

SIMD memory circuit and methodology to support upsampling, downsampling and transposition

Citation for published version (APA):

Kampen, van, D., Berkel, van, C. H., Goossens, S. L. M., Kloosterhuis, W. E. H., & Zissulecsu-Ianculescu, C. (2013). SIMD memory circuit and methodology to support upsampling, downsampling and transposition. (Patent No. US20130091339).

Document status and date:

Published: 11/04/2013

Document Version:

Publisher's PDF, also known as Version of Record (includes final page, issue and volume numbers)

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy

If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.



(19) **United States**

(12) **Patent Application Publication**
Van Kampen et al.

(10) **Pub. No.: US 2013/0091339 A1**

(43) **Pub. Date: Apr. 11, 2013**

(54) **SIMD MEMORY CIRCUIT AND
METHODOLOGY TO SUPPORT
UPSAMPLING, DOWNSAMPLING AND
TRANSPOSITION**

G06F 9/302 (2006.01)
G06F 9/38 (2006.01)
G06F 15/76 (2006.01)
G06F 9/305 (2006.01)

(75) Inventors: **David Van Kampen**, Eindhoven (NL);
Kees Van Berkel, Heeze (NL); **Sven
Goossens**, Eindhoven (NL); **Wim
Kloosterhuis**, Waalre (NL); **Claudiu
Zissulescu-Ianculescu**, Eindhoven (NL)

(52) **U.S. Cl.**
USPC **712/4**; 712/200; 712/E09.017; 712/E09.018;
712/E09.033; 712/E09.045; 712/E09.023;
712/E09.034

(73) Assignee: **ST-Ericsson SA**, Plan-les-Ouates (CH)

(21) Appl. No.: **13/253,754**


(22) Filed: **Oct. 5, 2011**

Publication Classification

(51) **Int. Cl.**
G06F 9/312 (2006.01)
G06F 9/315 (2006.01)

(57) **ABSTRACT**

An apparatus and method for creation of reordered vectors from sequential input data for block based decimation, filtering, interpolation and matrix transposition using a memory circuit for a Single Instruction, Multiple Data (SIMD) Digital Signal Processor (DSP). This memory circuit includes a two-dimensional storage array, a rotate-and-distribute unit, a read-controller and a write to controller, to map input vectors containing sequential data elements in columns of the two-dimensional array and extract reordered target vectors from this array. The data elements and memory configuration are received from the SIMD DSP.






