & \text{otherwise.} \end{cases}$$
(3)

Hence, our metric has a natural interpretation: If at step *i*, the ℓ -th path does not follow the direction $\delta(\mathsf{L}_n^{(i)}[\ell])$, it will be penalized by an amount of $\approx |\mathsf{L}_n^{(i)}[\ell]|$ (which is the reliability of $\mathsf{L}_n^{(i)}[\ell]$).

We devote the rest of this section to prove Theorem 1.

Lemma 2. If U_i is uniformly distributed in $\{0, 1\}$, then,

$$\frac{W_n^{(i)}(\mathbf{y}, \mathbf{u}^{i-1}|u_i)}{\mathbb{P}[\mathbf{U}^i = \mathbf{u}^i | \mathbf{Y} = \mathbf{y}]} = 2\mathbb{P}[\mathbf{Y} = \mathbf{y}].$$

Proof. Since $\mathbb{P}[U_i = u_i] = \frac{1}{2}$ for $\forall u_i \in \{0, 1\}$,

$$\frac{W_n^{(i)}(\mathbf{y}, \mathbf{u}^{i-1}|u_i)}{\mathbb{P}[\mathbf{U}^i = \mathbf{u}^i | \mathbf{Y} = \mathbf{y}]} = \frac{\mathbb{P}[\mathbf{Y} = \mathbf{y}, \mathbf{U}^i = \mathbf{u}^i]}{\mathbb{P}[U_i = u_i]\mathbb{P}[\mathbf{U}^i = \mathbf{u}^i | \mathbf{Y} = \mathbf{y}]}$$
$$= \frac{\mathbb{P}[\mathbf{Y} = \mathbf{y}]\mathbb{P}[\mathbf{U}^i = \mathbf{u}^i | \mathbf{Y} = \mathbf{y}]}{\mathbb{P}[U_i = u_i]\mathbb{P}[\mathbf{U}^i = \mathbf{u}^i | \mathbf{Y} = \mathbf{y}]} = 2\mathbb{P}[\mathbf{Y} = \mathbf{y}].$$

Proof of Theorem 1. We show that

$$\mathsf{PM}_{\ell}^{(i)} = -\ln\left(\mathbb{P}[\mathbf{U}^{i} = \hat{\mathbf{u}}^{i}[\ell]|\mathbf{Y} = \mathbf{y}]\right)$$
(4)

Having shown (4), Theorem 1 will follow as an immediate corollary to Lemma 2 (since the channel output **y** is fixed for all decoding paths). Since the path index ℓ is fixed on both sides of (1) we will drop it in the sequel. Let $\mu(u) \triangleq 1 - 2u$ and $\Lambda_n^{(i)} \triangleq \frac{W_n^{(i)}(\mathbf{y}, \hat{\mathbf{u}}^{i-1}|0)}{W_n^{(i)}(\mathbf{y}, \mathbf{u}^{i-1}|1)} = \frac{\mathbb{P}[\mathbf{Y}=\mathbf{y}, \mathbf{U}^{i-1}=\hat{\mathbf{u}}^{i-1}, U_i=0]}{\mathbb{P}[\mathbf{Y}=\mathbf{y}, \mathbf{U}^{i-1}=\hat{\mathbf{u}}^{i-1}, U_i=1]}$ (the last equality follows since $\mathbb{P}[U_i = 0] = \mathbb{P}[U_i = 1]$), and observe that showing (4) is equivalent to proving

$$\mathbb{P}[\mathbf{U}^{i} = \hat{\mathbf{u}}^{i} | \mathbf{Y} = \mathbf{y}] = \prod_{j=0}^{i} \left(1 + (\Lambda_{n}^{(j)})^{-\mu(\hat{u}_{j})}\right)^{-1}.$$
 (5)

