# Scalable Data Concentrator with Baseline Interconnection Network for Triggerless Data Acquisition Systems 

Wojciech Marek Zabołotny (

Institute of Electronic Systems, Faculty of Electronics and Information Technology, Warsaw University of Technology, Nowowiejska 15/19, 00-650 Warszawa, Poland; wojciech.zabolotny@pw.edu.pl


#### Abstract

Triggerless data acquisition systems (DAQs) require that the data stream is transmitted from multiple links and into the processing node. The short input data words must be concentrated and packed into the longer bit vectors the output interface (e.g., PCI Express) uses. In that process, the unneeded data must be eliminated, and a dense stream of useful DAQ data must be created. Additionally, the time order of the data should be preserved. This paper presents a new solution using the Baseline Network with Reversed Outputs (BNRO) for high-speed data routing. A thorough analysis of the network's operation enabled increased scalability compared to the previously published concentrator based on an $8 \times 8$ network. The solution may be scaled by adding additional layers to the BNRO network while minimizing resource consumption. Simulations were performed for four and five layers (16 and 32 inputs). The FPGA implementation and tests in the actual hardware were successfully performed for 16 inputs. The pipeline registers may be added in each layer independently, shortening the critical path and increasing the maximum acceptable clock frequency.


Keywords: FPGA; DAQ; data concentration; multistage interconnection networks; baseline network

Citation: Zabołotny, W.M. Scalable Data Concentrator with Baseline Interconnection Network for Triggerless Data Acquisition Systems. Electronics 2024, 13, 81. https:// doi.org/10.3390/electronics13010081

Academic Editor: Alexander Barkalov
Received: 14 November 2023
Revised: 17 December 2023
Accepted: 20 December 2023
Published: 23 December 2023


Copyright: © 2023 by the author. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https:// creativecommons.org/licenses/by/ 4.0/).

## 1. Introduction

The high-energy physics experiments generate vast volumes of data. It is not possible to archive all that data for later offline processing. Two approaches are used to solve that issue. In the triggered data acquisition systems, the incoming data are buffered until the "level one (L1) trigger" decides whether it is useful. The data must be buffered until this decision is taken and delivered. The L1 trigger decision must be elaborated quickly to keep the necessary capacity of data buffers reasonable, which usually requires high-performance FPGA chips. When elaborating the local trigger decision and waiting for the L1 trigger decision, the data may be prepared for sending to the data acquisition system (DAQ).

However, using the triggered approach is not always possible. The papers [1,2] predicted the increasing role of triggerless DAQs eight years ago. Some experiments already use or plan to use this approach: LHCb [3], AMBER [4], and CBM [5]. In that approach, the raw data from the detector are submitted only to elementary processing (e.g., zero suppression) and then delivered to the system responsible for identifying the interesting events-the event selector. The event selector may be built based on ATCA architecture [6] or as a tightly connected computer network [7,8]. The second solution may be better scalable and upgradable, using the standard servers as the main component. In that solution, it is still necessary to deliver the data from the front-end electronics (FEE) via the optical links to the computer being the entry nodes of the event building network. That may be accomplished with specialized FPGA-based data concentrator boards connected to the PCIe bus in the server $[9,10]$ and equipped with multiple optical transceivers.

The PCIe blocks in modern FPGA chips use a wide datapath to fully utilize the bandwidth offered by the PCIe interface at reasonable clock frequency. For example, the PCIe Gen3 block working with 8 lanes requires delivering 256 bits of data at a frequency of 250 MHz . For 16 lanes, the data should be delivered as 512-bit words at 250 MHz [11].

The overall volume of data should be limited. Therefore, the data about the physics events recorded by detectors (so-called "hit messages") are transmitted as possibly short data words. For example, in the CBM experiment [12], the STS detector readout chip SMX2 delivers the hit data as 24-bit words [13], which, after preprocessing and adding the source ID, will be extended to 32-bit words. Similarly, the detector data link (DDL) protocol, used in the ALICE experiment at CERN LHC [14], is oriented toward transmitting 32-bit words.

Concatenating the shorter words received from links into the wider word used by PCIe seems simple. In the case of eight 32-bit links concentrated to a 256-bit PCIe word, a constant group of bits could be allocated for each link. However, such a solution would be inefficient. The links do not transmit the hit data all the time. The data stream may contain additional words with time markers, control, status, and diagnostic information. If there are no data to be sent now, the "empty" words are transmitted. Generally, the data may be divided into "DAQ data", which should be transmitted to DAQ, and "non-DAQ data", which should be discarded. If the PCIe word consists of bit groups permanently assigned to particular links, such discarded data would create "holes" in the data transmitted via PCIe and then stored in the DMA memory buffers in the event selector computer. That results in wasting the PCIe bandwidth and memory and may reduce the performance of data processing.

The "non-DAQ" input data words should be removed before packing the data into the wide output words. However, in that case, the output word must be assembled from the received "DAQ" words in the process of data concentration. It is desired that the "DAQ" words are packed in the same time order as they are received from the links.

