

Reliable Linear, Sesquilinear and Bijective Operations On Integer Data Streams Via Numerical Entanglement

Mohammad Ashraful Anam and Yiannis Andreopoulos, *Senior Member, IEEE*

Abstract—A new technique is proposed for fault-tolerant linear, sesquilinear and bijective (LSB) operations on M integer data streams ($M \geq 3$), such as: scaling, additions/subtractions, inner or outer vector products, permutations and convolutions. In the proposed method, the M input integer data streams are linearly superimposed to form M numerically-entangled integer data streams that are stored in-place of the original inputs. A series of LSB operations can then be performed directly using these entangled data streams. The results are extracted from the M entangled output streams by additions and arithmetic shifts. Any soft errors affecting any single disentangled output stream are guaranteed to be detectable via a specific post-computation reliability check. In addition, when utilizing a separate processor core for each of the M streams, the proposed approach can recover all outputs after any single fail-stop failure. Importantly, unlike algorithm-based fault tolerance (ABFT) methods, the number of operations required for the entanglement, extraction and validation of the results is linearly related to the number of the inputs and does not depend on the complexity of the performed LSB operations. We have validated our proposal in an Intel processor (Haswell architecture with AVX2 support) via fast Fourier transforms, circular convolutions, and matrix multiplication operations. Our analysis and experiments reveal that the proposed approach incurs between 0.03% to 7% reduction in processing throughput for a wide variety of LSB operations. This overhead is 5 to 1000 times smaller than that of the equivalent ABFT method that uses a checksum stream. Thus, our proposal can be used in fault-generating processor hardware or safety-critical applications, where high reliability is required without the cost of ABFT or modular redundancy.

Index Terms—linear operations, sum-of-products, algorithm-based fault tolerance, silent data corruption, core failures, numerical entanglement

I. INTRODUCTION

THE current technology roadmap for high-performance computing (HPC) indicates that digital signal processing (DSP) routines running on such systems must become resilient to transient or permanent faults occurring in arithmetic, memory or logic units. Such faults can be caused by process variations, silent data corruptions (e.g., due to particle strikes, circuit overclocking or undervolting, or other

hardware non-idealities) [1], scheduling and runtime-induced faults [2], misconfigured application programming interfaces (APIs) [3] and opportunistic resource reservation [4] in cloud computing systems. For example, DSP systems such as: webpage or multimedia retrieval [5], object or face recognition in images [6], machine learning and security applications [7], transform decompositions and video encoding [8]–[16], and visual search and retrieval systems [17], are now deployed using Amazon Elastic Compute Cloud (EC2) spot instances with substantially-reduced billing cost (e.g., in the order of 0.01\$ per core per hour). However, Amazon reserves the right to terminate EC2 spot instances at any moment with little or no prior notice. In addition, transient service interruptions may occur at unpredictable intervals, since processor cores in EC2 spot instance reservations may not be solely dedicated to the cluster under consideration [3], [4].

System-induced faults in DSP routines manifest as [1], [2], [18]: (i) *transient faults*, where execution continues uninterrupted on all input data streams—albeit with corrupted data and possibly carrying out erroneous logic or arithmetic operations—or (ii) *fail-stop failures*, where the execution on one of the processor cores halts due to a fail-stop exception (e.g., overflow detection, memory leak assertion, etc.) or a system crash. In the first case, a highly-reliable system should be able to detect all faults (and possibly correct them); in the second case (which is analogous to an *erasure* in a communications system, since it is easily detectable), the system should be able to recover the results of the halted execution without requiring recomputation.

The compute- and memory-intensive parts within DSP routines comprise: transform decompositions, signal cross-correlation, inner or outer vector products, matrix products, etc. Such operations are linear, sesquilinear (also known as “one-and-half linear”) and bijective, collectively called LSB operations in this paper. Moreover, especially when considering multimedia data, these operations are typically performed using 32-bit or 64-bit integer arithmetic. Therefore, ensuring *highly reliable* integer LSB operations with minimal overhead against their *unreliable* equivalents is of paramount importance for signal processing systems.

A. Summary of Prior Work

Existing techniques that can ensure reliability to transient faults and/or fail-stop failures comprise three categories: (i) component-based reliability via error-correcting

Copyright (c) 2015 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org. This work was supported by EPSRC, project EP/M00113X/1. The authors are with the Electronic and Electrical Engineering Department, University College London, Roberts Building, Torrington Place, London, WC1E 7JE, Tel. +44 20 7679 7303, Fax. +44 20 7388 9325 (both authors), Email: {mohammad.anam.10, i.andreopoulos}@ucl.ac.uk. Parts of this paper have been presented at the 2015 IEEE International On-Line Testing Symposium (IEEE IOLTS 2015).

codes (ECC) [19], e.g., random access memory or cache memory chip designs with native ECC support for the detection/correction of bit flips; (ii) algorithm-based fault-tolerance (ABFT) [20]–[26], i.e., methods producing additional checksum inputs/outputs that are tailored to the algorithm under consideration; (iii) systems with double or triple modular redundancy (MR), where the same operation is performed in parallel in two or three separate processors (or threads) that cross-validate their results and recover from fail-stop failures [27].

While component-based methods can indeed mitigate many of the faults occurring in memory or processor components, their integration into hardware designs is known to incur substantial overhead [28]. Therefore, system designers tend to disable such functionalities for most applications, even at the risk of allowing faults to occur undetected. Moreover, in order to offer end-to-end guarantees on fault tolerance and robustness to fail-stop failures (i.e., from the input data streams to the results of a certain computation), component-based fault tolerance must be integrated with checkpointing and execution roll-back [21], which further increases the complexity of the final solution.

Similarly, it is well known that ABFT and MR solutions can lead to substantial processing overhead in hardware/software systems and can incur increased energy consumption. For example, Bosilca *et al.* [29], report that checkpointing the system state to detect a single fault per process (and rolling back to a previous state when faults are detected) leads to 9%–34% time overhead for an implementation using up to 484 processes. Similarly, Chen and Dongarra [21] report that, for detecting a single fault per subblock of a large matrix operation, 4%–9% execution time overhead is incurred in a ScaLAPACK implementation over a distributed computing system. Finally, Wunderlich *et al.* [30] report that ABFT for the generic matrix multiply (GEMM) routine incurs 18%–45% execution time overhead versus the unprotected GEMM on medium to large matrix dimensions under a GPU implementation. In conjunction with recent studies on soft errors in processors that indicate that hardware faults tend to happen in bursts [1], [19], [31], this shows that ABFT techniques may ultimately not be the best way to mitigate *arbitrary fault patterns* occurring in 32-bit or 64-bit data representations in memory, arithmetic or logic units of the utilized hardware. On the other hand, while MR approaches can indeed mitigate such faults with very high probability, it is well known that they incur a two-fold or three-fold penalty in execution time (or energy consumption) as well as substantial data transfers and latencies to synchronize and cross-check results [27].

B. Contribution

We propose a new method to mitigate transient faults or failures in LSB operations performed in integer data streams with integer arithmetic units. Examples of such operations are element-by-element additions and multiplications, inner and outer vector products, sum-of-squares and permutation operations. They are the building blocks of algorithms of foundational importance, such as: matrix multiplication [20],

[32], convolution/cross-correlation [33], template matching for search and motion estimation algorithms [13]–[16], [34]–[36], covariance calculations [6], [37], integer-to-integer transforms [9], [10], [12], [38] and permutation-based encoding systems [39], which form the core of the applications discussed earlier. Our method:

- 1) does not generate additional data in form of checksum or duplicate inputs, as done by ABFT-based or MR-based methods; instead, it performs pairwise linear superpositions within the numerical representation of the original inputs, thereby increasing their dynamic range, albeit in a controllable manner.
- 2) does not require modifications to the arithmetic or memory units, as done by component-based ECC approaches, and can be deployed in standard 32/64-bit integer units or even 32/64-bit floating-point units;
- 3) does not depend on the specifics of the LSB operation performed; in fact, it can also be used to detect silent data corruptions in storage systems, i.e., when no computation is performed with the data.

Beyond the analytic presentation of our proposal and the theoretical estimation of its complexity against ABFT, we also present performance results using fast Fourier transform (FFT) computation, cross-correlation and matrix product operations, thereby significantly advancing our early exposition [40] that summarized our initial findings on fail-stop failure mitigation for limited types of operations. The results show that, for the vast majority of cases, our method’s percentile overhead in execution against the fault-intolerant (i.e., conventional) realization is upper-bounded by 0.6%. This overhead is found to be 5 to 1000 times smaller than the one incurred by the ABFT approach that offers the same fault tolerance capability.

C. Paper Organization

In Section II, we outline ABFT, MR and the basic concept of the proposed approach for fault tolerance in numerical stream processing. In Section III we present the details for the newly-proposed concept of numerical entanglement and demonstrate its inherent reliability for LSB processing of integer streams. Section IV presents the complexity of numerical entanglements within integer linear or sesquilinear operations. Section V presents experimental comparisons and Section VI presents some concluding remarks.

II. ABFT/MR METHODS VERSUS NUMERICAL ENTANGLEMENT

Consider a series of M input streams of integers, each comprising N_{in} samples¹ ($M \geq 3$, $N_{\text{in}} \in \mathbb{N}^*$):

$$\mathbf{c}_m = [c_{m,0} \ \dots \ c_{m,N_{\text{in}}-1}], \ 0 \leq m < M. \quad (1)$$

These may be the elements of M rows of a matrix of integers, or a set of M input integer streams of data to be operated upon with an integer kernel \mathbf{g} . This operation is performed by:

$$\forall m : \mathbf{d}_m = \mathbf{c}_m \text{ op } \mathbf{g} \\ \text{op} \in \left\{ +, -, \times, \langle \cdot, \cdot \rangle, \otimes, \left(\begin{array}{c} \mathcal{J} \\ \mathcal{G} \end{array} \right), * \right\} \quad (2)$$

with \mathbf{d}_m the m th vector of output results and op any LSB operator such as element-by-element addition/subtraction/multiplication, inner/outer product, permutation² (i.e., bijective mapping from the sequential index set \mathcal{J} to index set \mathcal{G} corresponding to \mathbf{g}) and circular convolution or cross-correlation with \mathbf{g} . An illustration of the application of (2) is given in Fig. 1(a). Beyond the single LSB operator indicated in (2) and illustrated in Fig. 1(a), we can also assume *series* of such operators applied consecutively in order to realize higher-level algorithmic processing, e.g., multiple consecutive additions, subtractions and scaling operations with pre-established kernels followed by circular convolutions and permutation operations. Conversely, the input data streams can also be left in their native state (i.e., stored in memory), if $\text{op} = \{\times\}$ and $\mathbf{g} = 1$.