FIG. 1
PRIOR ART

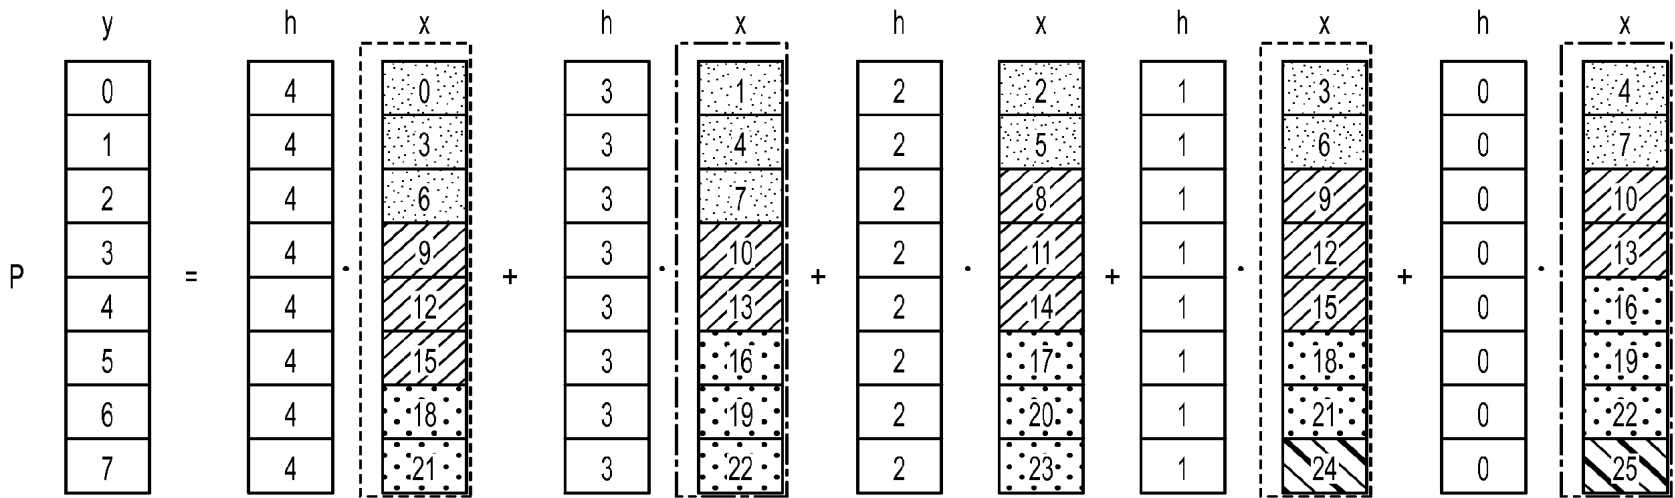


FIG. 2




FIG. 3




FIG. 4




FIG. 5

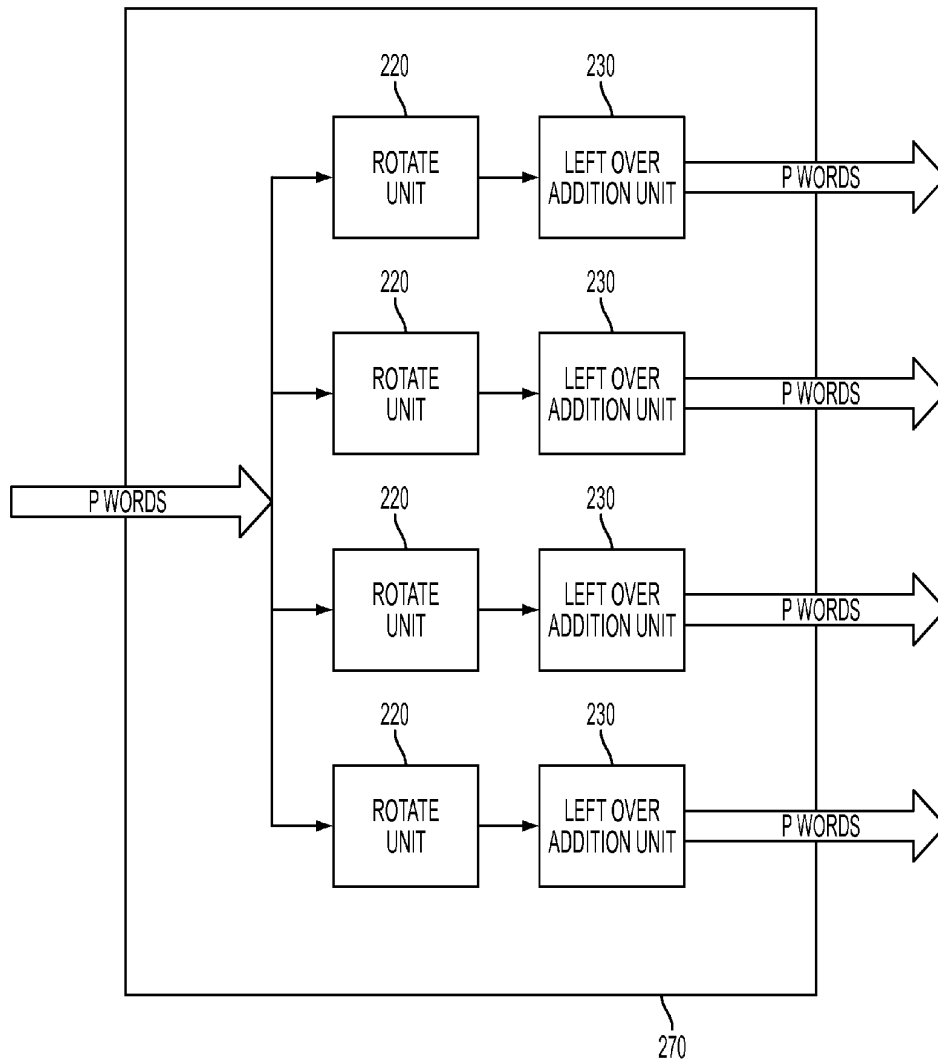


FIG. 6A




FIG. 6B




FIG. 6C

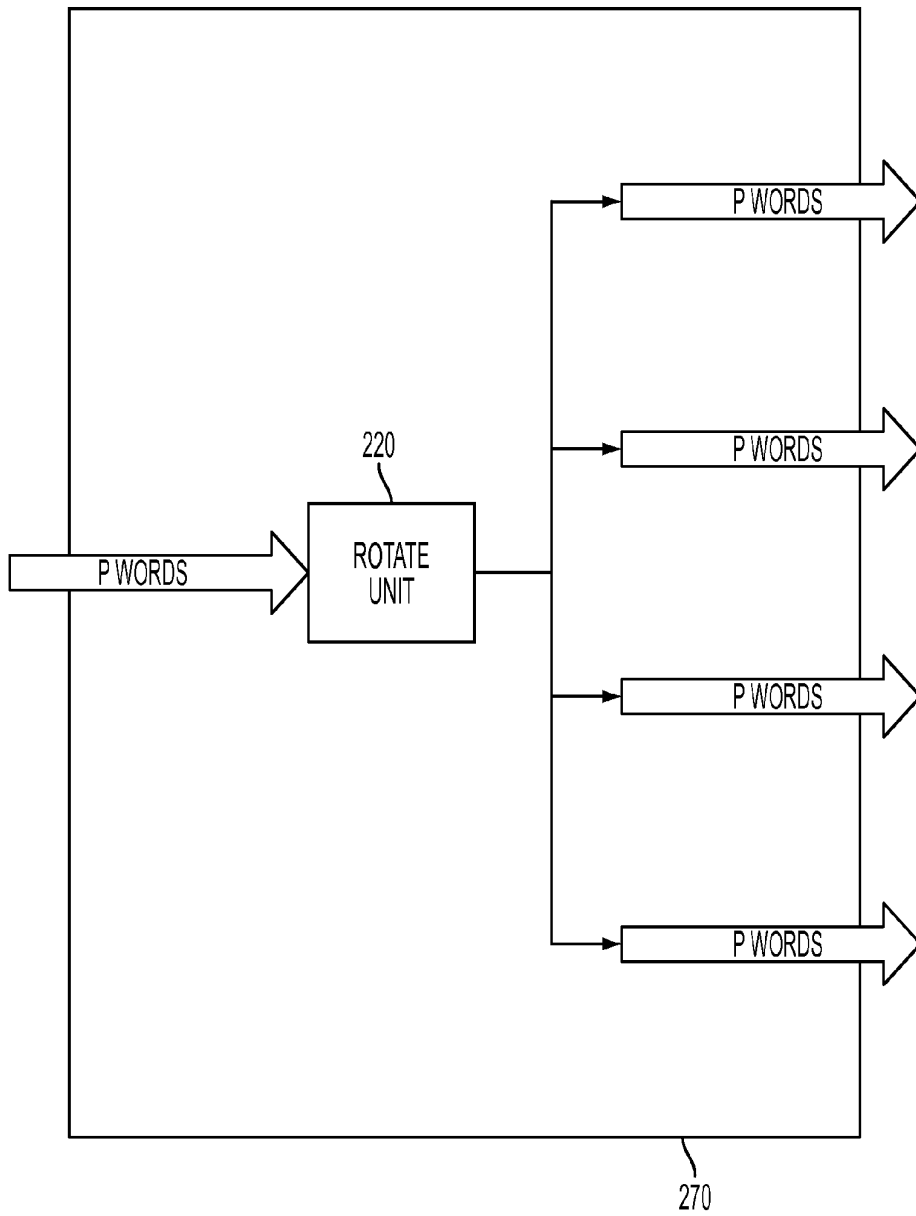


FIG. 6D

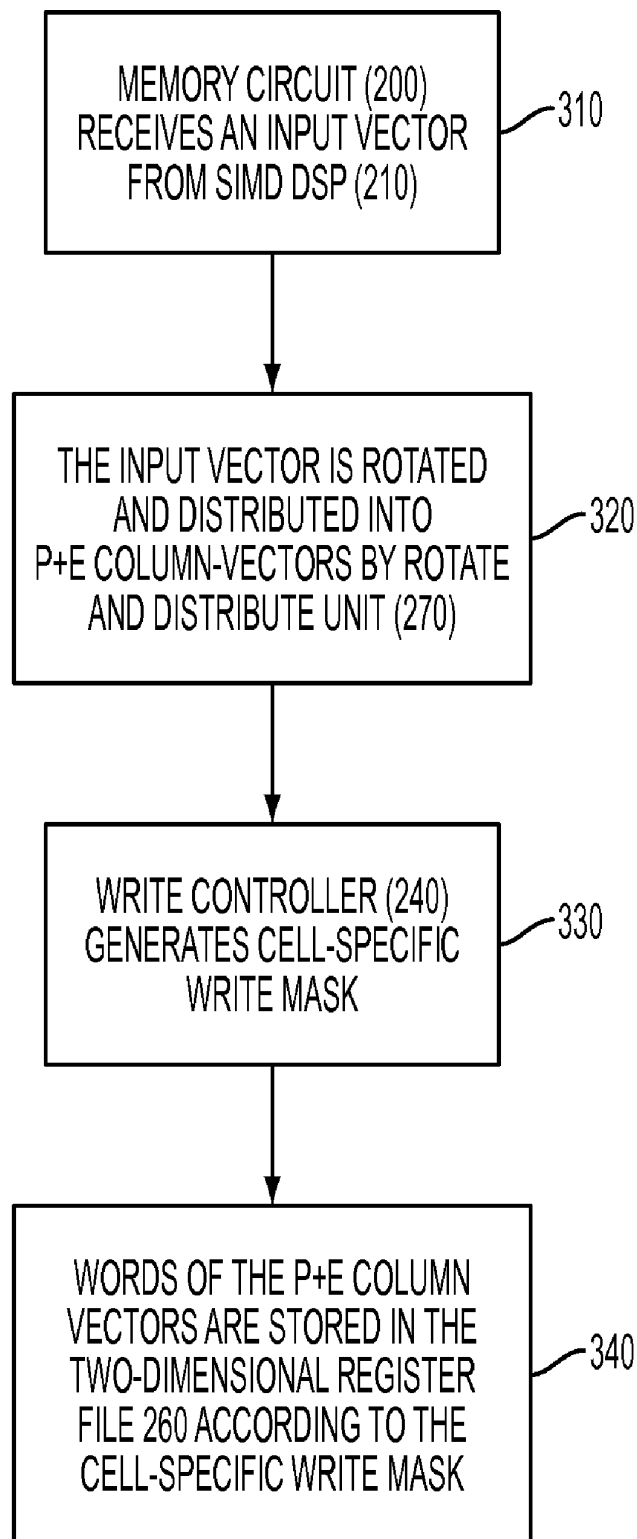


FIG. 7

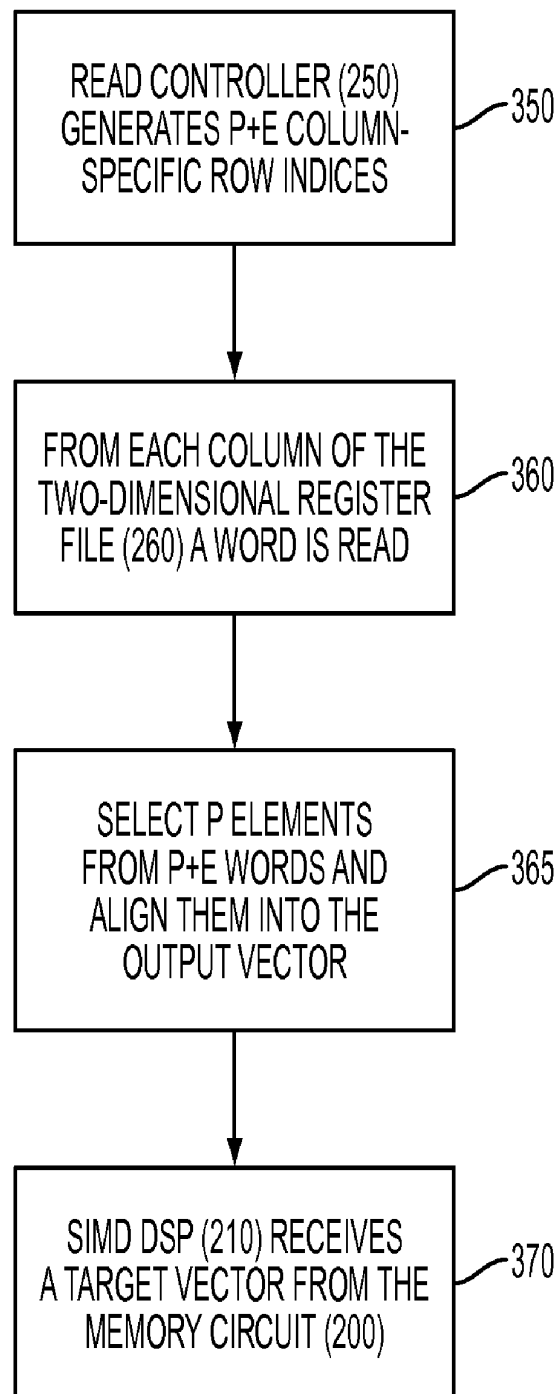


FIG. 8




FIG. 9




FIG. 10

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
1	1		6		11		16		21		26		31		36	
2		1		6		11		16		21		26		31		36
3	2		7		12		17		22		27		32		37	
4		2		7		12		17		22		27		32		37
5	3		8		13		18		23		28		33		38	
6		3		8		13		18		23		28		33		38
7	4		9		14		19		24		29		34		39	
8		4		9		14		19		24		29		34		39
9	5		10		15		20		25		30		35		40	
10		5		10		15		20		25		30		35		40
11																
12																
13																
14																
15																
16																

FIG. 11




FIG. 12




FIG. 13



FIG. 14




FIG. 15

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
1	1		9				17		25				33			
2		1		9				17		25				33		
3	2		10				18			26			34			
4		2		10				18			26			34		
5	3				11		19			27			35			
6		3				11		19			27			35		
7	4				12		20			28						
8		4				12		20			28					
9	5				13			21		29						
10		5				13			21		29					
11			6		14			22		30						
12				6		14			22		30					
13			7		15			23				31				
14				7		15			23					31		
15			8			16		24				32				
16				8			16		24					32		

1210 ↗

FIG. 16

**SIMD MEMORY CIRCUIT AND
METHODOLOGY TO SUPPORT
UPSAMPLING, DOWNSAMPLING AND
TRANSPOSITION**

TECHNICAL FIELD

[0001] The invention relates to a Single Instruction, Multiple Data (SIMD) memory circuit and method of use.