Before introducing the method proposed in [15] and further improved in this article, it is worth quickly reviewing alternative solutions. Their number is limited because, despite the high importance of efficient data concentration, it is very sparsely described in most papers about DAQ solutions.

Two alternative approaches are presented in [15]. The first one uses the scanning of $N$ input links at a frequency $N$ times higher than the data rate. All data may be checked and written to the appropriate location in the output word, or discarded before the new dataset arrives. That approach has been successfully used for the initial concentration of data in the CBM STS detector [12]. However, it cannot be used for high data rates and a large number of concentrated inputs.

The second method independently assembles the complete output records from the received "DAQ" data in each input channel. The complete records are then transmitted to the output with round-robin priority. In this method, data from low-rate channels may be significantly delayed (until their output record is filled) compared to the data from high-rate channels. The practical usage of this method is not explicitly described in the literature. But, as shown in [15], it may be found in publicly available firmware sources for the FELIX board used by the ATLAS experiment at CERN [16].

An interesting solution of concentrating eight 32-bit data streams into the 256-bit output has been developed for the LHCb Silicon Pixel Detector [17]. It uses a three-level binary tree of blocks called "2-to-1 encoders", creating the final "8-to-1 encoder" (see Figure 1). The data are concentrated first into four 64 -bit words, then into two 128 -bit words, and finally to a single 256-bit word. Each "2-to-1 encoder" has a relatively complex internal structure, as shown in Figure 2.

Eight 32-bit lines


Figure 1. Structure of an 8-to-1 encoder built from 2-to-1 encoders. Figure reproduced (redrawn) and caption copied from [17]: "A FPGA-Based Architecture for Real-Time Cluster Finding in the LHCb Silicon Pixel Detector," by G. Bassi, L. Giambastiani, K. Hennessy, F. Lazzari, M. J. Morello, T. Pajero, A. Fernandez Prieto, G. Punzi, in IEEE Transactions on Nuclear Science, vol. 70, no. 6, pp. 1189-1201, June 2023, doi: 10.1109/TNS.2023.3273600, CC BY.


Figure 2. Block diagram of a 2-to-1 encoder. R0, R1, R3, and State are registers, MUX0, MUX1, and MUX3 are multiplexers, and FSM is a finite state machine that manages hold, valid, and LE write signals. Figure reproduced (redrawn) and caption copied from [17]: "A FPGA-Based Architecture for Real-Time Cluster Finding in the LHCb Silicon Pixel Detector," by G. Bassi, L. Giambastiani, K. Hennessy, F. Lazzari, M. J. Morello, T. Pajero, A. Fernandez Prieto, G. Punzi, in IEEE Transactions on Nuclear Science, vol. 70, no. 6, pp. 1189-1201, June 2023, doi: 10.1109/TNS.2023.3273600, CC BY.

A relatively complex logic was possible because the concentration worked at a low clock frequency of 30 MHz .

The third method has been introduced in [15] and assumes concentration by properly routing data to the desired position in the output record. In that method, neither multiplication nor the prolonged buffering of data is required. Such routing of data may be performed with an interconnection network. It is a very mature technology [18], widely used for data concentration in telecommunications and data processing systems [19,20]. However, the data concentration in those applications aims to transmit the data via reduced channels [21] and differs from the data concentration in DAQ described earlier. Therefore, the problem studied in [15] and this paper is not a proposal of a new interconnection network but adopting the well-known network architecture for that specific task and proving its correctness.

## Data Concentration for DAQ with Interconnection Network

Example of data concentration with interconnection network is shown in Figure 3.


Figure 3. Example of the concentration of data from 8 inputs to 8 outputs. The "DAQ" words are denoted by " D " with the number in subscript. The "non-DAQ" words are denoted by " N ". (a) In the first concentration cycle, five DAQ words are delivered via links. They are stored in the output record in locations $0-4$. (b) In the second concentration cycle, the next 6 DAQ words are delivered. Three of them are stored in the output record in locations 5-7, but the next three must be stored elsewhere, as the locations 0-2 in the output record are still occupied. For that purpose, the auxiliary record is provided. The output record is completed and ready for sending to DAQ. (c) In the third concentration cycle, another six DAQ words are provided. Three DAQ words from the auxiliary record are copied to the output record. Therefore, filling the output records starts from location 3, and only 5 DAQ words may be stored. The sixth word must be stored again in the auxiliary record.

The data from the links are fed into the concentrating network, which removes the non-DAQ words and delivers the DAQ words to the consecutive locations in the output record. The output record may be partially filled at the end of the concentration cycle. Therefore, the location of the first DAQ word must be selectable. Suppose that, after filling the output record, there are still DAQ words at the link outputs to be received. In that case, they must be stored in an additional auxiliary record, from which they are copied at the beginning of the next concentration cycle. The concentrating network may be built as a multistage interconnection network [22], consisting of switches with two inputs and two outputs (see Figure 4), transmitting the data transparently ("bar" mode) or swapping them ("cross" mode).