A. Algorithm-based Fault Tolerance and Modular Redundancy

In their original (or “pure”) form, the input data streams of (1) are uncorrelated and one input element cannot be used to cross-check for faults in another without inserting some form of coding or redundancy. This is conventionally achieved via ABFT methods [20]–[26]. Specifically, one *additional* input stream is created that comprises *checksums* of the original inputs:

$$\mathbf{r} = [r_0 \ \dots \ r_{N_{\text{in}}-1}], \quad (3)$$

by using, for example, the sum of the input samples [25], [26] at the n th position in each of the M streams, $0 \leq n < N_{\text{in}}$:

¹Notations: Boldface uppercase and lowercase letters indicate matrices and vectors, respectively; the corresponding italicized lowercase indicate their individual elements, e.g. \mathbf{A} and $a_{m,n}$; calligraphic uppercase letters indicate operators; \mathbb{N}^* is the set of natural numbers excluding zero; \hat{d} denotes the recovered value of d after unpacking or disentanglement; all indices are integers. Basic operators: $\lfloor a \rfloor$ is the largest integer that is smaller or equal to a (floor operation); $\lceil a \rceil$ is the smallest integer that is larger or equal to a (ceiling operation); $\|\mathbf{a}\|$ is vector norm-2; $a \& b$ denotes binary AND operation between the bits of a and b , respectively; $a \ll b$ and $a \gg b$ indicate left and right arithmetic shift of integer a by b bits with truncation occurring at the most-significant or least significant bit, respectively; $a \bmod b = a - \lfloor \frac{a}{b} \rfloor b$ is the modulo operation; $a \leftarrow b$ assigns the value of variable or expression b to variable a .

²We remark that we consider LSB operations that are *not* data-dependent, e.g., permutations according to fixed index sets as in the Burrows-Wheeler transform [39], [41].

$$\forall n : r_n = \sum_{m=0}^{M-1} c_{m,n}. \quad (4)$$

Then the processing is performed in all input streams $\mathbf{c}_0, \dots, \mathbf{c}_{M-1}$ and in the checksum input stream \mathbf{r} by:

$$\begin{bmatrix} \mathbf{d}_0 \\ \vdots \\ \mathbf{d}_{M-1} \\ \mathbf{e} \end{bmatrix} = \begin{bmatrix} \mathbf{c}_0 \\ \vdots \\ \mathbf{c}_{M-1} \\ \mathbf{r} \end{bmatrix} \text{ op } \mathbf{g}, \quad (5)$$

Any transient faults in any single stream out of $M + 1$ output streams can then be detected by checking if:

$$\exists n : \sum_{m=0}^{M-1} d_{m,n} \neq e_n. \quad (6)$$

This process is pictorially illustrated in Fig. 1(b). Similarly, result recovery after any single fail-stop failure can take place by subtracting the results of all the remaining output streams from \mathbf{r} . As discussed in partitioning schemes for checksum-based methods and ABFT [25], [26], the recovery capability can be increased by using additional weighted checksums. However, this comes at the cost of increasing the number of checksum input streams, which leads to increased overhead. For this reason, practical approaches tend to use a single checksum stream [20], [21], [25], [29], [42]. At the other extreme, when M checksum streams are used, this corresponds to repeating the operation twice (dual modular redundancy) and any fault on the original computation can be detected if the results are compared with the results of the checksum set. In summary, the practical limitations of ABFT are:

- 1) The percentile implementation overhead (i.e., processing cycles, energy consumption, memory accesses) of ABFT is $\frac{1}{M} \times 100\%$.
- 2) The dynamic range of the computations with each of the checksum input streams is increased by $\lceil \log_2 M \rceil$ bits, as each of the checksum input data values comprises the sum of groups of M input samples, as shown in (4).
- 3) The overall execution flow changes as the total number of processed streams is changed from M to $M + 1$.

B. Numerical Entanglement

In our proposal, numerical entanglement mixes the inputs prior to linear processing using linear superposition and ensures the results can be extracted and validated via a mixture of shift-add operations. It is conceptualized in Fig. 2. As shown there, M ($M \geq 3$) input streams (each comprising N_{in} integer samples and denoted by \mathbf{c}_m , $0 \leq m < M$) become M entangled streams of integers (of N_{in} integer samples each), $\boldsymbol{\epsilon}_m$. Each element of the m th entangled stream, $\epsilon_{m,n}$ ($0 \leq n < N_{\text{in}}$), comprises the superposition of two input elements $c_{x,n}$ and $c_{y,n}$ from different input streams x and y , i.e., $0 \leq x, y < M$ and $x \neq y$. The LSB operation is carried out with the entangled streams, thereby producing the entangled output streams $\boldsymbol{\delta}_m$ (each comprising N_{out} integer samples). These can be disentangled to recover the final results $\hat{\mathbf{d}}_m$. Any transient faults that occurred on any single entangled output

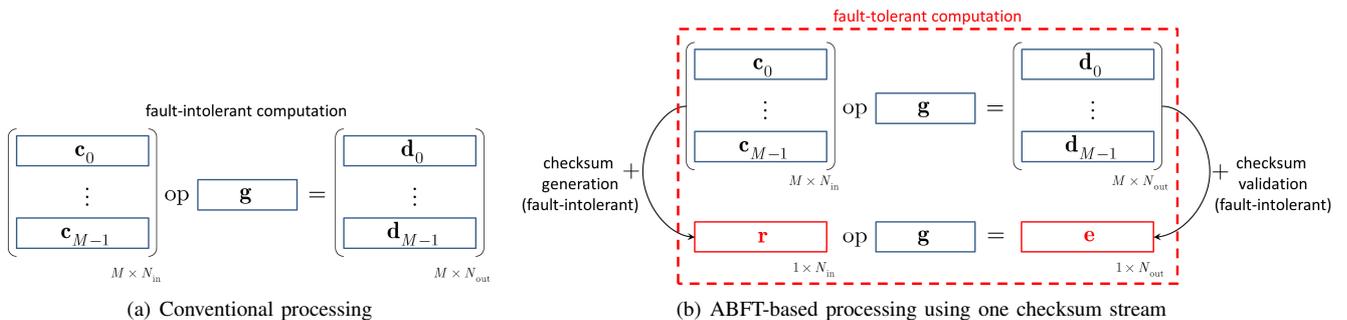


Figure 1. (a) kernel g applied to M streams of input integers via LSB operator op ; (b) corresponding application using one checksum input stream for transient fault detection or fail-stop failure recovery via ABFT.

stream out of M are detectable with a single test that utilizes additions and shift operations. In addition, any single fail-stop failure can be mitigated from the results of the remaining streams.

Unlike checksum or MR methods, numerical entanglement does not use additional checksum streams. Therefore, the complexity of entanglement, disentanglement (recovery) and fault checking does not depend on the complexity of the operator op , or on the length and type of the kernel (operand) g . The entangled inputs can be written in-place and no additional storage or additional operations are needed during the execution of the actual operation (albeit at the cost of reducing the dynamic range supported). In fact, with the exception of addition/subtraction operations with constants, no modifications are performed to the processors, software routines and arithmetic units performing the operation with kernel g and the computational system design remains unaware of the fact that entangled input streams are used instead of the original input streams. Thus, the entangled computation shown in Fig. 2 can be executed concurrently in M processing cores (that may be physically separate) and any memory optimization or other algorithmic optimization can be applied in the same manner as for the original computation. For example, if a fast Fourier transform (FFT) routine is used for the calculation of convolution or cross-correlation of each input stream c_m with kernel g , this routine can be used directly with the entangled input streams ϵ_m and kernel g . A summary of the features of each approach is presented in Table I. The comparison includes our previous work on numerical packing and duplicate execution [33], [36], [43]–[45], which can be seen as an alternative form of MR.

III. NUMERICAL ENTANGLEMENT

Numerical entanglement bears some resemblance to the concept of numerical packing proposed previously by Andreopoulos *et al.* [33], [35], [36], [43]–[45], Kadyrov and Petrou [34] and others [38], [46]. By using multiple packed representations and no overlap, it can be shown [45] that all faults occurring on a single description can be detected, at the cost of using a 64-bit integer representation that can accommodate up to 19-bit signed integer outputs in packed format. This illustrates that utilizing packing for fault detection may be a viable approach, but it comes at the cost of significant

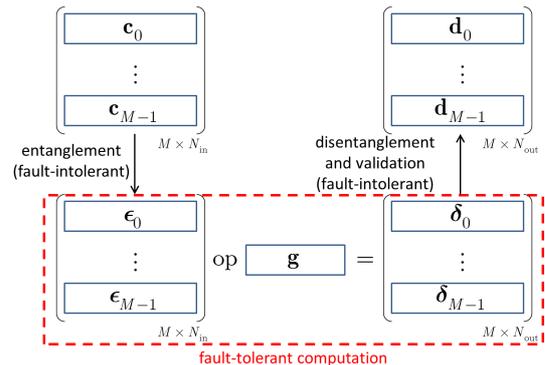


Figure 2. LSB processing via numerical entanglement, followed by result recovery and validation.

reduction in the dynamic range supported by the packed representation. The proposed approach overcomes this limitation by allowing inputs to superimpose each other, thereby creating an input representation where certain bits from pairs of input data samples are numerically entangled, as explained in the following.