BACKGROUND

[0002] In a relatively short period of time the demand for high speed, integrated wireless telecommunication services has increased rapidly. Currently, a number of third and fourth generation wireless communication standards have emerged, e.g. UMTS/TDD, UMTS/FDD, TD-SCDMA, etc. Single Instruction, Multiple Data (SIMD), is a class of parallel computers that has been used to facilitate wireless communications due to the large amount of data that must be accessed and manipulated. As illustrated in FIG. 1, computers with multiple parallel processor elements **100**, i.e. vector Digital Signal Processors (DSPs), simultaneously perform the same operation retrieved from a common instruction pool **120** on multiple data elements obtained from a data pool **130**. Thus, such SIMD machines exploit data level parallelism. The data pool **130** and common instruction pool **120** may be volatile or non-volatile memory or any data storage device that is readable by a computer system.

[0003] Application specific processors, i.e. Application-Specific Integrated Circuits (ASICs), for decimation or similar algorithms typically have only limited flexibility, i.e. limited support for variation of the algorithm's parameters, which make it difficult to deploy them in multi-standard transceivers. Further, their hardware cannot be reused for other functions. Still further, they are not updatable, which would be useful to debug their functionality or extend their functionality when standards or products evolve.

[0004] Classical DSP based solutions don't have the aforementioned limitations of ASICs. However, due to their processing of one element at a time, such DSPs would have to run at too high clock frequencies to still be able to support these applications in a power-efficient manner.

[0005] Vector DSPs can exploit data parallelism. The sample based decimation filtering approach has one important cause for reduced efficiency: the number of filter coefficients K is not necessarily a multiple of P . Further, intra-vector addition is also a required operation, which can be an issue. For block based filtering on a vector DSP, the creation of the required base set of M down-sampled vectors, e.g. through reordering, is a complex and time consuming task for a vector DSP. This task can be considered as overhead to the decimation filtering proper and therefore should be kept to a minimum.

[0006] A vector DSP having a memory accessible in two dimensions can be used to prepare M base vectors to which iteratively new samples are appended. They cannot, however, provide the desired target vectors for decimation filters with $K > P$ from their memory directly. Secondly, parallel filling of the memory and reading out for different software pipeline stages is only possible if both reading and writing is supported in both directions. An important loss of efficiency is seen since the same data samples are loaded multiple times.

[0007] To overcome the shortcomings of the related art an extension of a programmable vector processor is needed. A

programmable processor gives flexibility, typically at the expense of increased area and power consumption relative to dedicated hardware solutions.

[0008] Further, the problem of decimation filtering algorithms is their non-consecutive memory access. It needs vectors of a downsampled signal stream. These can be created in many ways with various (specialized) memory circuits or devices (e.g. using registers and shuffle units on the EVP itself). However, there is an increased tradeoff in complexity/flexibility of these memory circuits and the remaining overhead/load/costs for the processor in generating the target decimated vectors.

[0009] An article summarizing many alternative multi-access memory systems: (Park J. W., "Multiaccess memory system for attached SIMD computer", IEEE Transactions on Computers, Volume: 53, Issue: 4, April 2004, pages: 439-452 incorporated herein in its entirety by reference) provides various conflict-free memory access patterns for reading/writing vectors of data. A conflict-free access enables single cycle memory access. There are memory circuits that support conflict-free writing of a vector of P elements into a column of a virtual array of $P \times P$ elements. With additional support of conflict-free reading of a vector of P elements from a row of this virtual array, one can generate decimated vectors. Other memory circuits directly provide strided memory access (i.e. down sampling) in a virtual array of $1 \times (P \times P)$ elements.

[0010] Still further, decimation and interpolation filtering have another problem, symmetries in the filter coefficient sequence, which is typical for these filters cannot be exploited for any K . So these implementations will do maximally one filter coefficient per multiplication operation instead of two. Depending on the processor's VLIW resource configuration, the throughput is therefore about half of what is theoretically possible. Thirdly, source samples are typically reloaded from memory during the processing.

[0011] For block based filtering on a vector processor the creation of the required base set of M down-sampled vectors, e.g. through reordering, is a complex and time consuming task. With support of a two dimensionally accessible memory the efficiency already improves, by reducing the load of this overhead. However, through the present invention, this process can be made even more efficient than with a basic 2D memory. Still, exploiting symmetries of the filter coefficient array with any two dimensionally accessible memory of size $P \times P$ is a problem. This is because the target sample vectors of both mirror coefficients cannot easily be derived from the same base vector, as the shift distance in the down-sampled stream is larger than one sample. Only for $K \leq P$, symmetric filters can be efficiently supported as these target vectors are stored in the memory. Again, for $K > P$, depending on the processor's VLIW resource configuration, the throughput is therefore about half of what is theoretically possible.

[0012] The implementation of such memory circuits supporting the aforementioned memory operations are however complex and expensive. They typically have besides (at least) P independently writable/indexable banks both a rotator at the input/write port and a rotator per read port.

[0013] The above information disclosed is only for enhancement of understanding of the background of the invention and therefore it may contain information that does not form the prior art that is already known to a person of ordinary skill in the art.

SUMMARY

[0014] An exemplary embodiment of the present invention relates to a configurable memory circuit for a Single Instruction, Multiple Data (SIMD) Digital Signal Processor (DSP), for generating reordered output vectors from sequential input data vectors. This memory circuit includes a rotate-and-distribute unit to rotate and distribute an input data vector and a two-dimensional register file to receive and store vectors generated by the rotate-and-distribute unit. A write controller is used to control the mapping of input data elements in the two-dimensional register file, corresponding to a mapping configuration by configuring the rotate-and-distribute unit and generating write-enable signals for all cells of the two-dimensional register file. Finally, a read controller is used to select, on a per column basis, one data element to be received in the output vector, corresponding to the mapping configuration.

[0015] Further, the present invention provides for a method for generating reordered output vectors from sequential input data vectors utilizing a configurable memory circuit for a Single Instruction, Multiple Data (SIMD) Digital Signal Processor (DSP) having a rotate-and-distribute unit, a two-dimensional register file, a write controller and a read controller. This method begins by rotating and distributing an input data vector by the rotate-and-distribute unit generating output vectors and receiving and storing the target vectors generated by the rotate-and-distribute unit in the two-dimensional register file. Thereafter, the step of controlling the mapping of input data elements in the two-dimensional register file corresponding to a mapping configuration by configuring the rotate-and-distribute unit and generating write-enable signals for all cells of the two-dimensional register file is performed, and the selecting on a per column basis one data element to be received in the output vector, corresponding to the mapping configuration by the read controller.

BRIEF DESCRIPTION OF THE DRAWINGS

[0016] A more complete appreciation of the invention, and many of the attendant advantages thereof, will be readily apparent as the same becomes better understood by reference to the following detailed description when considered in conjunction with the accompanying drawings in which like reference symbols indicate the same or similar components, wherein:

[0017] FIG. 1 is a general systems diagram of a conventional Single Instruction, Multiple Data (SIMD) computer system;

[0018] FIG. 2 is an implementation of a block-based processing decimation algorithm depicting the manipulation of a two-dimensional data array systems which may be executed using the SIMD computer system illustrated in FIG. 1;

[0019] FIG. 3 is an implementation of an exemplary configuration of reordered target vectors created from sequential input vectors for decimation;

[0020] FIG. 4 is an implementation of an exemplary configuration of reordered target vectors created from sequential input vectors for interpolation;

[0021] FIG. 5 is a diagram illustrating the memory circuit 200 of an exemplary embodiment of the present invention;

[0022] FIGS. 6A-6D are alternative exemplary embodiments of the rotate-and-distribute unit 270 shown in FIG. 5 of the present invention;

[0023] FIG. 7 is a flowchart illustrating the operating steps of the write function executed by the memory circuit 200 shown in FIG. 5 in an exemplary embodiment of the present invention;

[0024] FIG. 8 is a flowchart illustrating the operating steps of the read function executed by the memory circuit 200 shown in FIG. 5 in an exemplary embodiment of the present invention;

[0025] FIGS. 9-10 are exemplary data structures generated by the memory circuit 200 shown in FIG. 5 executing the steps of the flowchart shown FIGS. 7-8 in an exemplary embodiment of the present invention for the rotate-and-distribute circuit 270 according to FIGS. 6A-6B;

[0026] FIG. 11 is a table illustrating the variable word width features of the present invention generated by the embodiment of the present invention shown in FIGS. 5, 7 and 8 of the present invention for rotate-and-distribute circuit 270 according to FIGS. 6A-6B;

[0027] FIG. 12 is an exemplary data structure generated by the memory circuit 200 shown in FIG. 5, executing the steps of the flowchart shown in FIGS. 6 and 7, in an exemplary embodiment of the present invention for the rotate-and-distribute circuit 270 according to FIG. 6C-6D;

[0028] FIG. 13 illustrates the mapping for symmetric decimation filtering belonging to the rotate-and-distribute circuit 270 according to 6C;

[0029] FIG. 14 illustrates the mapping for different software pipeline stages belonging to the rotate-and-distribute circuit 270 according to 6D;

[0030] FIG. 15 is an exemplary data structure generated by the memory circuit 200 shown in FIG. 5 executing the steps of the flowchart shown FIGS. 6 and 7 in an exemplary embodiment of the present invention for the rotate-and-distribute circuit 270 according to FIG. 6C-6D; and

[0031] FIG. 16 is a table illustrating the variable word width features generated by the embodiment of the present invention shown in FIGS. 5, 7 and 8 of the present invention for rotate-and-distribute circuit 270 according to FIGS. 6C-6D.