Now we have

$$\mathbb{P}[\mathbf{Y} = \mathbf{y}, \mathbf{U}^{i-1} = \hat{\mathbf{u}}^{i-1}] = \sum_{\hat{u}_i \in \{0,1\}} \mathbb{P}[\mathbf{Y} = \mathbf{y}, \mathbf{U}^i = \hat{\mathbf{u}}^i]$$
$$= \mathbb{P}[\mathbf{Y} = \mathbf{y}, \mathbf{U}^i = \hat{\mathbf{u}}^i] (1 + (\Lambda_n^{(i)})^{-\mu(\hat{u}_i)})$$

Therefore,

$$\mathbb{P}[\mathbf{Y} = \mathbf{y}, \mathbf{U}^{i} = \hat{\mathbf{u}}^{i}]$$
$$= \left(1 + (\Lambda_{n}^{(i)})^{-\mu(\hat{u}_{i})}\right)^{-1} \mathbb{P}[\mathbf{Y} = \mathbf{y}, \mathbf{U}^{i-1} = \hat{\mathbf{u}}^{i-1}]. \quad (6)$$

Repeated application of (6) (for i - 1, i - 2, ..., 0) yields

$$\mathbb{P}[\mathbf{Y} = \mathbf{y}, \mathbf{U}^{i} = \hat{\mathbf{u}}^{i}] = \prod_{j=0}^{i} \left(1 + (\Lambda_{n}^{(j)})^{-\mu(\hat{u}_{i})}\right)^{-1} \mathbb{P}[\mathbf{Y} = \mathbf{y}].$$

Dividing both sides by $\mathbb{P}[\mathbf{Y} = \mathbf{y}]$ proves (5).

3. SCL DECODER HARDWARE ARCHITECTURE

In the SCL decoder (SCLD) hardware architecture of [6] the SC decoder computations are implemented using pairs of log-likelihoods (LLs), $\ln W_n^{(i)}(\mathbf{y}, \mathbf{u}^{i-1}|u), u \in \{0, 1\}$. LLs provide some numerical stability and reduce the dynamic range of the involved quantities so that a fixed-point implementation is feasible, although a large number of quantization bits is still required for good performance. In this section, we present area and timing results for an LLR-based implementation of [6] in order to highlight the significant area savings and throughput gains that can be achieved by exploiting Theorem 1.

3.1. LL-based SCLD Hardware Architecture

The LL-based SCLD hardware architecture presented in [6] is mainly comprised of three components, namely the *metric computation* unit (MCU), the *path selection* component, and the *state memories* component. The MCU consists of L arrays of processing elements (PEs), which implement L parallel LL-based SC decoders. The state memories store the LLs, the paths $\hat{\mathbf{u}}[\ell]$, as well as the *partial sums* required by the SC decoders. The path selection component is responsible for sorting the 2L LL values, $\{\ln W_n^{(i)}(\mathbf{y}, \hat{\mathbf{u}}^{i-1}[\ell]|u_i) : \forall \ell \in \{0, 1, \ldots, L-1\}, \forall u_i \in \{0, 1\}\}$ and choosing the L most likely paths to follow in each step.

For a low-complexity hardware implementation, the LLs used by the SC decoders have to be quantized. Specifically, in [6] the channel LLs are quantized using an unsigned fixed-point representation with Q_i^{LL} integer and Q_f^{LL} fractional bits. The LL-based SC update rules involve additions of LLs [6], which all have the same sign. Thus, in order to prevent catastrophic overflows, at each SC decoding stage the number of integer bits is increased by one. Therefore, the LLs at each stage $s = 0, \ldots, n$, must be represented using $Q_i^{\text{LL}} + s$ integer bits. This necessity leads to large and very irregular LL memories, which are not well-suited for hardware implementation. Moreover, the PEs and the metric sorter need to support computations with the maximum bit-width, i.e., $Q^{\text{LL}} = Q_i^{\text{LL}} + Q_f^{\text{LL}} + n$ bits. It has been shown in [6] that in total

$$B_{\rm tot}^{\rm LL} = (2L+2)NQ^{\rm LL} + 2L(2N - n - Q^{\rm LL} - 2)$$
(7)

bits are required for storing the LLs in this scheme.