Figure 4. Switch used in the concentrating network. Figure based on [22].
To keep track of the current occupancy of the output and auxiliary records and to configure the concentrating network according to the availability of the DAQ words at inputs, an additional "controller" block is needed, which is not shown in Figure 3.

It counts the inputs delivering the DAQ words ("active inputs") and assigns them the consecutive free positions in the output record. After the last position is assigned, the next active inputs are assigned the positions in the auxiliary record.

The development of a data concentrator with eight inputs was presented in [15]. Initially, the $8 \times 8$ Beneš network was used as the concentrating network, as it can perform any data permutation. However, finding the proper configuration of switches for the particular permutation is a complex problem. Therefore, the brute force approach was used, and all needed configurations were stored in the table. During the development, it was found that the full Beneš network is not needed for data concentration. Therefore, finally, the concentrator based on an $8 \times 8$ "reduced Beneš network" was proposed as the optimal solution (see Figure 5).


Figure 5. Concentration of data from 8 inputs. Slightly modified figure from [15], according to CC BY license.

However, concentrating data from more input streams is needed in certain applications. An example may be a concentration of 32 -bit words to a 512 -bit output record (e.g., for 8-lane PCIe Gen 4 or 16-lane PCIe Gen 3) or a concentration of 16 -bit input words to a 256-bit output record.

Therefore, a concentrating network with a higher number of inputs and outputs is needed.

## 2. Attempt to Extend the Previous Concentrator

The basis for the attempt to directly extend the previous concentrator was the assumption that the "reduced Beneš network" (described in [15]) of any size could concentrate data. However, without formal proof, this hypothesis could only be checked by explicitly testing all possible settings of switches.

For a $16 \times 16$ network, the number of switches is 32 , so the number of possible configurations is ca. 4 billion. The tools developed in [15] have been appropriately extended, and all possible configurations have been checked. Indeed, it appeared that such a network can perform the concentration task. However, the problem is the considerable size of the table storing the necessary configurations of the switches, which must be used in the concentrator controller. There are 32 switches to be configured, and the configuration depends on the "DAQ word" flags of 16 inputs and the 4-bit number of occupied words in
the output record. Therefore, storing the configuration in the table requires the memory with $2^{16+4}=2^{20} 32$-bit words. Implementing it in the AMD/Xilinx FPGA would consume 1024 BRAMs or 128 UltraRAMs, which is unacceptable.

A solution could be transforming that table into a combinational function. An attempt was made to use the generated configuration table as the truth table and perform the logic synthesis and optimization. The synthesis was attempted with two commercially available tools: AMD/Xilinx Vivado [23], Intel/Altera Quartus [24], and the open source Yosys [25] environment. Unfortunately, none of those attempts resulted in combinational logic of reasonable complexity. Vivado 2022.2 implemented it as a combinational logic, consuming 524,189 LUTs in an Ultrascale + FPGA. Quartus 21.1 implemented it as a memory, consuming exactly $2^{20} \cdot 32=33,554,432$ memory bits. Therefore, this approach was abandoned and replaced with the one based on the analysis of the concentration network.

## 3. Analysis of the Concentrating Network

The concentrating network described in [15] as the "reduced Beneš network" is, in fact, the "baseline network" [22] with a bit-reversed order of outputs. Furthermore, it will be described as a "baseline network with reversed outputs", or BNRO in short.

The topology of this network for different numbers of layers (and what is related-different numbers of inputs and outputs) is shown in Figures 6-8.


Figure 6. Concentration of data from 8 inputs to 8 outputs. The binary numbers at the inputs of the switches describe from which network inputs the particular input may receive data. The binary numbers at the outputs of the switches describe to which network outputs the particular output may deliver data. The digits in bold are used to find the required mode of the specific switch.


Figure 7. Concentration of data from 16 inputs to 16 outputs. The binary numbers at the inputs of the switches describe from which network inputs the particular input may receive data. The binary numbers at the outputs of the switches describe to which network outputs the particular output may deliver data. The digits in bold are used to find the required mode of the specific switch.


Figure 8. Concentration of data from 32 inputs to 32 outputs. The binary numbers at the inputs of the switches describe from which network inputs the particular input may receive data. The binary numbers at the outputs of the switches describe to which network outputs the particular output may deliver data. The digits in bold are used to find the required mode of the specific switch.

The structure of the baseline network is regular. The network with $N+1$ layers may be created from two baseline networks with $N$ layers by adding a layer with switches where the output 0 is connected to the "upper network" and output 1 is connected to the "lower network" (see Figure 9).


Figure 9. Recursive construction of the BNRO network with $N+1$ layers from two smaller BNRO networks with $N$ layers. The numbers of outputs are in binary form to show that the switch output number in the added layer determines the 0th bit in the output number of the new network. Diagram based on [22].