A. Proposed Numerical Entanglement in Groups of Three Inputs ($M = 3$)

Numerical entanglement guarantees the detection of *any* transient fault occurring in any one out of the M entangled streams created. In addition, if a separate core is used for the processing of each stream, it allows for the recovery of all outputs after any single fail-stop core failure.

1) *Entanglement*: In the simplest form of entanglement ($M = 3$), each triplet of input samples [shown in Fig. 3(a)] of the three integer streams, $c_{0,n}$, $c_{1,n}$ and $c_{2,n}$, $0 \leq n < N_{in}$, produces the following entangled triplet via the linear superposition:

$$\begin{aligned} \epsilon_{0,n} &= \mathcal{S}_l \{c_{2,n}\} + c_{0,n} \\ \epsilon_{1,n} &= \mathcal{S}_l \{c_{0,n}\} + c_{1,n} \\ \epsilon_{2,n} &= \mathcal{S}_l \{c_{1,n}\} + c_{2,n} \end{aligned} \quad (7)$$

with

$$2l + k \leq w \quad (8)$$

Table I

SUMMARY OF FEATURES OF DIFFERENT METHODS FOR FAULT TOLERANCE WITHIN EACH GROUP OF M OUTPUTS OF M STREAMS, WITH EACH STREAM ELEMENT COMPRISING w BITS.

Method	ABFT	Dual Modular	Packing and Duplicate	Proposed
Feature	[20], [21], [29], [42]	Redundancy [27]	Execution [33], [36], [43]	Numerical Entanglement
In-place storage	No	No	No	Yes
% of redundant computations	$\frac{1}{M} \times 100\%$	100%	0%	0%
Output bandwidth reduction due to fault tolerance capability	$\lceil \log_2 M \rceil$ bits (only for checksum inputs)	0 bits	0 bits	$\lceil \frac{w}{M} \rceil$ bits
Guaranteed fault detection / fail-stop failure mitigation	1 fault in every M outputs	1 fault in every 2 outputs	1 fault in every 2 outputs	1 fault in every M outputs
Percentile overhead in execution	$\frac{1}{M} \times 100\%$	More than 100%	More than 100%	0.03% to 7% (decreases with operand length)

where $w \in \{32, 64\}$ for 32 or 64-bit integer representations, and

$$\mathcal{S}_b \{a\} \equiv \begin{cases} (a \ll b), & \text{if } b \geq 0 \\ [a \gg (-b)], & \text{if } b < 0 \end{cases} \quad (9)$$

the left or right arithmetic shift of a by b bits, with the maximum dynamic range supported for each signed input, $c_{0,n}$, $c_{1,n}$ and $c_{2,n}$, being proportional to $l+k$ bits.

The values for l and k are chosen such that $l+k$ is maximum within the constraint of (8) and $k \leq l$. Via the application of LSB operations, each ϵ_m entangled input stream ($0 \leq m < M$) is converted to the entangled output stream³ δ_m (which contains N_{out} values):

$$\forall m : \delta_m = (\epsilon_m \text{ op } \mathbf{g}). \quad (10)$$

A conceptual illustration of the entangled outputs after (7) and (10) is given in Fig. 3(b). The overlap of the arrows in the representations of $\delta_{0,n}$, $\delta_{1,n}$ and $\delta_{2,n}$ indicate the ‘‘entangled’’ region of k bits, where the final outputs, $d_{0,n}$, $d_{1,n}$ and $d_{2,n}$, numerically superimpose each other (in pairs) due to the shift and add operations of (7). Similarly as before, l bits of dynamic range are sacrificed in order to detect faults (or mitigate a fail-stop failure) during the computation of (10). Since it is assumed that the dynamic range of the inputs does not exceed $l+k$ bits, the entangled representation is contained within $2l+k$ bits and never overflows. As a practical instantiation of (7), we can set $w = 32$, $l = 11$ and $k = 10$ in a signed 32-bit integer configuration.

We now describe the disentanglement and result validation (or recovery) process. The reader can also consult Fig. 3.

2) *Disentanglement*: We can disentangle and recover the final results $\hat{d}_{0,n}$, $\hat{d}_{1,n}$ and $\hat{d}_{2,n}$ by ($0 \leq n < N_{\text{out}}$):

$$\begin{aligned} d_{\text{temp}} &= \delta_{2,n} - \mathcal{S}_l \{ \delta_{1,n} \} \\ \hat{d}_{2,n} &= \mathcal{S}_{-2(w-l)} \{ \mathcal{S}_{2(w-l)} \{ d_{\text{temp}} \} \} \\ \hat{d}_{0,n} &= \mathcal{S}_{-2l} \{ - (d_{\text{temp}} - \hat{d}_{2,n}) \} \\ \hat{d}_{1,n} &= \delta_{1,n} - \mathcal{S}_l \{ \hat{d}_{0,n} \} \end{aligned} \quad (11)$$

³For the particular cases of: $\text{op} \in \{+, -\}$, \mathbf{g} must also be entangled with itself via: $g_n \leftarrow \mathcal{S}_l \{ g_n \} + g_n$, in order to retain the homomorphism of the performed operation. All other operations occur without any modification in \mathbf{g} .

The first three parts of (11) assume a $2w$ -bit integer representation is used for the interim operations, as the temporary variable d_{temp} is stored in $2w$ -bit integer representation. However, all recovered results, $\hat{d}_{0,n}$, $\hat{d}_{1,n}$ and $\hat{d}_{2,n}$, require only $l+k$ bits.

Explanation of (11): The first part creates a temporary composite number, d_{temp} , comprising $\hat{d}_{0,n}$ in the $l+k$ most-significant bits and $\hat{d}_{2,n}$ in the $2l$ least-significant bits (therefore, d_{temp} requires $3l+k$ bits). The creation of d_{temp} is pictorially illustrated in Fig. 3(c). In the second part, $\hat{d}_{2,n}$ is extracted by: (i) discarding the $(2w-2l)$ most-significant bits of d_{temp} ; (ii) arithmetically shifting the output down to the correct range. The third part of (11) uses $\hat{d}_{2,n}$ to recover $\hat{d}_{0,n}$ from d_{temp} and, in the fourth part of (11), $\hat{d}_{0,n}$ is used to recover $\hat{d}_{1,n}$. Notice that (11) recovers all three $\hat{d}_{0,n}$, $\hat{d}_{1,n}$ and $\hat{d}_{2,n}$ without using $\delta_{0,n}$. This is a crucial aspect that leads to the fault tolerance characteristic of our proposal, which are discussed next.

B. Properties and Fault Tolerance Characteristics

Remark 1 (operations within w bits): To facilitate our exposition, the first three parts of (11) are presented under the assumption of a $2w$ -bit integer representation. However, it is straightforward to implement them via w -bit integer operations by separating d_{temp} into two parts of w bits and performing the operations separately within these parts.

Remark 2 (dynamic range): Bit $l+k$ within each recovered output $\hat{d}_{0,n}$, $\hat{d}_{1,n}$ and $\hat{d}_{2,n}$ represents its sign bit. Given that: (i) each entangled output comprises the addition of two outputs (with one of them left-shifted by l bits); (ii) the entangled outputs must not exceed $2l+k$ bits, the outputs of the LSB operations must not exceed the range

$$\forall n : d_{0,n}, d_{1,n}, d_{2,n} \in \{ - (2^{l+k-1} - 2^l), \dots, 2^{l+k-1} - 2^l \}. \quad (12)$$

Therefore, (12) comprises the range permissible for the LSB operations with the entangled representation.

Proposition 1. *If a separate core is used for each stream computation of (10) with $M = 3$, the disentanglement process of (11) can recover all results, $\hat{\mathbf{d}}_0$, $\hat{\mathbf{d}}_1$, and $\hat{\mathbf{d}}_2$, after any single fail-stop failure.*

Proof: See Appendix. ■

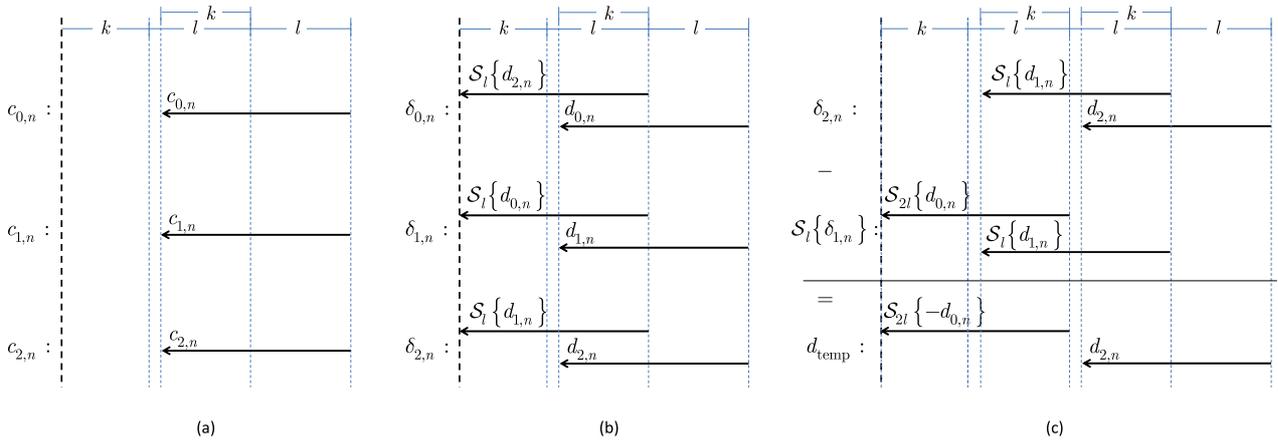


Figure 3. The arrows indicate the dynamic range of the represented data elements (from least-significant to most-significant bits). (a) Illustration of three input data samples at position n ; (b) entangled outputs after entangling and integer LSB processing; (c) illustration of the first part of (11) that produces the temporary variable d_{temp} .

The following proposition proves that, if all three entangled output streams are available, we can detect any transient fault in any single stream.