DETAILED DESCRIPTION

[0032] In the following detailed description, only certain exemplary embodiments of the present invention have been shown and described, simply by way of illustration. As those skilled in the art would realize, the described embodiments may be modified in various different ways, all without departing from the spirit or scope of the present invention. Accordingly, the drawings and description are to be regarded as illustrative in nature and not restrictive. Also, when an element is referred to as being "connected to" another element, it can be directly connected to another element or be indirectly connected to another element with one or more intervening elements interposed therebetween. Hereinafter, like reference numerals refer to like elements.

[0033] Further, whenever software or computer programs are referred to in reference to the drawings these items may include, but not limited to, algorithms, software, firmware, applets, programming code or languages, and code segments which are stored on a computer readable medium and executable on a computer or other processor-based apparatus.

[0034] Several factors are important in the design of such SIMD machines when used for wireless communications. Decimation, i.e. low-pass anti-alias filtering combined with down-sampling, is an important function in digital wireless communication. In digital front end (DFE) subsystems of

wireless receivers for modern standards like IEEE 802.11a-n, UMTS, LTE, the load of decimation is a dominant factor. Therefore, it is important that its implementation is power and cost efficient. For multi-standard transceiver solutions, flexibility of the decimation chain of the transceiver is another highly desired property.

[0035] The formula for the output of a decimating Finite Impulse Response (FIR) filter is:

$$y[m] = \sum_{i=0}^{K-1} h[K-1-i]x[mM+i] \quad \text{Equation 1}$$

[0036] In Equation 1, output element $y[m]$ is calculated by weighing a sequence of K samples of input signal $x[n]$ with coefficients $h[k]$ and summing them. Output elements $y[m]$, are based on an input sample sequence with an offset that is a multiple of M , where M is the decimation or down sampling factor.

[0037] Another important characteristic of implementations of a decimating FIR filter is the non-sequential data access patterns. Other algorithms found in the wireless communication domain with the same characteristic are: signal-up and down-sampling, interpolation, i.e. spectral image filtering combined with signal-up sampling, fractional sample-rate conversion and matrix transposition, where the latter can be considered as a special case of down-sampling.

[0038] Decimation and interpolation filters are typically characterized by an impulse response with linear phase. For digital filter processing this implies a symmetric filter coefficient sequence, i.e. $h[n]$ is equal to $h[K-1-n]$, which may be exploited in some implementations of the filtering algorithm: the number of multiplications can be halved, by replacing multiplications of two samples belonging to the same coefficient value with one addition and one multiplication.

[0039] Implementations of decimation and interpolation algorithms often have a polyphase filter structure, meaning the decomposition into M , respectively L , parallel subfilters. Each subfilter filters a down-sampled version of the input stream $x[n]$, with different start offset.

[0040] Decimation and similar algorithms can be implemented, either using application specific integrated circuits or using programmable DSPs. In addition, any type of processing device may be used. As discussed above, SIMD DSPs that operate on vectors of data can exploit the available data level parallelism in these algorithms. SIMD DSPs may have support for load operations of memory words that are not aligned on a vector boundary. With this capability, it is possible to implement decimation in a "sample based" approach, i.e. a direct implementation of equation 1, where in each iteration, a single output value of the low-pass filter operation is calculated, by filtering a consecutive array of K samples. A vector of P output elements is thus produced in P iterations. This implementation requires support for intra-vector addition, i.e. addition of all P elements from a vector.

[0041] Alternatively, vector DSPs can implement decimation in a "block based" approach, if they have some support for loading/creating vectors from non-sequential samples. In the filter step, a vector of P output elements is produced in K iterations, where in each iteration, a scalar filter coefficient is multiplied with a down-sampled signal vector, with a specific start offset. The corresponding equation is,

$$\begin{bmatrix} y_{m+0} \\ y_{m+1} \\ y_{m+2} \\ \vdots \\ y_{m+P-1} \end{bmatrix} = h_{K-1} \begin{bmatrix} x_{M(m+0)} \\ x_{M(m+1)} \\ x_{M(m+2)} \\ \vdots \\ x_{M(m+P-1)} \end{bmatrix} + \quad \text{Equation 2}$$

$$h_{K-2} \begin{bmatrix} x_{M(m+0)+1} \\ x_{M(m+1)+1} \\ x_{M(m+2)+1} \\ \vdots \\ x_{M(m+P-1)+1} \end{bmatrix} + \dots + h_0 \begin{bmatrix} x_{M(m+0)+(K-1)} \\ x_{M(m+1)+(K-1)} \\ x_{M(m+2)+(K-1)} \\ \vdots \\ x_{M(m+P-1)+(K-1)} \end{bmatrix}$$

[0042] The main advantages of block based decimation on a vector DSP are that the algorithm scales very well with K and that it does not need intra-vector addition.

[0043] As indicated in Equation 2, every M^{th} iteration uses a down-sampled version of the input signal $x[n]$ with the same start offset, only time shifted with the down-sample step M . This is visualized in FIG. 2 using shading/hashing of the boxes around input vectors with the same start offset. There are M down-sampled streams with a different start offset. A decimation algorithm may be implemented where at the start of the processing of a block M base vectors are created. Thereafter, for each filter iteration, a new sample must be appended to a sample vector and the oldest sample is removed.

[0044] FIG. 2 is an example of a block based decimation algorithm with $M=3$, $K=5$, $P=8$. Note that there are M base vectors that are a down-sampled version of the input signal with factor M and a different start offset. All the input vectors can be derived by shifting new samples into the base vector.

[0045] Further, as exemplified in FIG. 3, the formula for relating M target down-sampled vectors (with decimation factor M , variable start offset in the sample stream) from M consecutive input vectors (starting at an arbitrary point in the sequential input stream):

```

for i: 0..M-1
  for j: 0..P-1
    vector_index = floor((j*M+i)/P)
    element_index = (j*M+i)%P
    target_vector[i][j] = input_vector[s][vector_index]
    [element_index]
  end
end

```

[0046] In which $j*M+i$ is the index in the M consecutive input vectors combined into one array. Interpolation, i.e. up-sampling followed by image filtering, has SIMD implementation problems similar to decimation, because of non-sequential element ordering. The polyphase decomposed implementation uses L block-based (sub) filters of length K/L . Each of these L filters gives as output a decimated version of the desired output signal, only with a different start offset. The output vectors of the L subfilters have to be interleaved to form L output vectors of the desired output signal. This interleaving problem is similar to the generation of down-sampled vectors for decimation filtering.

[0047] As illustrated in FIG. 4, the interpolation algorithm receives an input vector where the target vectors can be related to input vector using the following algorithm:

```

for i: 0..M-1
  for j: 0..P-1
    vector_index = (i*P+j)%M
    element_index = floor((i*P+j)/M)
    target_vector[i][j] = input_vector[vector_index]
                        [element_index]
  end
end

```

[0048] For a more detailed explanation of the aforementioned interleaving problem, see S. L. M. Goossens—“An SIMD register file with support for dual-phase decimation and transposition”, (M. Sc. Thesis, TU Eindhoven. June 2010, conducted at ST-Ericsson/DSP-IC), herein after referred to as Goossens and incorporated herein by reference in its entirety.

[0049] Vector DSPs typically come without support for loads of vectors with non-sequential memory elements to make the target vectors for block based filtering. There can be support for vector element shuffling and mask based selection of elements from two vectors, which is sufficient to create the desired target vectors. An Embedded Vector Processor (EVP) is an example of one such processor element **100**.

[0050] Some vector DSPs have limited transposition support. Special register files for transposition have been proposed (see Shahbahrami, Juurlink, Vassiliadis; “Matrix register file and extended subwords: two techniques for embedded media processors”; Proceedings of the 2nd conference on Computing Frontiers, ACM, 2005: and Jung, Berg, Kim, Kim “A register file with transposed access mode”; Proceeding of International Conference on Computer Design, 2000, incorporated herein by reference in their entirety), having both a row and column wise access mode. In principle, having a memory that is accessible in two dimensions is sufficient to generate down-sampled target vectors. First, P vectors are stored in one dimension, each with P samples from a continuous input signal and a relative offset of M samples between their first elements. Second, P target vectors can be read out from the other dimension. Note that block level software pipelining can be supported if both reading and writing is supported in both row and column direction. Per iteration the direction in which data is stored and the direction in which it is read out is changed. Another specialized register file designed for transposition uses diagonal write and read patterns as further described in Hanounik, Hu; “Linear-time matrix transpose algorithms using vector register file with diagonal registers”; Proceedings of 15th International Parallel and Distributed Processing Symposium, incorporated herein in its entirety by reference. Further, Jayasena, Erez, Ahn, and Dally; “Stream register files with indexed access”; Proceedings of the 10th International Symposium on High Performance Computer Architecture, 2004, incorporated herein in its entirety by reference, proposes a conflict resolving register file used by multiple stream processors with indexed access modes. It exploits an SRAM like structure of register file for indexing per bank, but that does not translate well to vector processing.