3.2. LLR-based SCLD Hardware Architecture

For the LLR-based implementation, the MCUs are modified to implement the Min-Sum LLR-based SC update rules [3, Section VI]. Moreover, the path selection unit implements the approximated pathmetric update-rule of (3). Since the rule is iterative, the path selection unit contains a memory with L storage locations to store $\mathsf{PM}_{\ell}^{(i)}, \ell \in \{0, 1, \dots, L-1\}$. The state memory component is modified to support the following quantization scheme of the LLRs.

LLRs are quantized using a signed fixed-point representation with Q_i^{LLR} integer and Q_f^{LLR} fractional bits. The total number of bits per LLR is $Q^{\text{LLR}} = Q_i^{\text{LLR}} + Q_f^{\text{LLR}} + 1$. Since the LLRs are signed quantities and the Min-Sum update rules for SC decoding



Fig. 1. The performance of floating point LLR-based vs. fixed-point LLR-based and fixed-point LL-based SCL decoders (and that of SC decoder for comparison).

involve both additions and subtractions, the dynamic range of the LLRs used in LLR-based SCLD is generally smaller than that of the LLs used in the LL-based SCLD. Thus, one intuitively expects overflows to happen less frequently and that there is no need to increase the word size by one bit per decoding stage. This intuition is confirmed by our simulation results. This leads to an LLR memory with a fixed word size, which is more suitable for hardware implementation than the irregular memory required by the LL-based decoder. We can guarantee that there will be no overflows in the path metrics by using an unsigned representation with $Q_L^{\rm LLR} + n$ integer and $Q_f^{\rm LLR}$ fractional bits yielding totally $Q^{\rm M} = Q_L^{\rm LLR} + Q_f^{\rm LLR} + n$ bits per path-metric. It turns out that, in practice, much fewer bits are sufficient. Using an approach identical to that of [6, Section IV.C] one can verify that totally $(N + (N - 1)L)Q^{\rm LLR}$ bits will be needed to store the LLRs. Adding the $LQ^{\rm M}$ bits for storing the path-metrics, we see that this implementation will require

$$B_{\rm tot}^{\rm LLR} = (N + (N - 1)L)Q^{\rm LLR} + LQ^{\rm M}$$
(8)

bits for the storage of LLRs and path metrics in total.

3.3. Implementation Results

In Fig. 1 the frame error rate (FER) of an LLR-based floating-point implementation of an SCL decoder with exact SC decoding and metric update rules is compared to that of a fixed-point implementation of an SCL decoder with approximated SC and metric update rules for an N = 1024 polar code of rate 1/2 over a BAWGN channel².

The fixed-point LLR-based decoder uses $Q_i^{\text{LLR}} = 5$ and $Q_f^{\text{LLR}} = 0$ for quantizing the LLRs. Although 15 bit for quantizing each path-metric $\mathsf{PM}_{\ell}^{(i)}$ are required to guarantee no overflows,

²The code is optimized for $E_b/N_0 = 2$ dB and constructed using the Monte-Carlo method of [1, Section IX].

L	$B_{ m tot}^{ m LL}\left(7 ight)$	$B_{ m tot}^{ m LLR}$ (8)	Reduction
2	32,736 bits	22,532 bits	31.2~%
4	57,280 bits	38,920 bits	32.1~%
8	$106, 368 {\rm bits}$	71,696 bits	32.6~%
16	204,544 bits	137,248 bits	32.9~%

Table 1. Memory requirement for LL- and LLR-based decoder for $N = 1024, Q^{LL} = 4, Q^{LLR} = 6$ and $Q^{M} = 8$.