The rule for determining the output switch also defines the simple and unambiguous routing algorithm: To route the data through the $N$-layer BNRO network from input $k$ to output $m$, we:

- Start from the input $k$ of the network,
- Compare the bit 0 in the input number with the bit 0 of the target output number. The switch should be set to the "bar" mode if the bits are equal. If they differ, the switch should be set to the "cross" mode. After setting the switch in the suitable mode, find the right switch and its input in the first layer.
- Repeat the next steps in the loop until the output of the network is reached:
- In layer $l$, compare the $l$-th bit in the target output and the network input numbers. The switch should be set to the "bar" mode if the bits are equal. If they differ, the switch should be set to the "cross" mode. After setting the switch in the suitable mode, find the right switch and its input in the next layer.
- Increase $l$.

The above algorithm enables routing any input of the BNRO network to any of its outputs. However, it does not warrant that the BNRO network can perform any possible data permutation. The routing is impossible if the data delivered to two inputs of a particular switch should be routed to the same output of that switch, which results in a collision.

The results obtained in [15] and the experiments with $16 \times 16$ (4-layer) network in Section 2 suggest that such a collision should not happen in the BNRO network when used for data concentration. However, formal proof is needed.

## 4. Proof of the BNRO Suitability for Data Concentration

The suitability of the BNRO network to data concentration may be proven in various ways. Below, two of them are presented.

### 4.1. Proof Based on the Mathematical Induction

We can take the network with one layer $(N=1)$ and check that it can perform the concentration task. That is trivial for two inputs and outputs-the network can propagate the data without swapping them (corresponding to writing the output record from location 0 ) or swapping them (corresponding to writing the output record from location 1). The corresponding data may be simply ignored if one or both inputs deliver a non-DAQ word. So, the BNRO with $N=1$ layer can concentrate data (see Figure 10).


Figure 10. All possible cases of data concentration in 1-layer BNRO.
Then, we must prove that, if a BNRO network with $N$ layers can concentrate data, the network with $N+1$ layers, created as shown in Figure 9, can also do it.

Let us assume that the first DAQ word should be routed to the output $t<2^{n+1}$. If the $N+1$ network receives the data with $m<2^{n+1}$ DAQ words, that set of words may be split
into two subsets with sizes differing by no more than 1 . If $t$ is even, we route the first DAQ word to the upper subnetwork. If $t$ is odd, we route that word to the lower subnetwork. The next DAQ words (if available) must be routed alternately to both subnetworks. Such a routing is always possible.

Suppose that the currently handled DAQ word is connected to the switch with another input receiving the non-DAQ word. In that case, we can select the output and the upper or lower subnetwork as required.

Suppose that the currently handled DAQ word is connected to the switch where another input also receives the DAQ word. In that case, we may route it to the required network, and the next DAQ word will be automatically routed to the opposite network, fulfilling the requirement of the alternate routing.

As each subnetwork can perform the data concentration, both the even and odd DAQ words may be routed to the appropriate subnetworks' consecutive (in modulo $2^{N}$ sense) outputs.

If $t$ is even, we configure the upper subnetwork so that even DAQ words are routed to its outputs starting from the $t / 2$ output and the lower subnetwork so that the odd DAQ words are routed to its outputs starting from the $t / 2$ output. In the whole network, those data will appear at consecutive (in modulo $2^{N}$ sense) outputs starting from output $t$. An example of such concentration for 4-layer BNRO is shown in Figure 11.


Figure 11. Example of concentration of data in 4-layer BNRO assuming that 3-layer BNRO can concentrate a subset of data. The case with an even number of the first location in the output record $(t=8)$ is shown. As described in Section 4.1, the even data are passed to the upper subnetwork and routed starting from the output $t / 2=4$. The odd data are passed to the lower subnetwork and routed starting from output $t / 2=4$. After merging the outputs, the data are correctly concentrated, starting from output $t=8$. Please note how the output numbers have wrapped around. Data $D_{8}, D_{9}$, and $D_{10}$ are delivered to outputs 0,1 and 2 . For clarity, the outputs of the subnetworks are shown in the natural order (not in reversed order).

If $t$ is odd, we configure the lower subnetwork so that even DAQ words are routed to its outputs starting from $(t-1) / 2$ output and the upper subnetwork so that the odd DAQ words are routed to its outputs starting from $(t+1) / 2$ output. Those data will appear at consecutive (in modulo $2^{N}$ sense) outputs starting from output $t$ in the whole network. An example of such concentration for the four-layer BNRO is shown in Figure 12.


Figure 12. Example of concentration of data in 4-layer BNRO assuming that a 3-layer BNRO can concentrate a subset of data. The case with an odd number of the first location in the output record $(t=7)$ is shown. As described in Section 4.1, the even data are passed to the upper subnetwork and routed starting from the output $(t-1) / 2=3$. The odd data are passed to the lower subnetwork and routed starting from the output $(t+1) / 2=4$. After merging the outputs, the data are correctly concentrated, starting from the output $t=7$. Please note how the output numbers have wrapped around. Data $D_{9}$ and $D_{10}$ are delivered to outputs 0 and 1 . For clarity, the outputs of the subnetworks are shown in the natural order (not in reversed order).