Proposition 2. *Any transient fault occurring on a single entangled output stream during the computation of (10) with $M = 3$ is detectable.*

Proof: See Appendix. \blacksquare

C. Generalized Entanglement in Groups of M Inputs ($M \geq 3$)

We extend the proposed entanglement process to using M inputs and providing M entangled descriptions, each comprising the linear superposition of two inputs. This ensures that, for every n ($0 \leq n < N_{\text{out}}$), any single transient fault will be detected within each group of M output samples. Alternatively, if M separate cores are used for the computation of the M entangled outputs, any single fail-stop failure will be mitigated from the remaining $M - 1$ output entangled streams.

The condition for ensuring that overflow is avoided is

$$(M - 1)l + k \leq w \quad (13)$$

and the dynamic range supported for all outputs is ($\forall m, n$):

$$d_{m,n} \in \{-2^{(M-3)l+k} (2^{l-1} - 1), \dots, 2^{(M-3)l+k} (2^{l-1} - 1)\}. \quad (14)$$

The values for l and k are chosen such that $(M - 2)l + k$ is maximum within the constraint of (13) and $k \leq l$.

We now define the following circulant matrix operator comprising cyclic permutations of the $1 \times M$ vector $[1 \ 0 \ \dots \ 0 \ \mathcal{S}_l]$:

$$\mathcal{E} = \begin{bmatrix} 1 & 0 & \dots & 0 & \mathcal{S}_l \\ \mathcal{S}_l & 1 & \dots & 0 & 0 \\ & & \ddots & & \\ 0 & \dots & \mathcal{S}_l & 1 & 0 \\ 0 & \dots & 0 & \mathcal{S}_l & 1 \end{bmatrix}_{M \times M}. \quad (15)$$

Operator \mathcal{E} generalizes the proposed numerical entanglement process. Specifically, following the case of $M = 3$, in the generalized entanglement in groups of M streams, two inputs are entangled together (with one of the two shifted by l bits) to create each entangled input stream of data. Any LSB operation is then performed on these entangled input streams and we shall show that any transient fault occurring during the processing of a single entangled stream can be detected within each group of M outputs. Alternatively, any single fail-stop core failure can be mitigated.

For every input stream position n , $0 \leq n < N_{\text{in}}$, the entanglement vector performing the linear superposition of pairs out of M inputs is now formed by:

$$[\epsilon_{0,n} \ \dots \ \epsilon_{M-1,n}]^T = \mathcal{E} \{ [c_{0,n} \ \dots \ c_{M-1,n}]^T \} \quad (16)$$

Fig. 4 illustrates the entangled outputs after (16) and the LSB processing of (10) under $M = 4$. In addition, Fig. 5 illustrates the general case of M outputs, produced after M inputs were entangled to create M descriptions and (16) and (10) were performed.

After the application of (10), we can disentangle every output stream element $\delta_{m,n}$, $0 \leq m < M$, $0 \leq n < N_{\text{out}}$, as follows. Let us first identify the entangled output stream δ_r (with $0 \leq r < M$) that we shall not use in the disentanglement process, either because it is unavailable due a fail-stop failure, or because we would like to recreate it if we are checking for transient faults. We produce the $2w$ -bit temporary variable⁴ d_{temp} by:

$$d_{\text{temp}} = \sum_{m=0}^{M-2} (-1)^m \mathcal{S}_{(M-2-m)l} \{ \delta_{(r+1+m) \bmod M, n} \}. \quad (17)$$

An illustration of the result of (17) for $M = 4$ and $r = 0$ is given in Fig. 6. Notice that (17) does not use δ_r . We can then

⁴Due to the usage of $2w$ bits, (17)–(19) must be separated in two parts of w bits if the entire operation has to occur via w -bit integer arithmetic. Since this is an implementation issue, we do not illustrate this separation in our exposition.

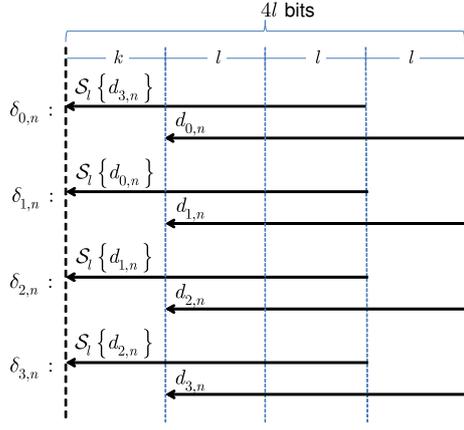


Figure 4. Illustration of entanglement of $M = 4$ outputs after integer LSB processing. In this case, for $w \in \{32, 64\}$, $k = l = \frac{w}{4}$.

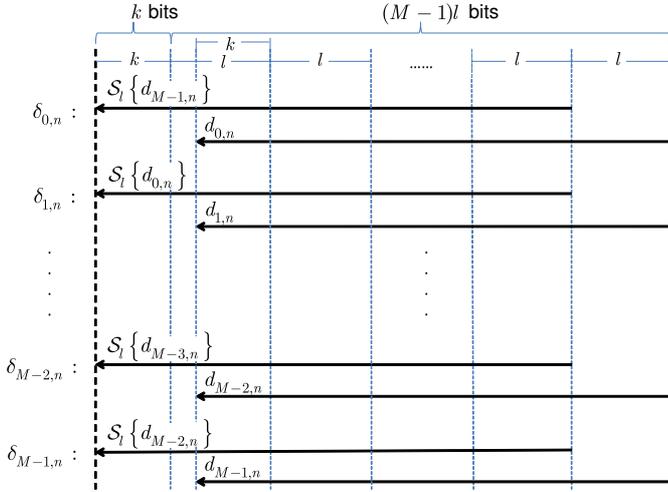


Figure 5. Illustration of the general form of entanglement of M outputs after integer LSB processing.

extract the value of $\hat{d}_{r,n}$ and $\hat{d}_{(r+M-1)\bmod M,n}$ directly from d_{temp} :

$$\hat{d}_{(r+M-1)\bmod M,n} = \mathcal{S}_{-[2w-(M-1)l]} \left\{ \mathcal{S}_{2w-(M-1)l} \left\{ d_{\text{temp}} \right\} \right\} \quad (18)$$

$$\hat{d}_{r,n} = \mathcal{S}_{-(M-1)l} \left\{ (-1)^M (d_{\text{temp}} - \hat{d}_{M-1,n}) \right\}. \quad (19)$$

The other outputs can now be disentangled by ($1 \leq m < M-2$):

$$\hat{d}_{(r+m)\bmod M,n} = \delta_{(r+m)\bmod M,n} - \mathcal{S}_l \left\{ \hat{d}_{(r+m-1)\bmod M,n} \right\}. \quad (20)$$

Proposition 3. *If a separate core is used for each stream computation of (10), the disentanglement process of (17)–(20) can recover all results, $\hat{\mathbf{d}}_0, \dots, \hat{\mathbf{d}}_{M-1}$, after any single fail-stop failure.*

Proof: See Appendix. ■

Proposition 4. *Any transient fault occurring on a single entangled stream during the computation of (10) with $M \geq 3$ is detectable.*

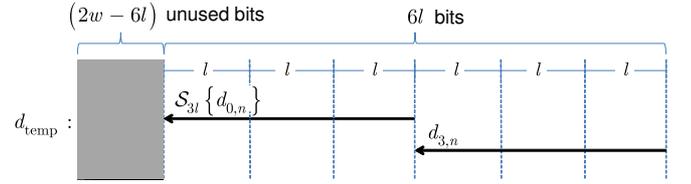


Figure 6. Temporary value d_{temp} produced during the disentanglement process of (17) within an integer representation, with $M = 4$. As shown in Table II, for $w \in \{32, 64\}$, $k = l$ for this case.

Proof: See Appendix. ■

Remark 3 (dynamic range of generalized entanglement and equivalence to checksum methods): Examples for the maximum bitwidth achievable for different cases of M are given in Table II assuming a 32-bit representation. We also present the dynamic range permitted by the equivalent checksum-based method [(3)–(6)] in order to ensure that its checksum stream does not overflow under a 32-bit representation. Evidently, for $M \leq 10$, the proposed approach incurs loss of 1 to 9 bits of dynamic range against the checksum-based method, while it allows for higher dynamic range than the checksum-based method for $M \geq 11$. At the same time, our proposal does not require the overhead of applying the LSB operations to an additional stream, as it “overlays” the information of each input onto another input via the numerical entanglement of pairs of inputs. Beyond this important difference, Propositions 3 and 4 show that our approach offers the exact equivalent to the checksum method of (3)–(6) for integer inputs.

Remark 4 (extensions): It is of theoretical interest to consider whether the proposed numerical entanglement approach can be extended to guarantee detection of multiple transient faults occurring in co-located positions in the entangled streams, or recover from more than a single fail-stop failure in M streams. In addition, while LSB operations cover a wide range of compute- and memory-intensive DSP systems, it would be interesting to investigate the applicability of our approach to data-dependent operations or non-linear operations, like modulo, binary operators, etc. Finally, beyond integer processing, the fault tolerance of numerical entanglement would be extremely beneficial to floating-point LSB operations.

Concerning the first point, a potential solution could be to create multiple tiers of entanglement, i.e., reapply the operator of (16) to $[\epsilon_{0,n} \dots \epsilon_{M-1,n}]^T$ to create K entanglement stages, and thereby examine if this leads to the possibility of resilience to K transient faults in M co-located entangled outputs (or recovery from K fail-stop failures in M streams) without requiring K checksum streams like ABFT. While we have some encouraging results in this direction, due to space limitations we plan to investigate and quantify this in another paper.

The second point (extension beyond LSB operations) is also a limitation of ABFT fault-tolerance methods. Therefore, it can be addressed in a similar way as carried out for such cases, i.e., disentangle all outputs before the non-LSB operation, carry out this operation with the disentangled outputs (using another form of fault tolerance—such as MR—for this process), and then re-entangle the results if no faults or failures are detected.