Cell Area	LLR-based	LL-based	Reduction	
List Size	L = 1 ([3, 7]	L = 1 ([3, Table IV], scaled to 90 nm)		
Total	$0.592 \ { m mm}^2$	n/a	n/a	
Memory	$0.554 \mathrm{~mm^2}$	n/a	n/a	
MCU	$0.034~\mathrm{mm^2}$	n/a	n/a	
Other	0.004 mm^2	n/a	n/a	
List Size		L = 2		
Total	$0.977~\mathrm{mm^2}$	$1.668 \mathrm{~mm^2}$	41.4%	
Memory	$0.749 \mathrm{~mm^2}$	1.126 mm^2	33.4%	
MCU	$0.205 \mathrm{~mm^2}$	$0.504 \; { m mm^2}$	59.3%	
Path Selection	$0.001 \mathrm{~mm^2}$	$0.002 \mathrm{~mm^2}$	50.0%	
Other	0.022 mm^2	0.036 mm^2	38.8%	
List Size		L = 4		
Total	$1.743 \mathrm{~mm^2}$	$3.708 \; { m mm}^2$	53.0%	
Memory	$1.303 \mathrm{~mm^2}$	2.348 mm^2	44.5%	
MCU	$0.348 \mathrm{~mm^2}$	$0.984 \mathrm{~mm^2}$	64.6%	
Path Selection	$0.019~\mathrm{mm^2}$	$0.030 \mathrm{~mm^2}$	36.7%	
Other	$0.073~\mathrm{mm^2}$	0.346 mm^2	78.9%	

Table 2. Cell area breakdown of the LL- and LLR-based SCL decoders for an N = 1024 polar code and that of the SC decoder implementation of [3] for comparison.

 $Q^{\rm M} = 8$ bits are sufficient; simulated performance for $Q^{\rm M} = 8$ and $Q^{\rm M} = 15$ are the same (while setting $Q^{\rm M} = 7$ degrades the performance). We observe that, using the aforementioned parameters, the performance loss of the fixed-point implementation with approximated updated rules is minimal with respect to that of the floating-point implementation with exact update rules. Moreover, the FER of this fixed-point implementation can be seen to be exactly equal to that of an LL-based fixed-point implementation of SCL decoding in [6] using $Q_i^{\rm LL} = 4$ and $Q_f^{\rm LL} = 0$. Thus, the following comparison between the two decoders is fair in terms of the FER.

Using these parameters, the PEs in the LL-based decoder have to support bit-widths of up to $Q_i^{\text{LL}} + Q_f^{\text{LL}} + 10 = 14$ bits, whereas for the LLR-based decoder supporting $Q_i^{\text{LLR}} + Q_f^{\text{LLR}} + 1 = 6$ bits suffice. Moreover, the sorting metric in the former implementation is 14 bits wide, while in the latter it is only $Q^{\text{M}} = 8$ bits wide. We have compared the memory requirement of two implementations (Equations (7) and (8)) as a function of the list size L in Table 1. We observe that the LLR-based representation is advantageous in terms of the memory requirements, in particular as the list-size increases³.

In Table 2 we compare synthesis area results for LL- and LLRbased SCL decoders with L = 2 and L = 4. All designs were synthesized using the same UMC 90nm library in the typical corner. We observe that by using LLRs the total area is reduced by 41%and 53% for an SCL decoder with L = 2 and L = 4, respectively. In absolute numbers, the largest gain in both cases comes from the reduced size of the state memories, which are the largest components

	LLR-based	LL-based	Speedup
List Size	L = 1 ([3, Table IV])		
Clock Frequency	$500 \mathrm{~MHz}$	n/a	n/a
Coded Throughput	$246 \mathrm{~Mbps}$	n/a	n/a
List Size		L = 2	
Clock Frequency	$558 \mathrm{~MHz}$	$427 \mathrm{~MHz}$	30.7%
Coded Throughput	$220 \mathrm{~Mbps}$	$168 \mathrm{~Mbps}$	30.7%
List Size		L = 4	
Clock Frequency	$412 \mathrm{~MHz}$	$386 \mathrm{~MHz}$	6.7%
Coded Throughput	$162 \mathrm{~Mbps}$	$152 \mathrm{~Mbps}$	6.7%

Table 3. Operating frequency results for the LL- and LLR-based SCL decoders for an N = 1024 polar code and that of the SC decoder implementation of [3] for comparison.

of both the LL- and the LLR-based decoders. In relative numbers the transition to LLRs is most beneficial for the MCU, which is reduced in size by approximately 60% in both cases. This number is in line with the approximately 60% reduction in bit-width of the quantities involved in the computations (i.e., from 14 bits to 6 bits). The path selection component benefits from the 40% bit-width reduction of the path metrics with an average area reduction of the same order.