Based on the mathematical induction principle, the above proves that the BNRO with any number of layers $N \geq 1$ can concentrate data.

### 4.2. Proof Based on Analysis of Collision Possibility

Analyzing the recursive building of the BNRO network and analyzing the topologies of BNRO for small numbers of layers (see Figures 6-8), we may find the rules describing the numbers of network inputs that may be connected to the particular input $i$ of the switch number $r$ located in layer $l$, and numbers of network outputs that may be connected to the particular output $j$ of that switch. These rules are presented in Figures 13 and 14.


Figure 13. Diagram showing the dependency between the number of layer $l$ and the number of switch $r$, and the numbers of the network inputs from which the data may be routed to the input $i$ of that switch. The bits creating the numbers of the accessible network inputs (the lower row) are selected from the vector consisting of the bits of the number of the switch $(r)$, the bit corresponding to the number of the switch input $(i)$, and variable bits $x$ depending on the configuration of the switches in the preceding layers (the upper row).


Figure 14. Diagram showing the dependency between the number of layers $l$ and the number of switches $r$, and the numbers of the network outputs to which the data may be routed from the input $j$ of that switch. The bits creating the number of the accessible network outputs (the lower row) are selected from the vector (the upper row) consisting of variable bits $x$ depending on the configuration of the switches in the following layers, the bit corresponding to the number of the switch output ( $j$ ), and the bits of the number of the switch $(r)$.

With the above rules, we may analyze the conditions resulting in impossible data routing due to a collision in a certain switch. The collision happens when data from both inputs of the particular switch should be routed to the same output, as shown in Figure 15.


Figure 15. Diagram showing a single switch with a collision in data routing. The collision (shown with the red line) occurs when the data delivered to both switch inputs should be routed to the same output.

The network inputs accessible from the individual input have the $2^{l}$ consecutive numbers. The group of numbers associated with inputs 0 and 1 are neighboring, so all inputs accessible from the switch have $2^{l+1}$ consecutive numbers (see Figure 16a). Therefore, the difference between the numbers of two associated inputs must be not higher than $2^{l+1}$.

The outputs accessible from the switch have numbers spread equally and differ by $N \cdot 2^{l+1}$. Therefore, the difference between the numbers of two accessible outputs must be not less than $2^{l+1}$ (see Figure 16b).

Suppose that the collision in the switch should happen. In that case, the data from two network inputs accessible from the switch (with numbers differing by no more than $2^{l+1}-1$ ) should be routed to different network outputs accessible from the same switch output. However, the numbers of those network outputs must differ by no less than $2^{l+1}$. The concentration of data, however, does not insert new data words into the data stream. It may only remove certain not needed data words from that stream. Therefore, during the concentration, for each pair of data words processed in the same concentration cycle, the difference in the numbers between the destination network outputs must be not higher than the difference between the numbers of their network inputs.

This means that the collision never happens when the analyzed network concentrates the data, and the BNRO network can perform the concentration task.


Figure 16. Diagram showing numbers of the network inputs from which data may be routed to the inputs of a particular switch and numbers of network outputs to which the data from its specific output may be routed. (a) The numbers of network inputs connected to the switch in layer $l$ may differ by no more than $2^{l+1}-1$. (b) The numbers of network outputs connected to that switch must differ at least by $2^{l+1}$.

## 5. Implementation

The concentrator based on the BNRO network was implemented in the VHDL language. Its block diagram is presented in Figure 17.


Figure 17. The block diagram of the scalable data concentrator based on an $N$-layer baseline network with reversed outputs (BNRO). It is an extended and scalable version of the concentrator shown in Figure 5.

Its sources are publicly available under the dual GPL/BSD license in the Gitlab repository [26].

The sources are highly parameterized. The user may define the number of layers in the concentrating network and, therefore, the number of inputs and outputs. It is also possible to define the type of data that are concentrated.

The repository contains the testbench, allowing the user to verify the correct operation of the concentrator. As the payload, the integer values are used. The DAQ words with consecutive integer values are delivered to the consecutive inputs. The user may set the probability of the data being accepted by an input. If the data are not accepted, the input is skipped, and the data are tried to be delivered to the next input. The data leaving the concentrator are written to the file. If the concentrator works correctly, the output file contains the consecutive integers. A dedicated Python script checks the content of the generated file.

Additionally, hardware testbenches for the KCU105 [27] and KCU116 [28] boards were created for the 16-input concentrator. Similarly, like in [15], the setup contains FIFOs for storing the input data to be fed into the concentrator and its output data. The block diagram of the test setup is shown in Figure 18.


Figure 18. Testbench for testing the concentrator in the hardware. Slightly modified figure from [15], according to CC BY license.

The testbench is controlled with a Python script using the standard uio_pci_generic driver. Both the script and the testbench have been extended to support a higher number of concentrated inputs.