[0051] As discussed, a basic 2D memory circuit, configured with P rows and P columns, in which the input data is strictly written in one column at a time, can be used to generate target vectors. From P input vectors written in columns of the 2D memory, P target vectors can be obtained, by reading rows in the 2D memory. Exemplary embodiments are

directed to dimensioning the number of columns of the incorporated 2D memory circuit as P rows and P+E columns and providing one or two extended read ports that are able to select P elements out of P+E elements and give it a proper alignment in the output vector. Thus more than P target vectors can be generated with the memory. For a basic 2D memory extended with E columns, M(1+E)+(P-M) target vectors can be extracted from P+E source vectors. For the memory circuits of the present invention, higher efficiencies can be obtained. An exemplary embodiment will be discussed where M(1+E) target vectors can be extracted from M+ceil(ME/P) source vectors. Even higher efficiencies are possible under certain conditions for other embodiments.

[0052] With the proposed extensions, for certain filter lengths ($K \leq M(1+E) + (P-M)$) for a basic 2D memory respectively $K \leq M(1+E)$ all desired target vectors for block based symmetric decimation filtering are available in the memory after the input vector store phase. Having two read ports allows the parallel reading out of two target vectors, specifically the target vectors that are to be multiplied with the coefficients $h[j]$ and $h[K-1-i]$, that have the same value because of symmetry. Both the storage capacity and the read port bandwidth enable a very efficient implementation of symmetric block based filtering on certain VLIW vector DSPs. In the key loop the reading out of two down-sampled target vectors, the addition of two target vectors (of the previous read step) and a multiply accumulate operation can go in parallel. By processing two filter coefficients per cycle, block based decimation filtering with the invented PE×P 2D memory, can have twice the throughput of implementations using a P×P 2D memory.

[0053] Hereinafter, configurations and operations of exemplary embodiments according to the present invention will be described in detail with reference with the accompanying drawings.

[0054] A parallel memory system (or multi-access memory circuit) for strided memory access, of which some embodiments are effectively rotator free is described herein.

[0055] According to exemplary embodiments, described in further detail in reference to specific figures discussed below, the memory circuit **200** (of FIG. 5) with data storage capability integrated, is accessible in two dimensions, with its controller and the underlying method of storing and loading data into it, supports efficient creation of down-sampled vectors for block based decimation filtering, interpolation and matrix transposition. Through M vector stores in memory circuit **200**, followed by M vector loads, M target vectors are obtained.

[0056] Referring to FIGS. 5 and 6A-6D, it should be noted that FIGS. 6A-6D illustrate alternative embodiments of a rotate and distribute unit **270** shown in FIG. 5.

[0057] FIG. 5 includes a number of SIMD DSP(s) **210**, a rotate and distribute unit **270**. A two-dimensional register file **260**, a read controller **250** and a write controller **240** are provided.

[0058] Vertical alignment of elements occurs at multiples of the decimation factor in memory circuit **200**. There is spreading of the input vector over multiple columns and this flow is repeated per write operation as seen from SIMD DSP **210**.

[0059] The rotation offset is a function of M and P, the software pipeline stage and the write column index. The same holds for the process of combining in the left-over addition unit **230**.

[0060] The method of generating the target vectors for decimation filtering, i.e. vectors of the down-sampled input streams with different start offsets, includes performing a special store operation, which stores samples from a vector in a rotated manner in one or multiple columns of a memory, according to a pattern specific for decimation factor M , the write iteration, P , and the block level software pipeline stage. Thereafter, the desired down-sampled vectors are obtained by indexed reading of one element per column of this two dimensional memory, basically a logical row. The write and read patterns are chosen such that all elements of each desired target vectors are distributed over different columns, so that they can be read out in parallel in a single operation. For this method, in combination with the memory circuit **200**, is that each input sample is only loaded once from an external memory (not shown).

[0061] The memory circuit **200** may support the creation of interpolation output vectors using a method similar to generating down sampled vectors. The main difference in the patterns is how data is stored. It does require P rotators, where for decimation only, $0.5P$ rotators are sufficient for $M > L$.

[0062] The memory circuit **200** may be designed to provide support for parallel load and store operations. This enables an implementation of block based decimation filtering for $M \leq 0.5P$, with software pipelined processing of sequential blocks: while reading out down-sampled vectors from the two dimensional memory directly used for filtering, the M input vectors needed for the next block/stage can be stored into it. The enabler of this pipelining is the method of control and therefore efficient filling of the memory (maximally MP cells are filled per stage, leaving room for $(P-M)P$ samples of other stages).

[0063] Assuming a constant vector width, the memory circuit **200** may be implemented to support mode switching between different number of words per vector P . For example, with a vector size of 256 bits, it can be configured to support 8 words of 32 bit, 16 words of 16 bits or 32 words of 8 bits, etc. The number of rotators is $0.5P$. (here use for decimation-filtering only is assumed with $M > 1$).

[0064] Referring to FIGS. 5 and 6A-6B, the memory circuit **200** of these exemplary embodiment is illustrated which accomplishes decimation. In this exemplary embodiment, only decimation filtering has to be supported with $M \geq 2$ having minimally 2 consecutive data elements. The store operation is done in a single cycle with data elements received from SIMD DSP **210**, with a minimum of $0.5P$ rotate units **220**, illustrated in FIGS. 6A and 6B, being needed since a maximum of $0.5P$ columns have to be filled using data elements from the same input vector.

[0065] Rotate units **220**, shown in FIGS. 6A and 6B, rotate the input vector received from SIMD DSP **210**, each with certain offset.

[0066] When block level software pipelining is required, the implementation may have for each row a write enable signal shared by all cells in the row generated by write controller **240**. Only if both row and column enable signals are positive can a cell be written. The memory controller **200** does block based write enabling: $M \times \text{roundup}(P/M)$ write enabled cells. Write controller **240** generates write enables signals for the rows (M rows are write enabled depending on the software pipeline stage) and columns (max $0.5P$ columns will be written to in case of decimation $M \geq 2$).

[0067] The left over addition units **230**, shown in the embodiment illustrated in FIG. 6A, enable the preservation of

samples from the previous store operation. The left over addition units **230** combine rotated vectors (received from rotate units **220**) of current and previous input vector (stored internally in left over addition unit **230**), such that relevant content of previous input vector is not discarded. However, the presence of the left over addition unit **230** is not required as shown in FIGS. 6B-6D. Therefore, normally left over addition is not required and data is usually passed directly to the two-dimensional register file **260** without further processing by the left over addition unit **230**, as illustrated in FIG. 6B. This processing is further discussed in detail in reference FIGS. 7 and 8. The read controller **250** first selects per column one out of P cells per read port based on the row select signal. Read controller **250** indexes a logical=physical row in the two-dimensional register file **260** that is to be read out. The physical row index is based on the software pipeline stage and index of the decimated target vector. The two-dimensional register file **260** provides the indexed (row) vector to the EVP for further use or storage. Thereafter, a selector unit **265** selects P elements from $P+E$ words and aligns them into the output vector. It should be noted that each output port of the two-dimensional register file **260** may have its own selector unit **265**. The resulting data array is returned to SIMD DSP **210**.

[0068] In this exemplary embodiment the vector width is 256 bits or P words. The memory circuit **200** may be configured for different values of P , i.e. different word widths W : 8 words of 32 bit, 16 words of 16 bits or 32 words of 8 bits. A smaller maximum for P , i.e. larger W , results in a less costly implementation: number of rotate units **220**, columns, etc., scale based on the value of P .

[0069] Regarding the embodiments of the rotate and distribute unit **270** shown in FIG. 6C, the rotate and distribute unit **270** would simply transmit the vectors received SIMD DSP **210** to the two-dimensional register file **260** without further processing.

[0070] Further, as illustrated in FIG. 6D, rather than a number of rotate units **220** provided as in FIGS. 6A and 6B, as single rotate unit may be configured to accomplish the function of giving a common rotation related to the software pipeline stage as that previously discussed in reference to FIG. 6B.

[0071] Referring to FIG. 7, an exemplary method of generating down-sampled target vectors for decimation filtering using the memory circuit **200** is detailed utilizing the rotate and distribute unit **270** shown in FIG. 5, as detailed in the embodiments shown in either FIG. 6A or 6B. In FIG. 7, in step **310**, memory circuit **200** (of FIG. 5) receives an input vector from SIMD DSP **210** (of FIG. 5). The source vectors (columns) **410** are written to two-dimensional register file **260** as vector **500** having subgroups **510** and then read into the target vectors (rows) **420**, illustrated in FIG. 9, are indicated for a use case with $M=3$, $P=8$, $E=0$. The target vectors **420** are identical to the content of the M rows of interest in the memory circuit **200** (the M rows belonging to a certain block level pipeline stage).

[0072] Further, as indicated in FIG. 9 is the data structure **500** contained within the two-dimensional register file **260** shown in FIG. 5. Data elements (subgroups **510**) marked with white boxes are those that the method needs support to prevent overwriting of elements from the previous vectors. This may be done in the left over addition units **230**, illustrated in FIG. 6A, for the rotation based implementation, by combining the rotated input vector **410** of the previous store operation with the rotated input vector of the current one.

[0073] However, in an alternative embodiment, cell-specific write enabling occurs per output of the rotate unit 220, so no longer is the write enable signal shared by all cells in a row. Thus the data of the previous vector is not overwritten.

[0074] For block level pipelining, all rotations may have a common offset, based on the pipeline stage. In this exemplary embodiment, the top M rows are filled, leaving P-M rows for other stages.

[0075] The reading out of target vectors 420 may be done by selecting a row in the two-dimensional register file 260.

[0076] Thereafter, in step 320, input vector 410 is rotated into P+E column-vectors by rotate and distribute unit 220. In step 320, it is further determined per column whether data from the prior input/source vector 410 exists.

