ESRA

Engeneering & Technology Letters

A Unified FPGA Address Generator for the IEEE 802.16e WiMAX Deinterleaver with Generalised Ncbps Support

Aisha Misbah*, **Geethanand** and **Sundari T****

Department of Electronics and Communication Engineering (2013–2014 Batch)
Atria Institute of Technology, Anandnagar, Bangalore 560024, India
**email: aishamisbah@gmail.com*
**Mentor

ARTICLE INFO

Received 1st September 2026, www.esrapublications.com

Revised & Accepted 29th October 2026

ABSTRACT

The channel deinterleaver is a critical block in the IEEE 802.16e (WiMAX) receive chain, and its area and throughput are dominated by the address generator. A direct hardware mapping of the two permutations defined by the standard needs a floor operator, an integer division and a modulo block, all of which are costly in silicon. This paper presents a single, unified FPGA address generator that serves every permitted combination of modulation order (QPSK, 16-QAM, 64-QAM) and number of coded bits per OFDM symbol (Ncbps). A two-dimensional reformulation of the permutation removes the floor operator, and a divide-by-sixteen front end derives the column count directly from Ncbps, so no configuration-indexed multiplexer is required. The Verilog design was simulated in Xilinx ISim and synthesised for a Spartan-3 XC3S400. It occupies 65 slices (1.8%), 16 flip-flops (0.2%) and 120 four-input LUTs (1.7%) and reaches a maximum clock frequency of 132.39 MHz. This is roughly half the slice and LUT usage of the closest published Spartan-3 design [10] at about 9% higher clock frequency, and a marginal improvement in clock frequency over a closely related design [11].

Index Terms: Address generator, channel deinterleaver, FPGA, IEEE 802.16e, OFDM, WiMAX, Verilog HDL.

1. INTRODUCTION

Worldwide Interoperability for Microwave Access, marketed as WiMAX, is a family of metropolitan-area wireless broadband technologies derived from the IEEE 802.16 standard [1], [2]. The mobile variant, 802.16e [3], specifies an OFDM-based physical layer with a concatenated Reed-Solomon and convolutional channel code [4], a block interleaver whose depth equals the number of coded bits per OFDM symbol (Ncbps), and a choice of QPSK, 16-QAM or 64-QAM constellations [5].

The receive-side deinterleaver is built around two single-port memory banks and an address generator that supplies read and write pointers. The two permutations defined by the standard involve a floor operator, an integer division and a small modulo comparison, whose direct hardware realisation is inefficient in both area and delay. Several prior works address this problem. Lookup-table (LUT) implementations [6] are simple but need large memories and scale poorly with Ncbps. Finite-state-machine generators [7] cover all permitted code rates and modulation schemes without storing full address tables. A VHDL address generator for the interleaver at a single code rate is reported in [9]. Two-dimensional algebraic reformulations [8],[10] express the permutation in closed form and eliminate the divider and the floor operator.

The most direct algorithmic predecessor of the present work is [10], which derives a closed-form two-dimensional mapping for the permitted modulation and Ncbps combinations, shares hardware between the QPSK, 16-QAM and 64-QAM paths, and reports a Spartan-3 implementation at 121.82 MHz with 3.49% slice utilisation. A subsequent amendment [11] replaces the configuration multiplexers of [10] with a divider-and-subtractor front end so that the generator accepts arbitrary interleaving depths, and reports 130.24 MHz. Building on this line of work, the present paper: (i). restates the two-dimensional mapping together with the explicit conditions under which it is exact, and checks it against the standard permutations; (ii). presents a unified datapath in which the row counter, column counter, multiply-by-

Copyright: The authors © 2026. This is an open-access article distributed under Creative Commons Attribution License, which permits unrestricted use, distribution and reproduction in any medium, provided the original author and source are credited.

16 stage and final adder are shared by all three modulation orders and a divide-by-sixteen block derives the column limit directly from Ncbps, so that no manual pre-computation is needed; and (iii). reports an implementation on a Xilinx Spartan-3 XC3S400 that occupies 65 slices and runs at 132.39 MHz, about half the area of [10] at a clock frequency about 9% higher.

The remainder of the paper is organised as follows. Section II summarises the WiMAX deinterleaver structure. Section III derives the two-dimensional form of the permutation for each modulation order and states its validity conditions. Section IV describes the unified hardware architecture. Section V reports simulation and synthesis results and compares them with prior work. Section VI concludes.

2. WIMAX DEINTERLEAVER BACKGROUND

The 802.16e channel deinterleaver is a block deinterleaver of depth Ncbps. It is assembled from two single-port memory banks M1 and M2 and an address generator. A select signal, sel, alternates the two banks between read and write phases every Ncbps cycles. While one bank is loaded with newly demodulated bits under the write address sequence, the other is read out under the deinterleave address sequence into the Viterbi/Reed–Solomon decoder.