## 6. Tests and Results

The simulation tests were performed for 4 layers ( 16 inputs) and for 5 layers (32 inputs). Configurations without pipeline registers, with pipeline registers in all layers, and with pipeline registers in selected layers were simulated. Simulations were repeated with different values of the probability of input data availability. All simulations confirmed that the concentrator works correctly, and the concentrated words are delivered in the correct order as a dense data stream without "holes".

For tests in the hardware, the synthesis and implementation with Vivado 2021.2 were performed for two variants-with pipeline registers inside switches and without them. For the KCU116 board, both variants obtained the correct timing closure for 250 MHz clock frequency. The resource consumption is summarized in Table 1.

Table 1. Resource consumption of the 16-input BNRO-based data concentrator together with the testbench for the KCU116 platform. Absolute and percentage (in parenthesis) consumption is given. Separate values for the testbench and concentrator themselves are given. The "non-registered switches" configuration does not use pipeline registers in the concentrating network. The "registered switches configuration" uses pipeline registers in switches in all layers. That increases the resource consumption and latency but may increase the maximum clock frequency.

|  | With Non-Registered Switches |  | With Registered Switches |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | LUTs | Flip Flops | Block RAMs | LUTs | Flip Flops | Block RAMs |
| Available | 216,960 | 433,920 | 480 | 216,960 | 433,920 | 480 |
| Whole testbench | $17,741(8.17 \%)$ | $19,921(4.59 \%)$ | $69(14.37 \%)$ | $17,767(8.19 \%)$ | $22,113(5.10 \%)$ | $69(14.37 \%)$ |
| Data concentrator | $3380(1.56 \%)$ | $1640(0.38 \%)$ | $0(0.0 \%)$ | $3410(1.57 \%)$ | $3830(0.88 \%)$ | $0(0.0 \%)$ |

For KCU105 board, only the version with pipeline registers could work with a 250 MHz clock. The timing closure could be obtained for clock frequency below 197 MHz in the version without pipeline registers. Without the pipeline registers, the whole concentrating network is implemented as a combinational logic, which results in a long critical path. Adding pipeline registers in switches in all layers or in selected ones shortens the combinational path reducing the propagation time between registers and thence increasing the maximum clock frequency. The resource consumption is summarized in Table 2.

Table 2. The resource consumption of the 16-input BNRO-based data concentrator together with the testbench for the KCU105 platform. Absolute and percentage (in parenthesis) consumption is given. Separate values for the testbench and concentrator itself are given. The "non-registered switches" configuration does not use pipeline registers in the concentrating network. The "registered switches configuration" uses pipeline registers in switches in all layers. That increases the resource consumption and latency but may increase the maximum clock frequency.

|  | With Non-Registered Switches |  |  | With Registered Switches |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | LUTs | Flip Flops | Block RAMs | LUTs | Flip Flops | Block RAMs |
| Available | 242,400 | 484,800 | 600 | 242,400 | 484,800 | 600 |
| Whole testbench | $9279(3.83 \%)$ | $10,342(2.13 \%)$ | $50(8.33 \%)$ | $9387(3.87 \%)$ | $11,619(2.40 \%)$ | $50(8.33 \%)$ |
| Data concentrator | $3273(1.35 \%)$ | $2714(0.56 \%)$ | $0(0.0 \%)$ | $3436(1.42 \%)$ | $3921(0.81 \%)$ | $0(0.0 \%)$ |

The three configurations capable of operating at a frequency of 250 MHz frequency were tested in the FPGA boards inserted into PCIe Gen $3 \times 8$ slots in a Linux PC. The testbench was controlled with the prepared Python script. As in simulations, the tests were repeated with different values of the probability of input data availability. All the tests confirmed the correct operation of the concentrator.

## 7. Discussion

The data concentration concept presented in [15] significantly improved the concentration performance compared to the previously used methods like high-frequency polling or width conversion in the input channels. This eliminated the necessity to use very high clock frequencies and disturbances in the time ordering of the concentrated data.

The disadvantage of that solution was its poor scalability. It was based on an $8 \times 8$ interconnection network. Therefore, natively, it supported up to eight inputs and the output word with a length eight times bigger than the input words.

Specific extensions have been proposed in [15], increasing the number of inputs to 9 or 12 but at the cost of introducing additional clocks with slightly higher frequency and higher design complexity. Attempts to use a bigger $16 \times 16$ network was blocked by problems with an excessively large lookup table storing the configuration of switches in the interconnection network.

The new solution described in this paper extends this concept with scalability. It proves that the $N$-layer (with $2^{N}$ inputs and outputs) baseline network with reversed outputs (BNRO) can correctly concentrate the data. It significantly simplifies the design of such a network, offering an efficient and synthesizable implementation of the BNRO controller. The new method enabled the easy implementation of a concentrator with 4 layers ( 16 inputs), confirmed in hardware, and with 5 layers ( 32 inputs), confirmed in simulations.