This requires a system-level approach to carefully leverage the cost of such modifications, so we opt to leave it as a topic for future work.

Finally, concerning extension to floating-point arithmetic, we do not foresee a direct way to achieve this with the proposed approach, as standard floating-point arithmetic does not allow for two “clean” top and bottom zones of bits in the way presented in Figs. 3, 4 and 5. Moreover, it is important to note that, in the case of floating-point arithmetic, even conventional ABFT approaches require more complex handling of fault detection based on appropriate thresholds, as the sum of the results of the M output streams will not be reproduced exactly by the results of the checksum stream due to the lossy and non-commutative nature of floating-point arithmetic. However, possibilities to incorporate kernel-adaptive companding and rounding (that convert floating-point inputs into integers for error-tolerant generic matrix multiplication kernels [36]) prior to entanglement, may be investigated in future work.

Table II

EXAMPLES OF l AND k VALUES AND BITWIDTH SUPPORTED FOR THE OUTPUT DATA UNDER $w = 32$ BITS AND: (i) DIFFERENT NUMBERS OF ENTANGLEMENTS; (ii) ABFT WITH $P = 1$. BOTH APPROACHES GUARANTEE THE MITIGATION OF A FAIL-STOP FAILURE (OR DETECT TRANSIENT FAULTS) IN ONE OUT OF M STREAMS.

M	l	k	Maximum bitwidth supported by proposed: $(M-2)l+k$	ABFT: $w - \lceil \log_2 M \rceil$
3	11	10	21	30
4	8	8	24	30
5	7	4	25	29
8	4	4	28	29
11	3	2	29	28
16	2	2	30	28
32	1	1	31	27

IV. COMPLEXITY IN LSB OPERATIONS WITH NUMERICAL ENTANGLEMENTS

We now turn our attention to the cost of performing numerical entanglement, result extraction and validation versus the cost of the LSB operation itself.

A. Complexity Analysis

Consider M input integer data streams, each comprising several samples and consider that an LSB operation op with kernel g is performed in each stream. This is the case, for example, under inner-products performed for GEMM or convolution/cross-correlation between multiple input streams for similarity detection or filtering applications or matrix-vector products in Lanczos iterations and iterative methods [37]. If the kernel g has substantially smaller length than the length of each input stream, the effective input stream size can be adjusted to the kernel length under overlap-save or overlap-add operation in convolution and cross-correlation [33] and several (smaller) overlapping input blocks can be processed independently. Similarly, block-major reordering is used in matrix products and transform decompositions for increased memory efficiency [9], [10], [12], [32], [36], [47]. Thus, in the remainder of this section we assume that N expresses *both* the

input data stream and kernel dimension under 32/64-bit integer representation.

The operations count (additions/multiplications) for stream-by-stream sum-of-products between a matrix comprising M subblocks of $N \times N$ integers and a matrix kernel comprising $N \times N$ integers (see [21], [32], [36], [42] for example instantiations within high-performance computing environments) is: $C_{\text{GEMM}} = MN^3$. For sesquilinear operations, like convolution and cross-correlation of M input integer data streams (each comprising N samples) with kernel g [see Fig. 1(a)], depending on the utilized realization, the number of operations can range from $O(MN^2)$ for direct algorithms (e.g., time-domain convolution) to $O(MN \log_2 N)$ for fast algorithms (e.g., FFT-based convolution) [33]. For example, for convolution or cross-correlation under these settings and an overlap-save realization for consecutive block processing, the number of operations (additions/multiplications) is [33]: $C_{\text{conv,time}} = 4MN^2$ for time domain processing and $C_{\text{conv,freq}} = M[(45N+15)\log_2(3N+1) + 3N+1]$ for frequency-domain processing.

As described in Section III, numerical entanglement of M input integer data streams (of N samples each) requires $O(MN)$ operations for the entanglement, extraction and validation (or recovery) per output sample. For example, ignoring all arithmetic-shifting operations (which take a negligible amount of time), based on the description of Section III the upper bound of the operations for numerical entanglement, extraction and validation/recovery is: $C_{\text{ne,conv}} = 2MN$. Similarly as before, for the special case of the GEMM operation using M subblocks of $N \times N$ integers, the upper bound of the overhead of numerical entanglement of all inputs is: $C_{\text{ne,GEMM}} = 2MN^2$. We present the percentile values obtained for $\frac{C_{\text{ne,GEMM}}}{C_{\text{GEMM}}} \times 100\%$, $\frac{C_{\text{ne,conv}}}{C_{\text{conv,time}}} \times 100\%$ and $\frac{C_{\text{ne,conv}}}{C_{\text{conv,freq}}} \times 100\%$ in Fig. 7 for typical values of N and M . For sesquilinear operations, the overhead of numerical entanglement, extraction and result validation/recovery in terms of arithmetic operations is below 0.3%. Moreover,

$$\lim_{N \rightarrow \infty} \frac{C_{\text{ne,GEMM}}}{C_{\text{GEMM}}} = \lim_{N \rightarrow \infty} \frac{C_{\text{ne,conv}}}{C_{\text{conv,time}}} = \lim_{N \rightarrow \infty} \frac{C_{\text{ne,conv}}}{C_{\text{conv,freq}}} = 0, \quad (21)$$

i.e., the overhead of the proposed approach approaches 0% as the dimension of the LSB processing increases.

For comparison purposes, Fig. 8 shows the percentile overhead of ABFT methods [Fig. 1(b) [21], [29], [42]] under: (i) the same range of values for N and M and (ii) the same fault tolerance capability⁵. Specifically, we examine the ratios: $\frac{C_{\text{ABFT,GEMM}}}{C_{\text{GEMM}}} \times 100\%$, $\frac{C_{\text{ABFT,conv,time}}}{C_{\text{conv,time}}} \times 100\%$ and $\frac{C_{\text{ABFT,conv,freq}}}{C_{\text{conv,freq}}} \times 100\%$, where $C_{\text{ABFT,GEMM}} = 2MN^2 + \frac{1}{M}C_{\text{GEMM}}$, $C_{\text{ABFT,conv,time}} = 2MN + \frac{1}{M}C_{\text{conv,time}}$ and $C_{\text{ABFT,conv,freq}} = 2MN + \frac{1}{M}C_{\text{conv,freq}}$ represent the overhead in terms of operations count (additions/multiplications) for each case. Given that time-domain and frequency-domain convolution exhibit the same percentile overhead as for the case of GEMM (with variation that is

⁵In order to keep our treatment generic, we consider as ABFT in GEMM the method that generates an additional (i.e., checksum) subblock for fault tolerance, instead of the row-column ABFT method of Huang and Abraham [20]. However, our experiments present a comparison of the proposed approach for GEMM against both approaches.

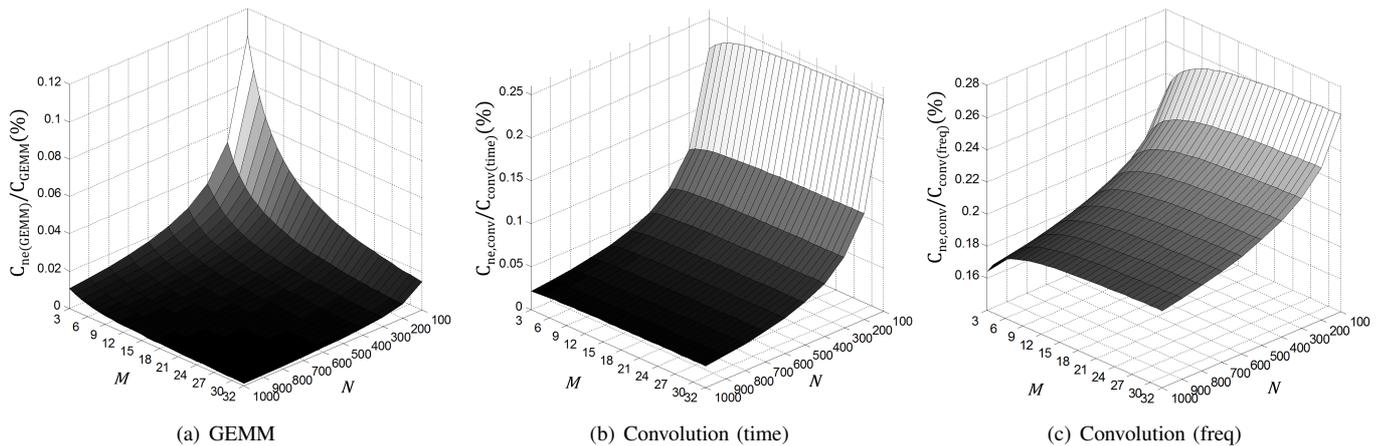


Figure 7. From left to right: Ratios of arithmetic operations for numerical entanglement, extraction and result validation/recovery versus the arithmetic operations of: generic matrix multiplication, time-domain convolution and frequency-domain convolution, with M the number of streams (or the number of subblocks within a GEMM operation) and N the kernel size (the dimension of each subblock in a GEMM operation).

limited to no more than 0.2%), Fig. 8 illustrates only the latter. As expected, the overhead of ABFT methods converges to $\frac{1}{M} \times 100\%$ as the dimension of the LSB processing operations increases, i.e.,

$$\begin{aligned} \lim_{N \rightarrow \infty} \frac{C_{\text{ABFT,GEMM}}}{C_{\text{GEMM}}} &= \lim_{N \rightarrow \infty} \frac{C_{\text{ABFT,conv,time}}}{C_{\text{conv,time}}} \\ &= \lim_{N \rightarrow \infty} \frac{C_{\text{ABFT,conv,freq}}}{C_{\text{conv,freq}}} \quad (22) \\ &= \frac{1}{M}. \end{aligned}$$

Therefore, ABFT leads to substantial overhead (above 10%) when high reliability is pursued, i.e., when $M \leq 8$. Finally, even for the low reliability regime (i.e., when $M > 8$), Fig. 8 shows that ABFT can incur more than 4% overhead in terms of arithmetic operations.