[0077] The condition that rotated input vectors of write iterations w and w-1 have to be combined is given by:

```

start_column[w] == end_column[w-1]
where
start_column[w] = floor(P*w/M)
end_column[w-1] = ceil(P*(w-1)/M)

```

[0078] If data from such a prior input/source vector 410 exists, then the rotated vectors are in one possible implementation combined with the previous rotated input vector by left over addition units 230 as shown in FIG. 6A. If left over addition is not required then data is passed directly to the two-dimensional register file 260 without further processing by the left over addition unit 230.

[0079] The algorithm applied in the left over addition units 230 for combining the rotated input vector with the rotated input vector of the previous write operation is given by:

```

for j: 0..P-1
  combined_vector[j] = mask[j]*rotated_vector[w][j] +
  (mask[j])*rotated_vector[w-1][j]
end

```

[0080] The algorithm for the generation of this cell specific mask vector for column i, under condition that combining is required, is given by:

```

for j:0..P-1
  mask[j] = 0
  if(j > ((P-1) - ((M*i)%P))) mask[j] = 1
end

```

[0081] In the alternative embodiment of the rotate and distribute unit 270, requiring left over addition units 230 shown in FIG. 6A, the write controller 240 would generate row and column enable signals in step 330 from which cell-specific write-enable signals are derived by doing a logical AND operation on the row and column signals specific to the cell. Entire blocks of the two dimensional register file are written. As explained, the content of the previously received input vector is preserved by the left-over-addition units.

[0082] In the alternative embodiment of the rotate and distribute unit 270, not requiring left over addition units 230 shown in FIG. 6B, the write controller 240 would generate cell-specific write-enable masks in step 330 to be applied by

the two-dimensional register file 260. Cells containing elements from the previous input vector will not be write-enabled, therefore they are not overwritten.

[0083] Thereafter, in step 340 words of the P+E column vectors are stored in the two-dimensional register file 260 according the cell specific mask created in step 330. The spreading process is executed at step 340 in which rotated and optionally combined vectors are stored in columns of the two-dimensional register file 260. The different shading/hashing used in FIGS. 9-11 are used to indicate the source vector of a data element.

[0084] As illustrated in FIG. 9, in step 320, the filling of M rows in the memory by rotating the M input vectors (as E=0) in a target column specific distance. Groups of (maximally) M consecutive elements from input vector are stored in different columns in one cycle. The rotation per group of elements is target column specific.

[0085] Regarding FIG. 9, the operation of left-over addition 230 is illustrated in which the input vector of the first write iteration (0, 1, . . . 7) is received by the rotator belonging to the third column. It is rotated with an offset corresponding to the third column, M=3 and the first write iteration: (6,7,0,1,2,3,4,5). Assuming row specific write enabling, only the elements of in the first 3 rows are written in this column: (6,7,0). The left over addition unit 230 stores the rotated input vector internally. The second input vector (8,9,10,11,12,13,14,15) is received by the rotator belonging to the third column. It receives a rotation corresponding to the third column, M=3 and second write operation: (14,15,8,9,10,11,12,13). The rotated versions of current and previous input vector are combined such that relevant data of the previous input vector is not overwritten: (6,7,8,X,X,X,X,X). Elements of this resulting vector 510 are stored into the third column of the two-dimensional register file 260, in the first M=3 write enabled rows, (6,7,8).

[0086] The mapping of input data elements of M input vectors into the two-dimensional register file 260 with E configured as 0, giving the column mapping function and (row) address function used during writing, is given by:

```

for i: 0..M-1
  for j: 0..P-1
    row_index = ((i*P+j)%M + offset[s])%P
    column_index = floor((i*P+j)/M)
    register_file[row_index][column_index] =
    input_vector[s][i][j]
  end
end

```

[0087] With regards to decimation filtering, the M rows containing the target vectors belonging to the software pipeline iteration of interest, relate to the input vectors similar as the algorithm that was explained in the background related to FIG. 3. In addition, there can be an offset applied to map the M target vectors in a free range of rows, for example depending on the software pipeline stage "s".

```

for i: 0..M-1
  for j: 0..P-1
    row_index = (i + offset[s])%P
    register_file[row_index][j] = target_vector[s][i][j]
  end
end

```

[0088] Paired with a rotate-and-distribute unit **270** of FIGS. **5** and **6A**, the write controller circuit **240** does simple block based write mask enabling (i.e. enable writing in R rows×C columns), as the rotate-and-distribute unit does combining and merging with the rotated vector of previous operation (in left-over addition units), which is required to keep its useful elements.

[0089] Referring to FIG. **8**, the read function of the read controller **250** will be discussed in detail. In step **350**, the read controller **250** generates P+E column-specific row indices. Thereafter, in step **360**, from each column of the two-dimensional register file **260** a word is read. Further, the selector unit **265** selects P elements from P+E column outputs in step **365** and aligns them in the correct position of the output vector. Finally, a SIMD DSP **210** receives a target vector from the memory circuit **200** in step **370**.

[0090] The method of generating target vector for interpolation will now be discussed in reference to FIGS. **7**, **8** and **10**. Input vectors **600** are used to generate L target vectors **620** in target table **610**. Each of these L input vectors **600** is a down-sampled version of the desired output signal with a different start offset. Using the memory circuit **200**, these L input vectors **600** may be interleaved, i.e. L reordered target vectors **620** of the output signal are generated. The elements of each input vector **600** are rotated according to the target vector row and written in the appropriate columns with distance L apart. After storing L input vectors **600**, L target vectors **620** may be read. It should be noted that the write patterns are such that from each column one element is read during read step **360**.

[0091] If interpolation has to be supported, in addition to decimation, P columns are written into during the store operation, implying that P rotate units **220** are required.

[0092] To support interpolation, in addition to decimation, a rotate-and-distribute unit, for example according to FIG. **6B** with P rotators, in combination with cell specific write control is required. Therefore, FIG. **10** illustrates input vectors **600** for the memory circuit **200**, which are down-sampled versions of the desired interpolation output signal each with different start offset. The data elements of each input vector **600** are rotated according to the target vector row and written in the appropriate columns **620** with distance L apart as target vectors **610**. After storing L input vectors **600**, L reordered target vectors **610** can be read out to the SIMD DSP **210**.

[0093] The mapping into the register file (with E=0) from the input vector perspective (column mapping function, row/cell addressing function) is given by the following algorithm (where i represents the input vector index and j represents the element index):

```

for i: 0..L-1
  for j: 0..P-1
    row_index = floor((j*L+i)/P)
    column_index = (j*L+i)%P
    register_file[row_index][column_index] = input_vector[i][j]
  end
end
end

```

[0094] FIG. **11** illustrates the variable word width feature of the exemplary embodiment shown in FIGS. **3-7**. FIG. **8** illustrates a physical register file **800** that can store 16×16 physical elements (N.B. E=0) of a certain width, say W_{min} . In this example, the number of logical words per vector is 8, where the word size W is $2 * W_{min}$. The register file **800** is concep-

tually split into 8×8 squares that can hold one logical word. FIG. **11**, illustrates the mapping for the decimation case with M=5 (N.B mapping specific to embodiment 1). The two elements that form a logical word are mapped diagonally into the square allocated to the logical word.

[0095] There are four significant benefits generated through the use of the exemplary embodiment shown in FIGS. **5-11**. First, flexible implementations of decimation filtering, interpolation, matrix transposition and similar algorithms on a processor elements **100** (i.e., vector DSP) are enabled. Second, using circuit memory **200**, down-sampled target vectors for decimation filtering can be efficiently generated in which only M unique input vectors are loaded for generating M target vectors. Similar output vectors for interpolation, matrix transposition and similar algorithms are also efficiently generated. Essential to this method is that each input sample is only loaded once per block from another external memory (not shown). Third, with the support for parallel load and stores provided by memory circuit **200**, block level software pipelining is enabled for both decimation and interpolation. Fourth, memory circuit **200** supports switching between a different number of elements P, with a certain P_{max} for both decimation and interpolation.

[0096] The memory circuit **200** according to other exemplary embodiments of the rotate-and-distribute and controller units shown in FIG. **6C-6D** will now be discussed in detail. The write controller **240** provides (in the general case) write-enable signals for P elements per column based on a pattern. The pattern can be either generated based on a user provided configuration and/or internal state specifying decimation factor M, P, the block-level pipelining stage and write iteration number (see Goossens for additional information). Alternatively, the pattern may be provided as an index/offset vector, i.e. for each source element, the target location/displacement in its row is specified. When only decimation is required and a pattern generator is used, the write controller only has to provide 0.5P write mask signals, each for P cells, (N.B. because of $M \Rightarrow 2$, the mask generators can be shared by two or more columns) and P+E column enable signals. To support interpolation, P write mask signals are needed (or P+E masks if there is no sharing). In this embodiment (FIG. **6D**), for the write phase an input vector rotator **220** is used to support block level software pipelining.

[0097] Unlike the exemplary embodiments related to FIGS. **6A-6B**, no rotate units **220** are required in the rotate-and-distribute unit **270** of this exemplary embodiment shown in FIGS. **12-15**. With the elimination of the rotate units **220** a further cost savings may be realized in the exemplary embodiment.