Currently, the maximum 512-bit size of the output word seems reasonable. Even the planned solutions for PCIe Gen5 assume that width [29]. With the 16-bit minimum width of concentrated data (it must contain the payload and the source ID), it gives a number of concentrated inputs equal to 32 .

In the concentrating network, nothing prevents adding the 6th layer and increasing the number of inputs to 64 if the width of the output word rises to 1024 in the future.

However, the number of inputs may affect the critical path in the network controller. In the current design, where to route the data from the particular input is elaborated in a single clock cycle. That operation includes counting the active inputs (those delivering the DAQ words). Then, the target output is calculated based on the active input number and the current occupancy of the output record (see Section 3). This results in the roughly linear dependency of the critical path in the controller on the number of inputs. However, the input stage of the concentrator may also be pipelined at the cost of additional latency. The input data may propagate through a required number of pipeline registers. This enables counting the active inputs and calculating the target outputs in a few cycles. Finding the optimal solution should be the object of future research.

## 8. Conclusions

This paper significantly improves and extends the data concentration method proposed in [15]. It proves that the baseline network with reversed outputs (BNRO) with any number of layers may concentrate data without the risk of collision. An efficient algorithm for finding the necessary configuration of switches has been formulated, eliminating the need for huge configuration tables. The reference implementation of the concentrator, including the BNRO and its controller, has been created as a parameterized VHDL code. Introducing optional pipeline registers in individual BNRO layers allows the user to find the optimal compromise between resource consumption, concentration latency, and maximum clock frequency. The design was tested in simulations in configurations with four and five layers. The four-layer version was implemented in two different FPGA boards and successfully tested in the actual hardware. The confirmed achieved throughput for sixteen 32-bitinputs at 250 MHz was $128 \mathrm{~Gb} / \mathrm{s}$.

The proposed concentrator may be a good solution for systems where the high-speed and low-latency concentration of multiple short-word data streams into a long-word output is necessary.

Thanks to the availability of sources in the public repository under a permissive double GPL/BSD license, the solution may be widely reused and modified for particular needs.

Funding: The work has been partially supported by the statutory funds of the Institute of Electronic Systems. This project has also received funding from the European Union's Horizon 2020 Research and Innovation Programme under grant agreement No. 871072, and from Polish Ministry of Education and Science programme "Premia na Horyzoncie 2".

Data Availability Statement: The source code of the described concentrator is available under a free and open dual GPL/BSD license in the public repository [26].

Conflicts of Interest: The author declares no conflicts of interest.