The standard [3] admits QPSK and 16-QAM at code rates 1/2 and 3/4, and 64-QAM at code rates 1/2, 2/3 and 3/4, with Ncbps values in the range 96–576 bits, all of which are integer multiples of 96. The permutations (1)–(2) below depend only on the modulation order and on Ncbps. A conformant receiver must support every permitted combination, so the address generator must either be re-synthesised per configuration (impractical) or be fully parameterised. The permutations originate in the fixed-WiMAX specification [12] and are carried over to 802.16e [3].

3. TWO-DIMENSIONAL PERMUTATION ALGORITHM

A. One-dimensional form

Let d = 16 be the number of rows of the block deinterleaver, and let s = max{1, Ncpc/2} be the modulation-dependent parameter, where Ncpc is the number of coded bits per subcarrier. Thus s = 1, 2, 3 for QPSK, 16-QAM and 64-QAM respectively (BPSK also gives s = 1 and is served by the QPSK path). The standard-mandated deinterleaver permutations [3], [12] for the bit at position n ($0 \le n <$ Ncbps) are

$$m_n = s\cdot\lfloor n/s \rfloor + (n + \lfloor d\cdot n/Ncbps \rfloor) \bmod s \tag{1}$$

$$k_n = d\cdot m_n - (Ncbps - 1)\cdot\lfloor d\cdot m_n/Ncbps \rfloor \tag{2}$$

A direct translation of (1)–(2) into hardware needs a divider, a floor unit and a modulo block, which together dominate the area and the critical path.

B. Two-dimensional reformulation

Let C = Ncbps/d be the number of columns and define the column and row indices

$$i = n \bmod C,\ j = \lfloor d\cdot n/Ncbps \rfloor \tag{3}$$

so that $n = j\cdot C + i$, with $0 \le i < C$ and $0 \le j < d$. The deinterleaved address then takes the common form

$$k\,(i, j) = d\cdot r_{i,j} + j \tag{4}$$

where $r_{i,j}$ is an effective column index that captures the intra-row permutation for the selected modulation. Analysing (1)–(2) for every permitted Ncbps gives the following three-case rules.

C. QPSK (s = 1)

No intra-row permutation is required, so

$$k_{QPSK}\,(i, j) = d\cdot i + j \tag{5}$$

D. 16-QAM (s = 2)

$$k_{16}\,(i, j) = d\cdot i + j, \qquad j \bmod 2 = 0$$

$$= d\cdot(i + 1) + j, \quad j \bmod 2 = 1, i \bmod 2 = 0$$
$$= d\cdot(i - 1) + j, \quad j \bmod 2 = 1, i \bmod 2 = 1 \tag{6}$$

E. 64-QAM (s = 3)

$$k_{64}(i, j) = d\cdot i + j, \; j \bmod 3 = 0$$
$$= d\cdot(i - 2) + j, \quad j \bmod 3 = 1, i \bmod 3 = 2$$
$$= d\cdot(i + 1) + j, \quad j \bmod 3 = 1, i \bmod 3 \neq 2$$
$$= d\cdot(i + 2) + j, \quad j \bmod 3 = 2, i \bmod 3 = 0$$
$$= d\cdot(i - 1) + j, \quad j \bmod 3 = 2, i \bmod 3 \neq 0 \tag{7}$$

Each branch of (5)–(7) requires only a multiplication by d = 16, a small add or subtract, and a modulo-2 or modulo-3 comparison. The floor operator has disappeared from the datapath.

F. Derivation and validity conditions

Substituting $n = j\cdot C + i$ gives $d\cdot n = j\cdot Ncbps + d\cdot i$, and since $d\cdot i < Ncbps$, $\lfloor d\cdot n/Ncbps \rfloor = j$, which confirms (3). For s = 1, (1) gives $m_n = n$, and (2) reduces to $d\cdot i + j$, which is (5). For s = 2 and s = 3, (1) moves n within its aligned group of s consecutive indices, $m_n = n - (n \bmod s) + ((n \bmod s) + j) \bmod s$. If s divides C, then $n \bmod s = i \bmod s$, so the group lies entirely within row j and $m_n = j\cdot C + i'$ with $0 \leq i' < C$. Here i' is i, $i \pm 1$ or $i \pm 2$ as listed in (6)–(7). In this case $\lfloor d\cdot m_n/Ncbps \rfloor$ is still j, and (2) yields $d\cdot i' + j$.