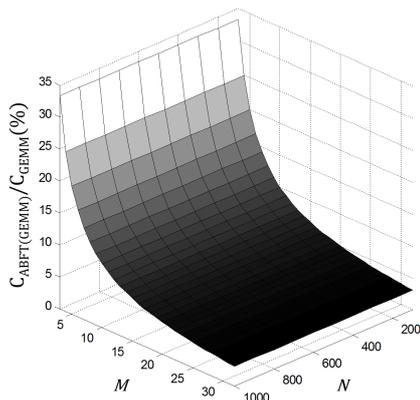


Figure 8. Ratio of arithmetic operations for ABFT generation and result validation/recovery versus the arithmetic operations of: generic matrix multiplication, with M the number of streams (i.e., GEMM subblocks) and N the kernel size (i.e., the dimension of GEMM subblocks).

B. Discussion

The comparison between Fig. 7 and Fig. 8 is illustrative for the capabilities unleashed by the proposed highly-reliable

numerical entanglement. Evidently, in the proposed approach, the *most-efficient operational area* is the leftmost part of the plots, i.e. small values of M and large values of N (small-size grouping of long streams of high-complex LSB operations). This area corresponds to the *least-efficient operational area* of ABFT. The comparison between the two figures demonstrates that, for the same fault tolerance capability (e.g., detection of one transient fault in every three co-located outputs or recovery from one fail-stop failure in three processors, which corresponds to $M = 3$), the overhead of the proposed approach is three orders of magnitude smaller than that of ABFT. Conversely, the least-efficient operational area for our approach is the rightmost part of the plots of Fig. 7 and Fig. 8, i.e. large values of M and small values of N (large-size grouping of short streams of low-complex LSB operations). This area corresponds to the *most-efficient operational area* of high-redundant ABFT methods. Nevertheless, the comparison between the two figures demonstrates that, for the same fault tolerance capability (e.g., detection of one transient fault in every 32 co-located outputs or recovery from one fail-stop failure out of 32 processors, which corresponds to $M = 32$), the overhead of the proposed approach is still one to two orders of magnitude less than that of ABFT. Overall, our approach is maximally beneficial when high reliability is desired for complex LSB operations with very low implementation overhead.

V. EXPERIMENTAL VALIDATION

All our results were obtained using an Intel Core i7-4700MQ 2.40GHz processor (Haswell architecture with AVX2 support, Windows 8 64-bit system, Microsoft Visual Studio 2013 compiler). Entanglement, disentanglement and fault detection/recovery mechanisms were realized using the Intel AVX2 SIMD instruction set for faster processing. For all cases, we also present comparisons with ABFT-based fault tolerance, the checksum elements of which were also generated using AVX2 SIMD instructions.

First, we consider the low-complexity case of basic frequency-domain processing of M streams (e.g., antialiasing, frequency-selective systems, etc.) comprising: (i) FFT of

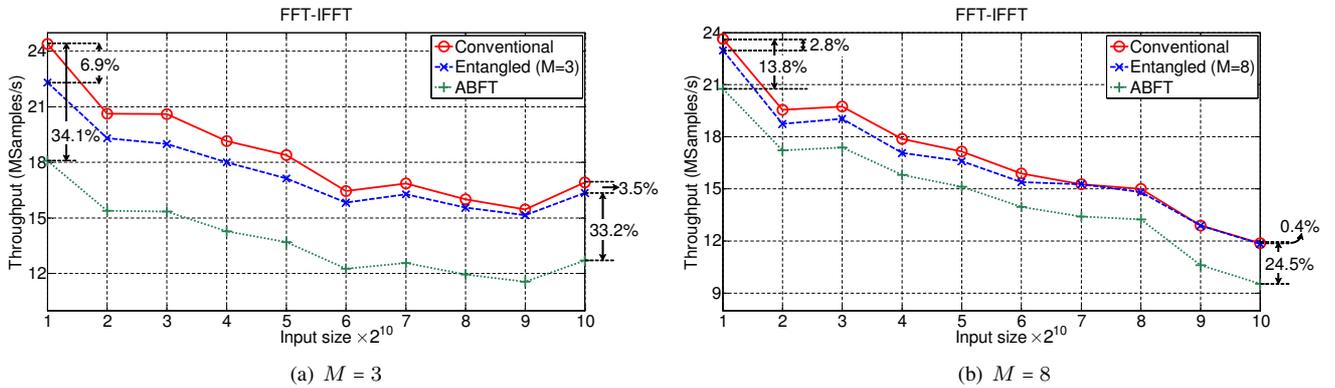


Figure 9. Throughput results for FFT-IFFT of M integer streams. “Conventional” refers to conventional (fault-intolerant) FFT realization using FFTW 3.3.3 and it is used as a benchmark under (a) $M = 3$; (b) $M = 8$.

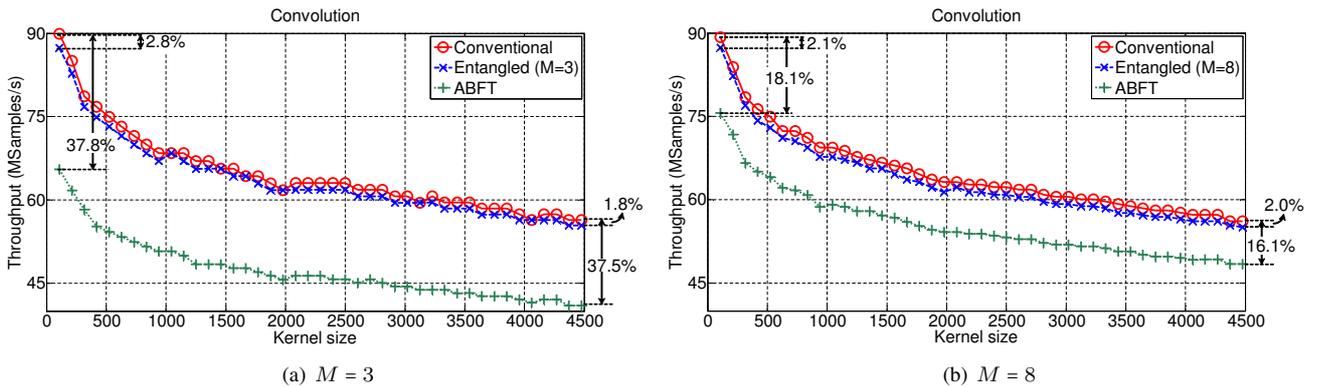


Figure 10. Throughput results for convolution of M integer streams. “Conventional” refers to conventional (fault-intolerant) convolution realization using Intel IPP 7.0 and it is used as a benchmark under (a) $M = 3$; (b) $M = 8$.

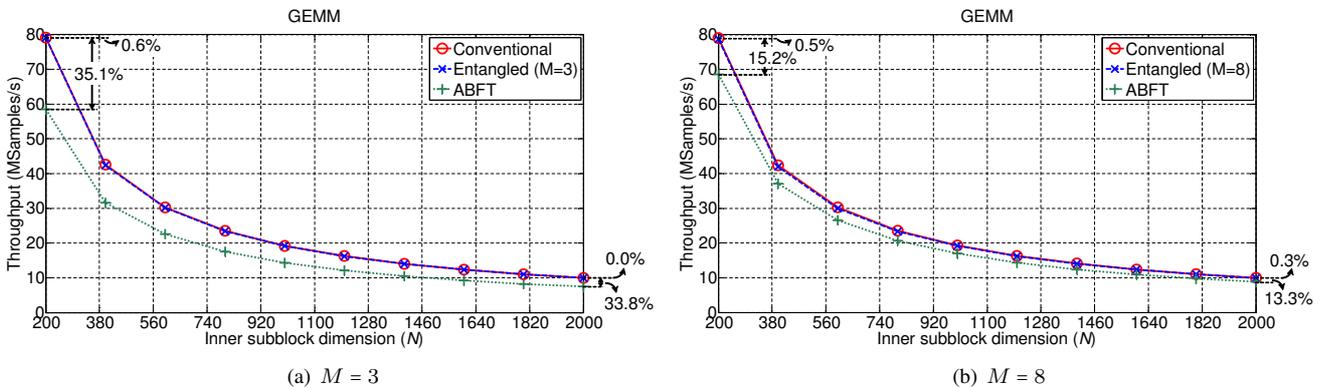


Figure 11. Throughput results for M GEMM subblocks, each comprising a $2000 \times N$ by $N \times 1200$ integer matrix product. “Conventional” refers to conventional (fault-intolerant) GEMM realization using Intel MKL 11.0 and it is used as a benchmark under (a) $M = 3$; (b) $M = 8$.

the input streams; *(ii)* attenuating a random subset of the resulting frequency components of each stream to zero by weighting with the Fourier transform of an integer frequency-selective filter; *(iii)* inverse FFT (IFFT) to obtain M integer streams. We present results for input blocks with dimension between $N \in \{2^{10}, \dots, 10 \times 2^{10}\}$ and select two sizes for number of input streams, M , which represent different fault tolerance capabilities and operational complexity for the FFT-IFFT multi-stream realization. Fig. 9 presents representative results based on the FFTW 3.3.3 library [48]. The throughput results are given in Mega-samples per second (Msamples/s).

The figure demonstrates that, under the same fault tolerance capability, our approach leads to only 0.4% to 6.9% decrease in throughput, while ABFT incurs throughput loss of 13.8% to 34.1%.

We then extend this case by performing experiments with general convolution operations of integer streams. We used Intel’s Integrated Performance Primitives (IPP) 7.0 [49] convolution routine `ippsConv_64f` that can handle the dynamic range required under convolutions with 32-bit integer inputs. We experimented with: input size of $N_{\text{in}} = 10^6$ samples, several kernel sizes between $N_{\text{kernel}} \in [100, 4500]$ samples.

Representative results are given in Fig. 10 under two settings for the number of input streams M . The results demonstrate that the proposed approach substantially outperforms ABFT, while allowing for same level of fault tolerance. In line with the theoretical predictions of Fig. 7(b) and Fig. 8(b), the decrease in throughput for the proposed approach is only 1.8% to 2.8%, while ABFT incurs 16.1% to 37.8% throughput loss against the fault-intolerant realization of convolution.