## References

1. Neufeld, N. Future of DAQ Frameworks and Approaches, and Their Evolution towards the Internet of Things. J. Phys. Conf. Ser. 2015, 664, 082038. [CrossRef]
2. Meschi, E. The DAQ needle in the big-data haystack. J. Phys. Conf. Ser. 2015, 664, 082032. [CrossRef]
3. Colombo, T.; Amihalachioaei, A.; Arnaud, K.; Alessio, F.; Brarda, L.; Cachemiche, J.P.; Cámpora, D.; Cap, S.; Cardoso, L.; Cindolo, F.; et al. The LHCb Online system in 2020: Trigger-free read-out with (almost exclusively) off-the-shelf hardware. J. Phys. Conf. Ser. 2018, 1085, 032041. [CrossRef]
4. Zemko, M.; Ecker, D.; Frolov, V.; Huber, S.; Jary, V.; Konorov, I.; Kveton, A.; Novy, J.; Veit, B.M.; Virius, M. Triggerless data acquisition system for the AMBER experiment. In Proceedings of the 41st International Conference on High Energy Physics-PoS(ICHEP2022), Bologna, Italy, 6-13 July 2022; p. 248. [CrossRef]
5. Gao, X.; Emschermann, D.; Lehnert, J.; Müller, W.F. Throttling strategies and optimization of the trigger-less streaming DAQ system in the CBM experiment. Nucl. Instruments Methods Phys. Res. Sect. A Accel. Spectrometers Detect. Assoc. Equip. 2020, 978, 164442. [CrossRef]
6. Korcyl, K.; Konorov, I.; Kühn, W.; Schmitt, L. Modeling event building architecture for the triggerless data acquisition system for PANDA experiment at the HESR facility at FAIR/GSI. J. Phys. Conf. Ser. 2012, 396, 012027. [CrossRef]
7. Cuveland, J.D.; Lindenstruth, V.; The CBM Collaboration. A First-level Event Selector for the CBM Experiment at FAIR. J. Phys. Conf. Ser. 2011, 331, 022006. [CrossRef]
8. Albrecht, J.; Campora, D.; Gligorov, V.; Jost, B.; Marconi, U.; Neufeld, N.; Raven, G.; Schwemmer, R. Event building and reconstruction at 30 MHz using a CPU farm. J. Instrum. 2014, 9, C10029. [CrossRef]
9. Paramonov, A. FELIX: The Detector Interface for the ATLAS Experiment at CERN. EPJ Web Conf. 2021, 251, 04006. [CrossRef]
10. Cachemiche, J.; Duval, P.; Hachon, F.; Gac, R.L.; Réthoré, F. The PCIe-based readout system for the LHCb experiment. J. Instrum. 2016, 11, P02013. [CrossRef]
11. DMA/Bridge Subsystem for PCI Express Product Guide (PG195). Available online: https:/ / docs.xilinx.com/r/en-US/pg195 -pcie-dma (accessed on 16 December 2023).
12. CBM Collaboration. Technical Design Report for the CBM Online Systems - Part I, DAQ and FLES Entry Stage; GSI Helmholtzzentrum fuer Schwerionenforschung, GSI: Darmstadt, Germany , 2023. [CrossRef]
13. Kasinski, K.; Szczygiel, R.; Zabolotny, W.; Lehnert, J.; Schmidt, C.; Müller, W. A protocol for hit and control synchronous transfer for the front-end electronics at the CBM experiment. Nucl. Instruments Methods Phys. Res. Sect. A Accel. Spectrometers Detect. Assoc. Equip. 2016, 835, 66-73. [CrossRef]
14. Carena, F.; Carena, W.; Barroso, V.C.; Costa, F.; Chapeland, S.; Delort, C.; Dénes, E.; Divià, R.; Fuchs, U.; Grigore, A.; et al. DDL, the ALICE data transmission protocol and its evolution from 2 to $6 \mathrm{~Gb} / \mathrm{s}$. J. Instrum. 2015, 10, C04008. [CrossRef]
15. Gumiński, M.; Kruszewski, M.; Zabołotny, B.M.; Zabołotny, W.M. Beneš Network-Based Efficient Data Concentrator for Triggerless Data Acquisition Systems. Electronics 2023, 12, 1437. [CrossRef]
16. Wu, W. FELIX: The New Detector Interface for the ATLAS Experiment. IEEE Trans. Nucl. Sci. 2019, 66, 986-992. [CrossRef]
17. Bassi, G.; Giambastiani, L.; Hennessy, K.; Lazzari, F.; Morello, M.J.; Pajero, T.; Prieto, A.F.; Punzi, G. A FPGA-Based Architecture for Real-Time Cluster Finding in the LHCb Silicon Pixel Detector. IEEE Trans. Nucl. Sci. 2023, 70, 1189-1201. [CrossRef]
18. Beneš, V.E. (Ed.) Mathematical Theory of Connecting Networks and Telephone Traffic (Mathematics in Science and Engineering); Academic Press: New York, NY, USA; London, UK, 1965; Volume 17.
19. Sinha, U.K.; Jha, S.K.; Mandal, H. Data Concentration on Torus Embedded Hypercube Network. In Advanced Computing, Networking and Informatics-Volume 2; Kumar Kundu, M., Mohapatra, D.P., Konar, A., Chakraborty, A., Eds.; Springer International Publishing: Cham, Switzerland, 2014; Volume 28, pp. 525-532. [CrossRef]
20. Jain, T. Nonblocking On-Chip Interconnection Networks. Ph.D. Thesis, Technische Universität Kaiserslautern, Kaiserslautern, Germany, 2020. Available online: https:/ /nbn-resolving.de/urn:nbn:de:hbz:386-kluedo-59761 (accessed on 22 December 2023).
21. Jain, T.; Schneider, K. Verifying the concentration property of permutation networks by BDDs. In Proceedings of the 2016 ACM/IEEE International Conference on Formal Methods and Models for System Design (MEMOCODE), Kanpur, India, 18-20 November 2016; pp. 43-53. [CrossRef]
22. Wu, C.-L.; Feng, T.-Y. On a Class of Multistage Interconnection Networks. IEEE Trans. Comput. 1980, C-29, 694-702. [CrossRef]
23. Vivado Design Suite. Available online: https:/ /www.xilinx.com/products/design-tools/vivado.html (accessed on 22 December 2023).
24. Intel Quartus Prime Design Software. Available online: https://www.intel.com/content/www/us/en/products/details/fpga/ development-tools/quartus-prime/resource.html (accessed on 22 December 2023).
25. Yosys Open Synthesis Suite. Available online: https:/ / yosyshq.net/yosys/ (accessed on 12 November 2023).
26. Xconcentrator-Scalable Data Concentrator Based on Baseline Interconnection Network for Triggerless DAQ-Sources. Available online: https: / / gitlab.com/WZabISE / xconcentrator (accessed on 12 November 2023).
27. AMD Kintex UltraScale FPGA KCU105 Evaluation Kit. Available online: https:/ / www.xilinx.com/products/boards-and-kits/ kcu105.html (accessed on 12 November 2023).
28. AMD Kintex UltraScale + FPGA KCU116 Evaluation Kit. Available online: https:/ /www.xilinx.com/products/boards-and-kits/ ek-u1-kcu116-g.html (accessed on 12 November 2023).
29. Versal ACAP Integrated Block for PCI Express. Available online: https://docs.xilinx.com/r/en-US/pg343-pcie-versal (accessed on 16 December 2023).

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.