The closed forms (5)–(7) are therefore exact if and only if Ncbps is a multiple of 16 and C is a multiple of s; that is, Ncbps is a multiple of 32 for 16-QAM and of 48 for 64-QAM. This holds for every Ncbps permitted by the standard. It does not hold for arbitrary depths; for example, (6) fails for Ncbps = 48. An exhaustive software comparison of (5)–(7) against (1)–(2) for every Ncbps that is a multiple of 96 up to 2304 found no mismatches. The "generalised Ncbps support" of this work thus means any depth satisfying these divisibility conditions, within the datapath width.

4. UNIFIED HARDWARE ARCHITECTURE

Equations (5)–(7) share the same backbone: a row counter, a column counter, a multiply-by-d stage and a final adder. The QPSK path multiplies the column index i; the 16-QAM path multiplies one of {i, i+1, i−1}; the 64-QAM path multiplies one of {i, i+1, i+2, i−1, i−2}. A single adder/subtractor pair computing $i \pm 1$ serves both 16-QAM and 64-QAM, so one set of operators services all three modes. A two-bit modulation-select signal drives a 3:1 multiplexer that selects the correct branch of (5)–(7); for 16-QAM and 64-QAM the branch selection is further refined by modulo-2 or modulo-3 comparisons on i and j (the modulo-2 test is simply the least-significant bit).

A second refinement concerns the column count C. In [10] the column count is supplied from an external configuration word through a multiplexer with one input per permitted Ncbps; [11] replaces such multiplexers with a divider and a subtractor. The present design follows the same principle within the unified datapath. Because C = Ncbps/d and d = 16 is a power of two, the column limit $C - 1$ is obtained from Ncbps by a right shift of four bits (pure wiring) followed by a decrement. This replaces the multiplexer and its decoder, removes the configuration-decode layer from the column-count path, and lets the generator accept Ncbps directly, so no pre-computation or stored configuration word is needed.

The complete unified generator therefore comprises: (i). a multiply-by-d stage; (ii). a common adder and subtractor used to compute $i \pm 1$; (iii). a dedicated (i + 2) adder and (i − 2) subtractor used only by the 64-QAM path; (iv). modulo-2 and modulo-3 comparators on i and j; (v). the branch-selection multiplexers; (vi). a row counter and a column counter driven by the derived column limit; and (vii). a divide-by-sixteen unit that derives $C - 1$ from Ncbps. All arithmetic is unsigned. Since $Ncbps \leq 576$ and the largest address is 575, a 10-bit datapath suffices. The subtractions are selected only when $i \geq 1$ ($i \geq 2$ for $i - 2$), so no negative intermediate value occurs.

5. RESULTS

A. Functional simulation

The Verilog RTL was simulated in Xilinx ISim for every permitted combination of modulation order and Ncbps. Reference address sequences were pre-computed in Python from the one-dimensional equations (1)–(2) and compared cycle by cycle against the DUT output. All configurations produced bit-exact matches over a full period of Ncbps addresses.

B. FPGA synthesis

The design was synthesised and placed-and-routed for the Xilinx Spartan-3 XC3S400, speed grade −5, using Xilinx ISE 14.7. Post-implementation utilisation is given in Table I; device capacities are taken from the data sheet [13].

TABLE I. DEVICE UTILISATION ON XILINX SPARTAN-3 XC3S400

Resource	Used	Available	Utilisation (%)
Slices	65	3584	1.81
Slice flip-flops	16	7168	0.22
4-input LUTs	120	7168	1.67
Bonded IOBs	27	141	19.1
MULT18×18	1	16	6.3
GCLKs	1	8	12.5

Percentages are computed as used/available; the synthesis tool itself reports truncated integer percentages (1%, 0%, 1%, …).

The maximum clock frequency reported by the static timing analyser after place-and-route is 132.39 MHz. At one address per clock this corresponds to 132.39 million addresses per second.

C. Comparison with prior work

Table II compares the present design with the base design [10], the amended design [11] and a LUT-based generator, whose figures are those reported alongside [10]. [10] and the LUT-based generator were synthesised on the same Spartan-3 XC3S400 part. [11] does not state its target device explicitly (it reports ISE 12.3), so the comparison with it is indicative only, and its slice and LUT figures are printed as integer percentages.

TABLE II. COMPARISON WITH PRIOR ART ON SPARTAN-3 XC3S400

Metric	This work	[10]	[11]	LUT-based
Slices (%)	1.81	3.49	1 (integer)	17.66
Flip-flops (%)	0.22	0.50	0.153	0.78
4-input LUTs (%)	1.67	3.35	1 (integer)	17.15
f_max (MHz)	132.39	121.82	130.24	62.51

Relative to [10], the unified design uses about half the slices (1.81% against 3.49%) and LUTs (1.67% against 3.35%), and its clock frequency is 8.7% higher (132.39 MHz against 121.82 MHz). Relative to the LUT-based generator it uses about ten times fewer slices (17.66% against 1.81%) and runs more

than twice as fast. Against [11], the clock frequency is 1.7% higher (132.39 MHz against 130.24 MHz), while [11] reports fewer flip-flops (0.153% against 0.22%); its integer-rounded slice and LUT figures cannot resolve any difference from the present design.