Thirdly, considering generic matrix multiplication, out of several sets of experiments performed, we present results for M matrix subblock products of $2000 \times N$ by $N \times 1200$ each, with $N \in [200, 2000]$. Fig. 11 presents results for the decrease in throughput (in Mega samples per second) against the conventional GEMM kernel realization based on the Intel MKL GEMM subroutine [47]. In line with the theoretical predictions of Fig. 7(a) and Fig. 8(a) and for the same level of fault tolerance capability, our approach is able to detect transient faults (or recover from fail-stop failures) with only 0.03% to 0.6% decrease in throughput, while ABFT requires 13.3% to 35.1% loss of throughput to perform result validation within the final GEMM result. Therefore, ABFT is 2 to 3 orders of magnitude less efficient than the proposed approach.

Finally, given that ABFT can be applied within individual GEMM subblocks following the row-wise and column-wise checksum generation within the input subblocks (as originally proposed by Huang and Abraham [20]), we also carried out a comparison against such an ABFT framework for GEMM (termed as ‘‘ABFT RC-check’’). The results demonstrated that, for the same subblock sizes as for the experiments of Fig. 11, the overhead of ABFT RC-check was between 3.5% to 5.5%, which is still 5 to 1000 times higher than the corresponding overhead of the proposed approach. Importantly, while ABFT RC-check offers guaranteed detection of transient faults in up to three outputs out of every 2000×1200 outputs, the proposed approach offers guaranteed detection/recovery of up to one fault out of every M ($M \in [3, 8]$) outputs. Therefore, in this case, the significantly-increased efficiency is also coupled with increased reliability for GEMM computations.

VI. CONCLUSION

We propose a new approach for highly-reliable LSB processing of integer data streams that is based on the novel concept of numerical entanglement. Under M input streams ($M \geq 3$), the proposed approach provides for: (i) guaranteed detection of transient faults within any single input/output stream; (ii) guaranteed recovery from any single fail-stop failure if each input stream is processed by a different core; (iii) complexity overhead that depends only on M and not on the complexity of the performed LSB operations, thus, quickly becoming negligible as the complexity of the LSB operations increases. These three features demonstrate that the proposed solution forms a *third family* of approaches for fault tolerance in data stream processing (i.e., beyond ABFT and MR) and offers unique advantages, summarized in Table I. As such, it is envisaged that it will find usage in a multitude of systems that require guaranteed reliability against transient or permanent faults in hardware with very low implementation overhead.

APPENDIX

A. Proof of Proposition 1

Proof: Notice that (11) does not use $\delta_{0,n}$. Therefore, full recovery of all outputs takes place even with the loss of δ_0 . This occurs because, for every n , $0 \leq n < N_{\text{out}}$, $\delta_{1,n}$ and $\delta_{2,n}$ contain $\hat{d}_{0,n}$ and $\hat{d}_{2,n}$ [which are recovered via (11)—this link is pictorially illustrated in Fig. 3]. Therefore, δ_1 and δ_2 suffice to recover all three output streams $\hat{\mathbf{d}}_0$, $\hat{\mathbf{d}}_1$ and $\hat{\mathbf{d}}_2$. Since the entanglement pattern is cyclically-symmetric, it is straightforward to rewrite the disentanglement process of (11) for recovery from any two out of δ_0 , δ_1 and δ_2 . ■

B. Proof of Proposition 2

Proof: Given that we can create any single entangled output stream from the other two streams, we can apply this and compare the recreated entangled output stream with the available one in order to check for occurrences of transient faults. For example, we create $\hat{\delta}_0$ starting from the recovered $\hat{\mathbf{d}}_0$ and $\hat{\mathbf{d}}_2$ (which were extracted solely from δ_1 and δ_2) and then compare it with the original values of δ_0 to detect transient faults. Specifically, for every n , $0 \leq n < N_{\text{out}}$, if

$$\delta_{0,n} - [\mathcal{S}_l \{ \hat{d}_{2,n} \} + \hat{d}_{0,n}] \neq 0 \quad (23)$$

holds, then a transient fault occurred in the triplet of $\{\delta_{0,n}, \delta_{1,n}, \delta_{2,n}\}$. Since the recovered outputs $\hat{d}_{0,n}$ and $\hat{d}_{2,n}$ stem from the $3l+k$ least-significant bits of d_{temp} , which in turn stem from the $2l+k$ bits of $\delta_{1,n}$ and $\delta_{2,n}$, the check of (23) includes all the bits of $\{\delta_{0,n}, \delta_{1,n}, \delta_{2,n}\}$. Transient faults may still remain undetected if, and only if, there exists a stream position n ($0 \leq n < N_{\text{out}}$) for which *two* or *all three* out of $\{\delta_{0,n}, \delta_{1,n}, \delta_{2,n}\}$ are corrupted in a manner that the check of (23) does not hold. Thus, we conclude that, for integer outputs with range bounded by (12), the check of (23) is *necessary and sufficient* for the detection of any transient fault in one of $\delta_{0,n}, \delta_{1,n}, \delta_{2,n}$ for all stream positions n , $0 \leq n < N_{\text{out}}$. ■

C. Proof of Proposition 3

Proof: For every output position n , $0 \leq n < N_{\text{out}}$, we are able to recover *all results of all M streams without using $\delta_{r,n}$* in (17)–(20). Therefore, the proposed method is able to recover from a single fail-stop failure in one of the M entangled streams. ■

D. Proof of Proposition 4

Proof: Given that we can create any single entangled output stream from the other $M-1$ streams, we can apply this and compare the created entangled stream with the available one to check for the occurrences of transient faults. For example, if we apply this for the r th entangled output stream, δ_r , ($1 \leq r \leq M$, $0 \leq n < N_{\text{out}}$):

$$\delta_{r,n} - [\hat{d}_{(r+M) \bmod M, n} + \mathcal{S}_l \{ \hat{d}_{(r+M-1) \bmod M, n} \}] \neq 0 \quad (24)$$

then, if (24) holds, a fault occurred in one of the M entanglements at stream position n . Since the recovered outputs $\hat{d}_{(r+m) \bmod M, n}$ and $\hat{d}_{(r+M-1) \bmod M, n}$ stem

from the $2(M-1)l + k$ least-significant bits of d_{temp} , which in turn stem from the $(M-1)l + k$ bits of $\{\delta_{0,n}, \dots, \delta_{r+1,n}, \delta_{r+1,n}, \dots, \delta_{M-1,n}\}$, the check of (24) includes all the bits of $\{\delta_{0,n}, \dots, \delta_{M-1,n}\}$. Transient faults may still remain undetected if, and only if, there exists a stream position n ($0 \leq n < N_{\text{out}}$) for which two or more out of $\{\delta_{0,n}, \dots, \delta_{M-1,n}\}$ are corrupted in a manner that the check of (24) does not hold. Thus, we conclude that, for integer outputs with range bounded by (12), the check of (24) is *necessary and sufficient* for the detection of any transient fault in one of $\{\delta_{0,n}, \dots, \delta_{M-1,n}\}$ for all stream positions n , $0 \leq n < N_{\text{out}}$. ■