[0098] In the exemplary embodiment shown in FIG. **9**, as illustrated in FIG. **5**, the vector width is 256 bits or P words. The memory circuit **200** may be configured for different values of P, i.e. different word widths W: 8 words of 32 bit, 16 words of 16 bits or 32 words of 8 bits. A smaller maximum for P, i.e. larger W, results in a less costly implementation of the present invention, as the number of columns, etc. change in scale according to the value of P.

[0099] An input vector rotator **220** (when software pipelining is needed) receives the input vector from the SIMD DSP **210**. It rotates the input sample vector and writes the rotated vector to all columns of the two-dimensional register file **260**. The column cells where the data is actually written depends on the write masks and column enable signals.

[0100] An example of the rotation for software pipelining is provided for stage **1** the first M elements are written (column-wise) starting from the top-left element of FIG. **13**. The second set of M elements is written in the second column starting in row M . The next stage starts to be written in the $M+1$ th element of the first row, i.e. a rotation of the input vector of M . In the second row the second set of M elements of the second vector gets a rotation of $2M$ (with wrap around of the data elements in the vector) before it is written. This vertical rotation/displacement is applied to all elements of this stage. Alternative to rotating the input to enable software pipelining, this support may also be implemented with rotators in the read port.

[0101] The write controller **240** generates per column, per element a write enable mask. These masks support writing with a certain distributed access pattern based on M , P , the software pipeline stage, the column index and the write iteration. For decimation, masks can be shared between two (or multiple if $E>0$) columns. Then there are additional column enable signals (equal to the number of columns) that enable one of the columns that share a mask.

[0102] The read controller **250** generates per read port per column a cell index, i.e. physical row index, which is based on the target vector index and software pipeline stage. This element is returned by the read circuit of the two-dimensional register file **260**. From $P+E$ columns P elements are selected and aligned in the output vector. The combined P elements are provided/output to the EVP.

[0103] The method of generating down-sampled target vectors using the memory circuit **200** of the exemplary embodiment shown in FIG. **5**, with rotate-and-distribute units according to FIG. **6C** or **6D**, will now be discussed in detail with reference to FIGS. **3**, **4** and **12-14**.

[0104] FIGS. **7** and **8** illustrate the flowchart detailing the operation of the memory **200** shown in FIG. **5**. At the outset, memory circuit **200** receives input vector from SIMD DSP **210** in step **310**. Thereafter, in step **320** the input vector is rotated by the sample rotator **230** of the rotate-and-distribute unit (a circuit prepared for software pipelining is considered in this case, see FIG. **6D**). In step **330**, cell specific write enable signals are generated by the write controller. In step **340**, elements of each rotated vector are stored in the two-dimensional register file **260** stored in a cell having the same row index as the input element, with a maximum of M consecutive elements per column.

[0105] FIG. **12** illustrates the source vectors (columns) **410** and target vectors **420** used in an example case with $M=3$, $P=8$, $E=0$. The target vectors **420** are stored into the two-dimensional register file **260** according to a certain pattern based on the decimation factor M , the write iteration, and the software pipeline stage. In Goossens, a mathematical description of the aforementioned patterns may be found. The write controller **240** provides cell-specific write enable signals according to the desired write pattern. It should be noted that in this exemplary embodiment, the source vectors **410** do not need a different rotation per-column of the two-dimensional register file as in previous embodiments. In FIG. **12**, an example pattern is shown for $M=3$.

[0106] The mapping/writing into the register file from the input vector perspective (column mapping function, row/cell addressing function) is given by the following algorithm. Here i represents the input vector index (up to M as $E=0$), j the element index.

```

for i: 0..M-1
  for j: 0..P-1
    row_index = j
    column_index = floor((i*P+j)/M)
    register_file[row_index][column_index] = input_vector[i][j]
  end
end

```

[0107] To extract the target vectors **420**, a pattern with one cell index per column is used, which is based on the desired virtual/logical row vector. In FIG. **11**, data elements in bold text (0, 3, 6, 9, 12, 15, 18 and 21) are the data elements belonging to the first target vector **420**.

[0108] The read control for reading out of target vectors from the two-dimensional register file (with $E=0$) is given by the following algorithm,

```

for i: 0..M-1
  for j: 0..P-1
    row_index = (j*M+i)%P
    column_index = j
    target_vector[i][j] = register_file[row_index][column_index]
  end
end

```

[0109] Referring to FIGS. **6C** and **13**, generating target vectors for symmetric decimation filtering may be accomplished as described below. The memory circuit **200** and control supports writing into $P+E$ columns, with $E>0$. There are one or two read ports, as illustrated in FIG. **5**, that may select P elements out of the $P+E$ columns from the input vectors **600** and do simple shift for alignment. Referring to FIG. **13**, the method of generating $M(1+E)$ target vectors **620** using the two-dimensional register file **260** with E additional columns is filled column-wise in $\text{ceil}(M+M^*E/P)$ steps, with input vectors, whose relative offset is P samples. The content of two dimensional register file then contains $M(1+E)$ unique target vectors **620** of length P . Reading out a vector of P samples is based on for example the desired target vector index given M , P and E .

[0110] In FIG. **13**, target vectors **620** are stored in the two-dimensional register file **260** by the rotate and distribute unit **270** shown in FIG. **6C**. Each block of M target vectors are the time-shifted versions (time step equals M) of the down-sampled sequences of the first block of M target vectors, having each 1 new element in the last column. So each additional element in the extra columns adds another target vector. Per column there are M extra elements. The number of target vectors is therefore the original M base vectors+ M^*E extra target vectors.

[0111] Using this memory, downsampled target vectors for decimation filtering and especially symmetric decimation filtering can be efficiently generated. $\text{Ceil}(M+M^*E/P)$ input vectors are stored into this memory and allow the generation of $M(1+E)$ downsampled vectors. With a basic 2D memory, P vectors are stored to generate maximally P downsampled vectors. With the invented memory circuits having rotate-and-distributed circuits according to FIGS. **6A-D** and multi-column distribution of input samples, M vectors are stored to generate maximally M downsampled vectors. With this invention applied in the discussed exemplary embodiment with simple selectors in the read ports, the overhead of mak-

ing M base vectors is identical to the unextended memory circuits, but the overhead and resource utilization of making the following $M \cdot E$ vectors can be reduced, by doing another $\text{ceil}(M \cdot E / P)$ stores into the two-dimensional register file **260**.

[0112] Symmetric filtering for this embodiment can be supported efficiently up to $K = M(1+E)$. Exploiting symmetries means the calculation of two filter coefficients per multiply-accumulation operation.

[0113] With two read ports (illustrated in FIG. 5) an efficient schedule for decimation filtering on a VLIW vector DSP (with support for parallel addition and multiply-accumulate operations) is enabled. This means in the key loop of the decimation filter (maximally) one multiply-accumulation operation per cycle, i.e. maximally **2** filter coefficients per cycle.

[0114] The processing involved in the symmetric decimation filtering, for an exemplary embodiment with simple selectors in the read ports, as illustrated in FIG. 13 will be described by the writing and reading algorithms. The mapping/writing into the register file from the input vector perspective (column mapping function, row/cell addressing function) is given by the following algorithm. Here i represents the input vector index, j the element index. Note that the algorithm is identical to the earlier one for $E=0$, except that the iteration bound is not $M-1$ but $\text{ceil}(M + M \cdot E / P) - 1$, as additional input vectors can be stored in the extra columns.

```

for i: 0..ceil(M+M*E/P)-1
  for j: 0..P-1
    row_index = j
    column_index = floor((i*P+j)/M)
    register_file[row_index][column_index] = input_vector[i][j]
  end
end

```

[0115] The read control for reading out of $M(1+E)$ target vectors from the two-dimensional register file with $E > 0$ and simple selectors in the read ports is given by the following algorithm,

```

for i: 0..M(1+E)-1
  for j: 0..P-1
    row_index = (j*M+i)%P
    column_index = j + floor(i/M)
    target_vector[i][j] = register_file[row_index][column_index]
  end
end

```

The difference in the algorithm for $E > 0$ relative to $E = 0$ that more target vectors can be read, changing the loop iteration bound to $M(1+E) - 1$. Further the column index, increments by one for every next set of M target vectors.

[0116] An exemplary method of supporting block level software pipeline stages with this memory circuit **200** the data structure shown in FIG. 14 will now be discussed. In this aspect of this exemplary embodiment, to enable block level software pipelining, a (shared) single rotate unit **220** (according to the embodiment illustrated in FIG. 6D) gives a rotation to align the target vectors of this block to M unused virtual rows. Alternatively, it is possible to store another block level pipeline stage with an offset in the column direction, instead of giving a rotation in the row direction, so no input vector

rotation is required. However, this alternate approach does require selection (with rotation) support in the read port (not shown).

[0117] Therefore, FIG. 14 illustrates two input source vectors **1010** for the next block ($P=8$, $E=0$). Pre-rotation of the input source vectors is done by rotate unit **220** of the rotate-and-distribute unit to give an offset of target vectors **1020** of M rows. Thus, filling of an array in the two-dimensional register file **260** according to pattern specific for M , using write iteration and block level software pipeline stage.

[0118] FIG. 14 extends FIG. 12 for support of block-level software pipelining. A software pipeline iteration specific rotation is applied to the logical row index to get the physical row index for reading, as follows (for $E=0$):

$$\text{row_index} = (j * M + i + \text{offset}[s]) \% P$$

where $\text{offset}[s]$ for example equals $\text{offset}[s] = s * M$.

[0119] Referring to FIG. 15, the method of interleaving vectors for interpolation is an extension of the previously discussed method of generating down sampled vectors for decimation filtering, previously discussed.