Two factors plausibly account for the gain over [10]: the three modulation-specific datapaths share their common operators, so the effective cost per modulation is lower than for independent instantiations; and replacing the column-count multiplexer with a divide-by-sixteen block removes the configuration-decode layer from the critical path. No per-factor ablation was performed, so the individual contributions are not quantified. The single MULT18×18 block in Table I is noteworthy, since multiplication by d = 16 is a four-bit shift and could in principle be realised without an embedded multiplier.

6. CONCLUSIONS

A unified, area- and delay-efficient FPGA address generator for the IEEE 802.16e WiMAX deinterleaver has been presented. A two-dimensional reformulation of the standard-mandated permutation eliminates the floor operator from the datapath, and its exactness conditions have been stated and checked. Structural sharing across the three modulation orders compresses the hardware footprint, and a divide-by-sixteen front end removes the need for a configuration-indexed multiplexer at the column-count stage. Implemented on a Xilinx Spartan-3 XC3S400 device, the design occupies 65 slices and 120 LUTs and runs at 132.39 MHz. This is about half the area and 9% faster than [10], well ahead of a LUT-based generator, and 1.7% faster than the closely related design [11]. Future work will extend the unified architecture to the LTE sub-block interleaver and characterise the design in a 40 nm ASIC flow.

Declaration

AI-assisted tools: AI-assisted tools have been used to support organization and language refinement. The authors have independently verified all claims, citations, originality and final text.

REFERENCES

[1] W. Konhauser, “Broadband wireless access solutions-Progressive challenges and potential value of next generation,” *Wireless Pers. Commun*., vol. 37, no. 3–4, pp. 243–259, May **2006**.

[2] J. G. Andrews, A. Ghosh, and R. Muhamed, Fundamentals of WiMAX: Understanding Broadband Wireless Networking. Upper Saddle River, NJ, USA: Prentice Hall, **2007**.

[3] IEEE Standard for Local and Metropolitan Area Networks-Part 16: Air Interface for Fixed and Mobile Broadband Wireless Access Systems-Amendment 2: Physical and Medium Access Control Layers for Combined Fixed and Mobile Operation in Licensed Bands, and Corrigendum 1, IEEE Std 802.16e-2005 and IEEE Std 802.16-2004/Cor 1-2005, Feb. **2006**.

[4] S. Lin and D. J. Costello, Jr., Error Control Coding, 2nd ed. Upper Saddle River, NJ, USA: Prentice Hall, **2004**.

[5] R. Prasad, OFDM for Wireless Communications Systems. Boston, MA, USA: Artech House, **2004**.

[6] Y.-N. Chang and Y.-C. Ding, “A low-cost dual-mode deinterleaver design,” in Proc. IEEE Int. Conf. Consumer Electronics (ICCE), Las Vegas, NV, USA, Jan. **2007**, pp. 1–2.

[7] B. K. Upadhyaya, I. S. Misra, and S. K. Sanyal, “Novel design of address generator for WiMAX multimode interleaver using FPGA based finite state machine,” in Proc. 13th Int. Conf. Computer and Information Technology (ICCIT), Dec. **2010**, pp. 153–158.

[8] R. Asghar and D. Liu, “2D realization of WiMAX channel interleaver for efficient hardware implementation,” in Proc. World Acad. Sci. Eng. Technol., Hong Kong, **2009**, vol. 51, pp. 25–29.

[9] A. A. Khater, M. M. Khairy, and S. E.-D. Habib, “Efficient FPGA implementation for the IEEE 802.16e interleaver,” in Proc. Int. Conf. Microelectronics (ICM), Marrakech, Morocco, Dec. **2009**, pp. 181–184.

[10] B. K. Upadhyaya and S. K. Sanyal, “Efficient FPGA implementation of address generator for WiMAX deinterleaver,” IEEE Trans. Circuits Syst. II, Exp. Briefs, vol. 60, no. 8, pp. 492–496, Aug. **2013**.

[11] O. Rafique and S. L. Gangadharaiah, “Hardware efficient WiMAX deinterleaver capable of address generation for random interleaving depths,” arXiv:1404.5885, Apr. **2014**.

[12] IEEE Standard for Local and Metropolitan Area Networks-Part 16: Air Interface for Fixed Broadband Wireless Access Systems, IEEE Std 802.16-2004, Oct. **2004**.

[13] Spartan-3 FPGA Family Data Sheet, DS099 (v3.1), Xilinx, Inc., San Jose, CA, USA, Jun. **2013**.