REFERENCES

- [1] M. Nicolaidis, L. Anghel, N-E Zergainoh, Y. Zorian, T. Karnik, K. Bowman, J. Tschanz, S.-L. Lu, C. Tokunaga, A. Raychowdhury, et al., "Design for test and reliability in ultimate cmos," in *IEEE Design, Automation & Test in Europe Conference & Exhibition (DATE), 2012*, pp. 677–682.
- [2] Y. Andreopoulos, "Error tolerant multimedia stream processing: There's plenty of room at the top (of the system stack)," *IEEE Trans. on Multimedia*, vol. 15, no. 2, pp. 291–303, Feb 2013.
- [3] Q. Lu, L. Zhu, L. Bass, X. Xu, Z. Li, and H. Wada, "Cloud api issues: an empirical study and impact," in *Proc. 9th Int. ACM SIGSOFT Conf. on Quality of Software Arch. ACM, 2013*, pp. 23–32.
- [4] D. Poola, K. Ramamohanarao, and R. Buyya, "Fault-tolerant workflow scheduling using spot instances on clouds," *Procedia Comp. Sci.*, vol. 29, pp. 523–533, 2014.
- [5] B. Carterette, V. Pavlu, H. Fang, and E. Kanoulas, "Million query track 2009 overview," in *Proc. TREC, 2009*, vol. 9.
- [6] J. Yang, D. Zhang, A. F Frangi, and J.-Y. Yang, "Two-dimensional PCA: a new approach to appearance-based face representation and recognition," *IEEE Trans. Patt. Anal. and Machine Intel.*, vol. 26, no. 1, pp. 131–137, 2004.
- [7] G. Bradski and A. Kaehler, *Learning OpenCV: Computer vision with the OpenCV library*, O'Reilly Media, Incorporated, 2008.
- [8] I Andreopoulos et al., "A hybrid image compression algorithm based on fractal coding and wavelet transform," in *Proc. IEEE Int. Symp. Circuits and Systems, 2000 (ISCAS 2000)*. IEEE, 2000, vol. 3, pp. 37–40.
- [9] Y. Andreopoulos, A. Munteanu, et al., "A new method for complete-to-overcomplete discrete wavelet transforms," in *Int. Conf. Dig. Signal Process., 2002, (DSP 2002)*. IEEE, 2002, vol. 2, pp. 501–504.
- [10] Y. Andreopoulos, N. Zervas, et al., "A local wavelet transform implementation versus an optimal row-column algorithm for the 2d multilevel decomposition," in *Proc. IEEE Int. Conf. Image Processing, 2001, ICIP, 2001*, vol. 3, pp. 330–333.
- [11] A. Munteanu et al., "Control of the distortion variation in video coding systems based on motion compensated temporal filtering," in *Proc. IEEE Int. Conf. Image Process., ICIP 2003*. IEEE, 2003, vol. 2, pp. II–61.
- [12] Y. Andreopoulos et al., "High-level cache modeling for 2-d discrete wavelet transform implementations," *Journal of VLSI signal processing systems for signal, image and video technology*, vol. 34, no. 3, pp. 209–226, 2003.
- [13] Y. Andreopoulos and M. Van der Schaar, "Adaptive linear prediction for resource estimation of video decoding," *IEEE Trans. on Circ. and Syst. for Video Technol.*, vol. 17, no. 6, pp. 751–764, 2007.
- [14] N. Kontorinis et al., "Statistical framework for video decoding complexity modeling and prediction," *IEEE Trans. on Circ. and Syst. for Video Technol.*, vol. 19, no. 7, 2009.
- [15] B. Foo et al., "Analytical rate-distortion-complexity modeling of wavelet-based video coders," *IEEE Trans. on Signal Processing*, vol. 56, no. 2, pp. 797–815, 2008.
- [16] J Barbarien et al., "Scalable motion vector coding," *IET Electr. Lett.*, vol. 40, no. 15, 2004.
- [17] H. Jégou, F. Perronnin, M. Douze, J. Sánchez, P. Pérez, and C. Schmid, "Aggregating local image descriptors into compact codes," *IEEE Trans. Pattern Anal. and Mach. Intell.*, vol. 34, no. 9, pp. 1704–1716, 2012.
- [18] D. Fiala, F. Mueller, C. Engelmann, R. Riesen, K. Ferreira, and R. Brightwell, "Detection and correction of silent data corruption for large-scale high-performance computing," in *Proc. of the Int. Conf. on High Perf. Comp., Netw., Storage and Anal.* IEEE Computer Society Press, 2012, p. 78.
- [19] A. R. Alameldeen, I. Wagner, Z. Chishti, W. Wu, C. Wilkerson, and S.-L. Lu, "Energy-efficient cache design using variable-strength error-correcting codes," in *Proc. 38th IEEE Int. Symp. Computer Archit. (ISCA), 2011*. IEEE, 2011, pp. 461–471.
- [20] K.-H. Huang and J. A. Abraham, "Algorithm-based fault tolerance for matrix operations," *IEEE Trans. Comput.*, vol. 100, no. 6, pp. 518–528, 1984.
- [21] Z. Chen, G. E Fagg, E. Gabriel, J. Langou, T. Angskun, G. Bosilca, and J. Dongarra, "Fault tolerant high performance computing by a coding approach," in *Proc. 10th ACM SIGPLAN Symp. Princip. and Pract. Paral. Prog.*, 2005, pp. 213–223.
- [22] Franklin T Luk, "Algorithm-based fault tolerance for parallel matrix equation solvers," *SPIE, Real-Time Signal processing VIII*, vol. 564, pp. 631–635, 1985.
- [23] V. K. Stefanidis and K. G. Margaritis, "Algorithm based fault tolerance: Review and experimental study," in *International Conference of Numerical Analysis and Applied Mathematics*. IEEE, 2004.
- [24] V.S.S Nair and J.A. Abraham, "General linear codes for fault tolerant matrix operations on processor arrays," in *Int. Symp. Fault Tolerant Comput.* IEEE, 1988, pp. 180–185.
- [25] J. Sloan, R. Kumar, and G. Bronevetsky, "Algorithmic approaches to low overhead fault detection for sparse linear algebra," in *Dependable Systems and Networks (DSN), 2012 42nd Annual IEEE/IFIP International Conference on*. IEEE, 2012, pp. 1–12.
- [26] N.K. Rexford, J.; Jha, "Algorithm-based fault tolerance for floating-point operations in massively parallel systems," in *Proceedings., 1992 IEEE International Symposium on Circuits and Systems*. IEEE, May 1992, vol. 2, pp. 649,652.
- [27] C. Engelmann, H. Ong, and S. L. Scott, "The case for modular redundancy in large-scale high performance computing systems," in *Proc. IASTED Int. Conf.*, 2009, vol. 641, p. 046.
- [28] M. Fatica, "Accelerating Linpack with CUDA on heterogenous clusters," in *Proc. 2nd Workshop General Purpose Processing on Graphics Processing Units*. ACM, 2009, pp. 46–51.
- [29] G. Bosilca, R. Delmas, J. Dongarra, and J. Langou, "Algorithm-based fault tolerance applied to high performance computing," *Elsevier J. Paral. and Distrib. Comput.*, vol. 69, no. 4, pp. 410–416, 2009.
- [30] H.-J. Wunderlich, C. Braun, and S. Halder, "Efficacy and efficiency of algorithm-based fault-tolerance on GPUs," in *IEEE Internat. On-Line Testing Symp., 2013*. IEEE, 2013, pp. 240–243.
- [31] H. M Quinn, A. De Hon, and N. Carter, "CCC visioning study: system-level cross-layer cooperation to achieve predictable systems from unpredictable components," Tech. Rep., Los Alamos National Laboratory (LANL), 2011.
- [32] K. Goto and R. A Van De Geijn, "Anatomy of high-performance matrix multiplication," *ACM Trans. Math. Soft.*, vol. 34, no. 3, pp. 12, 2008.
- [33] M. A. Anam and Y. Andreopoulos, "Throughput scaling of convolution for error-tolerant multimedia applications," *IEEE Trans. Multimedia*, vol. 14, no. 3, pp. 797–804, 2012.
- [34] A. Kadyrov and M. Petrou, "The "invaders" algorithm: Range of values modulation for accelerated correlation," *IEEE Trans. Patt. Anal. and Machine Intel.*, vol. 28, no. 11, pp. 1882–1886, 2006.
- [35] D. Anastasia and Y. Andreopoulos, "Software designs of image processing tasks with incremental refinement of computation," *IEEE Trans. Image Process.*, vol. 19, no. 8, pp. 2099–2114, 2010.
- [36] D. Anastasia and Y. Andreopoulos, "Throughput-distortion computation of generic matrix multiplication: Toward a computation channel for digital signal processing systems," *IEEE Trans. Signal Process.*, vol. 60, no. 4, pp. 2024–2037, 2012.
- [37] G. H Golub and C. F Van Loan, *Matrix computations*, Johns Hopkins University Press, 1996.
- [38] C. Lin, B. Zhang, and Y. F. Zheng, "Packed integer wavelet transform constructed by lifting scheme," *IEEE Trans. Circ. and Syst. for Video Technol.*, vol. 10, no. 8, pp. 1496–1501, 2000.
- [39] P. M. Fenwick, "The Burrows–Wheeler transform for block sorting text compression: principles and improvements," *The Comp. J.*, vol. 39, no. 9, pp. 731–740, 1996.
- [40] M. A. Anam and Y. Andreopoulos, "Failure mitigation in linear, sesquilinear and bijective operations on integer data streams via numerical entanglement," in *Proc. 2015 IEEE 21th Int. On-Line Testing Symposium (IOLTS)*, July 2015.
- [41] M.-Y. Ko, C. Zissulescu, S. Puthenpurayil, S. S. Bhattacharyya, B. Kienhuis, and E. F. Deprettere, "Parameterized looped schedules for compact representation of execution sequences in dsp hardware and software implementation," *IEEE Trans. on Signal Processing*, vol. 55, no. 6, pp. 3126–3138, 2007.

- [42] D. G Murray and S. Hand, "Spread-spectrum computation," in *Proc. USENIX 4th Conf. Hot Top. in Syst. Dependab.*, 2008, pp. 5–9.
- [43] D. Anastasia and Y. Andreopoulos, "Linear image processing operations with operational tight packing," *IEEE Signal Process. Lett.*, vol. 17, no. 4, pp. 375–378, 2010.
- [44] I. Anarado, M.A. Anam, D. Anastasia, F. Verdicchio, and Y. Andreopoulos, "Highly-reliable integer matrix multiplication via numerical packing," in *Proc. 2013 IEEE 19th Int. On-Line Testing Symposium (IOLTS)*, July 2013, pp. 19–24.
- [45] I. Anarado and Y. Andreopoulos, "Mitigation of fail-stop failures in integer matrix products via numerical packing," in *Proc. 2015 IEEE 21th Int. On-Line Testing Symposium (IOLTS)*, July 2015.
- [46] J. D Allen, "An approach to fast transform coding in software," *Elsevier Signal Process.: Image Comm.*, vol. 8, no. 1, pp. 3–11, 1996.
- [47] MKL Intel, "Intel math kernel library," 2007.
- [48] M. Frigo and S. G Johnson, "Fftw: Fastest fourier transform in the west," *Astrophysics Source Code Library*, vol. 1, pp. 01015, 2012.
- [49] S. Taylor, *Intel Integrated Performance Primitives: How to Optimize Software Applications Using Intel IPP*, 2003.



Mohammad Ashraf Anam obtained the PhD in Electronic Engineering from University College London (Lombardi Prize for the Best PhD thesis in Electronic Engineering) and is currently post-doctoral research associate in the Department of Electronic and Electrical Engineering, University College London, London, UK. His research interests are in error tolerant computing, and reliable cloud computing.



Yiannis Andreopoulos (M'00-SM'14) obtained the PhD in Applied Sciences from the Vrije Universiteit Brussel, Belgium. He is Reader (Assoc. Professor) in Data and Signal Processing Systems in the Department of Electronic and Electrical Engineering of University College London (U.K.). His research interests are in wireless sensor networks, error-tolerant computing and multimedia systems. He received the 2007 Most-Cited Paper Award from the Elsevier EURASIP Signal Processing: Image Communication journal and a best paper award from

the 2009 IEEE Workshop on Signal Processing Systems. He was Special Sessions Co-Chair of the 10th International Workshop on Image Analysis for Multimedia Interactive Services (WIAMIS 2009) and Programme Co-Chair of the 18th International Conference on Multimedia Modeling (MMM 2012) and the 9th International Conference on Body Area Networks (BODYNETS 2014). He has been an Associate Editor of the IEEE Transactions on Multimedia, the IEEE Signal Processing Letters and Image and Vision Computing (Elsevier).