The corresponding post-synthesis timing results are presented in Table 3. We observe that for L = 2, the LLR-based decoder can achieve a 31% higher clock frequency than the corresponding LLbased decoder. This significant improvement comes not only from the bit-width reduction of the MCUs, but also from the highly reduces size of the LLR storage memory. In the L = 4 LLR-based decoder the signal path with the highest delay in the hardware goes through the metric sorter contained in the path selection component and not through the MCU. The 40% bit-width reduction of the metrics helps the comparators used in the radix-2L sorter [6], but it does not improve the logic gate tree that follows these comparators and combines their results. Thus, the increase in clock frequency is 7%.

Remark. The transition to LLRs can significantly reduce the area and increase the operating frequency of an implementation of SCL decoding. To be fair, we mention one disadvantage of the iterative metric update, namely that the simplified SC decoding proposed in [8] can no longer be taken full advantage of, since the LLRs for all bits are required to keep the metric updated. It can still be applied to the first group of consecutive frozen bits, which is usually large.

4. CONCLUSION

In this paper, we derived an LLR-based implementation of the successive cancellation list decoder using a path metric based on which the paths can be ranked according to their likelihoods. This metric can also be used for any other tree-search decoding algorithm that compares the paths according to their likelihoods, such as SC stack decoding [9]. Moreover, we demonstrated the advantages of our implementation by comparing synthesis results for an LL- and an LLR-based hardware successive cancellation list decoder architecture. Specifically, the decoder area was reduced by up to 53% and the clock frequency was increased by up to 31%.

In addition to the gains in the hardware cost, most processing blocks, such as channel equalizers and demodulators, in practical receivers process data in the form of LLRs. Hence, the presented implementation of decoder can readily be integrated in existing systems while the LL-based decoder would require an extra preprocessing stage to convert the channel LLRs to LLs.

³Our simulation results confirm that the same quantization scheme leads to a fair comparison for L = 8 and L = 16 as well.

5. REFERENCES

- E. Arıkan, "Channel polarization: A method for constructing capacity-achieving codes for symmetric binary-input memoryless channels," *IEEE Transactions on Information Theory*, vol. 55, no. 7, pp. 3051–3073, July 2009.
- [2] C. Leroux, I. Tal, A. Vardy, and W. Gross, "Hardware architectures for successive cancellation decoding of polar codes," in 2011 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), 2011, pp. 1665–1668.
- [3] C. Leroux, A. Raymond, G. Sarkis, and W. Gross, "A semiparallel successive-cancellation decoder for polar codes," *IEEE Transactions on Signal Processing*, vol. 61, no. 2, pp. 289–299, January 2013.
- [4] C. Zhang and K. Parhi, "Low-latency sequential and overlapped architectures for successive cancellation polar decoder," *IEEE Transactions on Signal Processing*, vol. 61, no. 10, pp. 2429–2441, March 2013.

- [5] I. Tal and A. Vardy, "List decoding of polar codes," in *Proceedings of IEEE International Symposium on Information Theory* (ISIT), 2011, 2011, pp. 1–5.
- [6] A. Balatsoukas-Stimming, A. Raymond, W. Gross, and A. Burg, "Tree search architecture for list successive cancellation decoding of polar codes," *IEEE Transactions on Circuits* and Systems II: Express Briefs (submitted), 2013. [Online]. Available: http://arxiv.org/abs/1303.7127
- [7] E. Arıkan and E. Telatar, "On the rate of channel polarization," in *Proceedings of IEEE International Symposium on Information Theory (ISIT)*, 2009, Jul. 2009, pp. 1493–1495.
- [8] A. Alamdar-Yazdi and F. Kschischang, "A simplified successive-cancellation decoder for polar codes," *IEEE Communications Letters*, vol. 15, no. 12, pp. 1378–1380, October 2011.
- [9] K. Niu and K. Chen, "Stack decoding of polar codes," *Electron*ics Letters, vol. 48, no. 12, pp. 695–697, June 2012.