[0120] In FIG. 15, input vectors for the memory circuit, which are down-sampled versions of the desired interpolation output signal each with different start offset are illustrated. The data elements of each input vector **600** are written in a strided manner into the appropriate columns with a distance L apart and stored in two-dimensional register file **260** (configured as $E=0$). After storing L input vectors **600**, L reordered target vectors **620** may be read out.

[0121] The algorithm for storing into the two-dimensional register file **260** (column mapping function, row/cell addressing function), as seen from the input vector perspective is as follows:

```

for i: 0..L-1
  for j: 0..P-1
    row_index = j
    column_index = (j*L+i)%P
    register_file[row_index][column_index] = input_vector[i][j]
  end
end

```

[0122] The algorithm for reading target vectors from the two-dimensional register file **260** (configured as $E=0$) is as follows:

```

for i: 0..L-1
  for j: 0..P-1
    row_index = floor((i*P+j)/L)
    column_index = j
    target_vector[i][j] = register_file[row_index][column_index]
  end
end

```

[0123] Referring to FIG. 16, the data structure shown illustrates switching between number of elements P per vector, i.e. switching between element word width W . For a chosen P_{max} (or W_{min}), the W_{min} is considered as the granularity to store the data elements for any P ($2nP = P_{max}$). Words of size W are divided in sub-elements of size W_{min} and stored in a diagonal manner. It is noted that the memory circuit **200** is designed so that all sub-elements that have to be read out into the same target vector **1210** are stored in different columns. Further, the

relative offset between the sub-elements in the column direction is not changed during the write operation. In FIG. 16, $M=5$, $E=0$, $P_{max}=16$, and $P=8$.

[0124] FIG. 15 discusses the support of a variable number of words P in a vector, with a certain maximum P_{max} , where $nP=P_{max}$ and n is even. So a logical vector of P elements can be partitioned into P sets of n sub-elements. The mapping algorithms as discussed in relation to FIG. 15 apply to logical vectors of length P into a logical two-dimensional array of $P \times P$. The mapping of each cell element (having n sub-elements) of the logical array on the physical array of $P_{max} \times P_{max}$ is given as follows:

```

for i: 0..P-1 //logical row index
  for j: 0..P-1 //logical column index
    for e: 0..n //subelement index
      physical_register_file[*n+e][j*n+e] =
        logical_register_file[i][j][e]
    end
  end
end
end

```

[0125] The benefits provided by the embodiment of the present invention with a rotate-and-distribute unit according to FIG. 6C or 6D with mappings shown in FIGS. 12-15 include those provided by the exemplary embodiment with a rotate-and-distribute unit according to FIG. 6A or 6B as discussed related to FIGS. 9-11. Further, this exemplary embodiment is more cost efficient. This exemplary embodiment replaces P (or $0.5P$ if only decimation has to be supported) vector rotate units 220 in the rotate-and-distribute unit 270 by cell specific write control. However, this exemplary embodiment may require a single sample rotator 220 depending on the chosen manner of supporting block level software pipelining. This exemplary embodiment with generated patterns avoids the overhead of calculating them (in the typically resource limited filtering phase). Further, this exemplary embodiment also enables a user-provided write/read patterns allowing for more flexible memory access schemes.

[0126] Therefore, exemplary embodiments of the present invention provide for a parallel memory circuit/system (or multi-access memory circuit) supporting strided memory access patterns, that is rotator free in certain exemplary embodiments. The embodiments of the present invention have a particular memory bank assignment function/strategy (in which columns to write/read) and they have a particular addressing function (in/from which row=element to write/read; or how to rotate). In addition, the embodiments of the present invention support various word widths. Therefore, the embodiments of the present invention provide and for lower cost implementation.

[0127] While the present invention has been described in connection with certain exemplary embodiments, it is to be understood that the invention is not limited to the disclosed embodiments, but, on the contrary, is intended to cover various modifications and equivalent arrangements included within the spirit and scope of the appended claims, and equivalents thereof.

What is claimed is:

1. A configurable memory circuit for a Single Instruction, Multiple Data (SIMD) Digital Signal Processor (DSP), for generating reordered output vectors from sequential input data vectors, comprising:

a rotate-and-distribute unit to rotate and distribute an input data vector;

a two-dimensional register file to receive and store vectors generated by the rotate-and-distribute unit;

a write controller to control the mapping of input data elements in the two-dimensional register file, corresponding to a mapping configuration by configuring the rotate-and-distribute unit and generating write-enable signals for all cells of the two-dimensional register file; and

a read controller to select on a per column basis one data element to be received in the output vector, corresponding to the mapping configuration.

2. The memory circuit of claim 1, wherein the rotate-and-distribute unit distributes the rotated input vector to all columns of the two-dimensional register file.

3. The memory circuit of claim 2, wherein the rotate-and-distribute unit further comprises:

a single rotate unit that performs vertical translation of the mapping a pattern in the two-dimensional memory, based on a software pipeline stage.

4. The memory circuit of claim 1, wherein the rotate-and-distribute unit further comprises:

a plurality of rotate units, each corresponding to one or multiple columns of the two-dimensional register file, to rotate the input vector with a column-specific rotation.

5. The memory circuit of claim 4, further comprising:

a plurality of left-over-addition units connected to the plurality of rotate units to preserve data elements from a previous store operation when data elements from a previous store operation exist.

6. The memory circuit of claim 5, wherein the plurality of left over addition units combines a previously rotated input vector with a current rotated input vector.

7. The memory circuit of claim 5, wherein the write controller generates row and column enable signals, from which the cell-enable signals are derived by doing a logical AND operation on the row and column signals specific to the cell.

8. The memory circuit of claim 4, wherein the read controller selects a row in the two-dimensional register file, to be received in the output vector.

9. The memory circuit of claim 1, wherein the data elements are of a same word length.

10. The memory circuit of claim 1, wherein the data elements are of differing word length.

11. The memory circuit of claim 1, wherein the input vectors are of a constant vector length.

12. The memory circuit of claim 1, wherein the read and write controllers have state-machines for generating their respective control signals, based on parameters received from the SIMD DSP.

13. The memory circuit of claim 1, wherein the read and write controllers translate mapping configurations patterns received from the SIMD DSP to the respective control patterns.

14. The memory circuit of claim 13, wherein the mapping configuration patterns are preloaded into a configuration memory.

15. The memory circuit of claim 6, wherein the read controller selects P elements from $P+E$ column outputs stored in the two-dimensional register file and aligns them in the output vector.

16. The memory circuit of claim 15, wherein the rotate-and-distribute circuits has $P+E$ (logical) column output vec-

tors of which elements are stored in the P+E columns of the two-dimensional register file, according to a predetermined mapping pattern.

17. A method for generating reordered output vectors from sequential input data vectors utilizing a configurable memory circuit for a Single Instruction, Multiple Data (SIMD) Digital Signal Processor (DSP) having a rotate-and-distribute unit, a two-dimensional register file, a write controller and a read controller, comprising:

rotating and distributing an input data vector by the rotate-and-distribute unit generating output vectors;

receiving and storing the target vectors generated by the rotate-and-distribute unit in the two-dimensional register file controlling the mapping of input data elements in the two-dimensional register file corresponding to a mapping configuration by configuring the rotate-and-distribute unit and generating write-enable signals for all cells of the two-dimensional register file; and

selecting on a per column basis one data element to be received in the output vector, corresponding to the mapping configuration by the read controller.

18. The method of claim **17**, wherein the step of rotating and distributing an input data vector by the rotate-and-distribute unit generating output vectors further comprises:

distributing the rotated input vector to all columns of the two-dimensional register file.

19. The method of claim **18**, wherein the step of rotating and distributing an input data vector by the rotate-and-distribute unit generating output vectors further comprises:

vertical translating a mapping of a pattern in the two-dimensional memory, according to a software pipeline stage.

20. The method of claim **17**, wherein the step of rotating and distributing an input data vector by the rotate-and-distribute unit generating output vectors further comprises:

rotating the input vector with a column-specific rotation each corresponding to one or multiple columns of the two-dimensional register file using a plurality of rotate units.

21. The method of claim **20**, further comprising: preserving the data elements from a previous store operation when the data elements from a previous store operation exist using a plurality of left-over-addition units connected to the plurality of rotate units.

22. The method of claim **21**, wherein the plurality of left over addition units combines a previously rotated input vector with a current rotated input vector.

23. The method of claim **22**, further comprising: generating row and column enable signals from which the cell-enable signals are derived by doing a logical AND operation on the row and column signals specific to the cell by the write controller.

24. The method of claim **20**, wherein the read controller selects a row in the two-dimensional register file, to be received in the output vector.

25. The method of claim **17**, wherein the data elements are of a same word length.

26. The method of claim **17**, wherein the data elements are of differing word length.

27. The method of claim **17**, wherein the input vectors are of a constant vector length.

28. The method of claim **17**, wherein the read and write controllers have state-machines for generating their respective control signals, based on parameters received from the SIMD DSP.

29. The method of claim **17**, wherein the read and write controllers translate mapping configurations patterns received from the SIMD DSP to the respective control patterns.

30. The method of claim **29**, where the mapping configuration patterns are preloaded into a configuration memory.

31. The method of claim **22**, wherein the read controller selects P elements from P+E column outputs stored in the two-dimensional register file and aligns them into the output vector.

32. The method of claim **31**, wherein P+E (logical) column output vectors of which elements are stored in the P+E columns of the two-dimensional register file, according to a predetermined mapping pattern.

* * * * *