

SYSTEM, METHOD, AND APPARATUS FOR WAFER-SCALE MEMORY

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] The present application claims the benefit of and priority to U.S. Provisional Patent Application No. 63/518,988, filed on August 11, 2023, entitled "INTEGRATED CIRCUIT HAVING MEMORIES AND A SHARED WRITE PORT", identified by Docket Number P23-133-US-PSP, the entire contents of which is incorporated herein by reference in its entirety.

[0002] The present application claims the benefit of and priority to U.S. Provisional Patent Application No. 63/602,733, filed on November 27, 2023, entitled "METHOD AND SYSTEM FOR KNOWN-GOOD-DIE TESTABILITY OF FACE-TO-FACE BONDED CHIPLETS", the entire contents of which is incorporated herein by reference in its entirety.

[0003] The present application claims the benefit of and priority to U.S. Provisional Patent Application No. 63/602,737, filed on November 27, 2023, entitled "SYSTEM AND METHOD FOR HAVING CORRECT-BY-CONSTRUCTION TIMING CLOSURE", the entire contents of which is incorporated herein by reference in its entirety.

[0004] The present application claims the benefit of and priority to U.S. Provisional Patent Application No. 63/567,649, filed on March 20, 2024, entitled "ASSEMBLY HAVING A FACE-TO-FACE BONDED CHIPLET", identified by Docket Number P24-052-US-PSP, the entire contents of which is incorporated herein by reference in its entirety.

[0005] The present application claims the benefit and priority to U.S. Provisional Patent Application No. 63/637,742, filed on April 23, 2024, entitled "INTEGRATED CIRCUIT HAVING MICROVAULT MEMORIES", identified by Docket Number P24-081-US-PSP, the entire contents of which is incorporated herein by reference in its entirety.

[0006] The present application claims the benefit of and priority to U.S. Provisional Patent Application No. 63/637,764, filed on April 23, 2024, entitled "FEFET STRUCTURES ON INTEGRATED CIRCUITS", identified by Docket Number P24-082-US-PSP, the entire contents of which is incorporated herein by reference in its entirety.

[0007] The present application claims the benefit of and priority to U.S. Provisional Patent Application No. 63/674,471, filed on July 23, 2024, entitled "SYSTEM, METHOD, AND APPARATUS FOR WAFER-SCALE MEMORY", identified by Docket Number P24-135-US-PSP, the entire contents of which is incorporated herein by reference in its entirety.

BACKGROUND*Relevant Field*

[0008] The present disclosure relates to integrated circuits. More particularly, the present disclosure relates to wafer-scale memory.

Description of Related Art

[0010] Wafers serve as substrates in semiconductor manufacturing and play a role in the production of integrated circuits (ICs). They are typically thin slices of semiconductor material, such as silicon, which is used due to its electrical properties, cost-effectiveness, and the availability of high-quality single crystals. Wafers can vary in size, with diameters ranging from a few millimeters to 450 millimeters, and their thickness is typically controlled to ensure uniformity during processing.

[0011] The manufacturing process for semiconductor devices on wafers may involve multiple stages, including doping, oxidation, photolithography, etching, and deposition. These processes collectively form intricate and precise patterns of transistors, capacitors, resistors, and other electronic components that constitute an integrated circuit. Modern semiconductor devices often incorporate millions to billions, or more, of these components to provide the desired functionality.

[0012] As the demand for higher performance, increased storage capacity, and faster data access speeds continues to grow, wafer-scale integration and memory technologies have evolved to address these requirements. Wafer-scale integration involves fabricating multiple dies on a single wafer and interconnecting them to function as a unified system. This approach can improve performance by reducing latency, increasing bandwidth, and enabling more efficient power distribution compared to conventional packaging methods.

[0013] SUMMARY

[0014] An integrated circuit is disclosed herein that may be part of a semiconductor device. A method of manufacturing or a method of writing and reading data to modules are disclosed herein and can be used with all examples, embodiments, and aspects as described herein.

[0015] In some embodiments, a semiconductor package may include a memory wafer comprising a plurality of modules groups, where each modules group includes multiple modules such as a first module and a second module. The package is designed with a write port configured to write to the entire modules group, along with a first read port dedicated to reading data from the first module and a second read port dedicated to reading data from the second

module. Additionally, an application wafer is bonded to the memory wafer and is in operative communication with the plurality of modules groups.

[0016] In some embodiments, the semiconductor package may include a memory wafer and an application wafer where the outer perimeter of the memory wafer and the outer perimeter of the application wafer are coextensive relative to each other.

[0017] In some embodiments, the semiconductor package may involve an application wafer that optionally includes one or more of various processing elements. These processing elements may include an FPGA core, a microcontroller, a GPU processing element, an ASIC, a CPU core, a neural network accelerator, a digital signal processor (DSP), a reconfigurable processing unit, a machine learning inference engine, a tensor processing unit (TPU), a system-on-chip (SoC), a network processing unit (NPU), an embedded processor, a security module, and a programmable logic device. The semiconductor package may include a shared write port configured to write to both the first module and the second module within the modules group. The semiconductor package may feature a shared write port that is configured to write to an address space, where the shared write port is arranged to write to the first module through a first portion of the address space and to the second module through a second portion of the address space. The semiconductor package may include a write port that is specifically a shared write port, wherein this shared write port is configured to write solely to the modules group.

[0018] The semiconductor package may include a plurality of modules groups that comprise at least one write-group. Each write-group within these modules groups includes a respective read port and a common write port. The semiconductor package may include at least one write-group, which can be further divided into a first write group and a second write group. The first write group can comprise approximately the first half of the plurality of modules groups, while the second write group comprises approximately the second half of the plurality of modules groups.

[0019] The semiconductor package may be configured such that the first write group is the modules group. In some embodiments, a semiconductor package may feature a modules group designated as a write group, where the write port serves as a common write port for the entire write group. The semiconductor package may include a write-group that encompasses all the modules of the memory wafer. This configuration allows the write-group to span across the entirety of the memory wafer, ensuring that all modules are included within the write-group.

[0020] In some embodiments, the semiconductor package includes a memory wafer that is cut into a rectangle. The semiconductor package may include a memory wafer that is cut into a square. The semiconductor package may include a memory wafer consisting of a

plurality of modules groups, wherein each module within these groups is designated as a microvault. The semiconductor package may feature each module within the plurality of modules having an independent read port. This configuration allows for concurrent reading via a respective independent read port for any of the modules. Specifically, the first read port is configured to read data from the first module while simultaneously, the second module can be read using the second read port.

[0021] The semiconductor package may feature a modules group designed to have a single write address space. This configuration allows for writing to the plurality of modules within the modules group using this single write address space. The semiconductor package may feature a first read port that is configured to have a first read address space, along with a second read port that is configured to have a second read address space. The semiconductor package may be configured such that the first read address space numerically overlaps with the second read address space.

[0022] The semiconductor package may be configured such that the first read address space is coextensive with the second read address space. This means that the read address space designated for the first read port and the read address space designated for the second read port encompass the same range of addresses. The semiconductor package may be configured such that the first read address space is contiguous with the second read address space.

[0023] The semiconductor package may include an application wafer that integrates a plurality of processing modules. The semiconductor package may further include a redistribution layer on either the memory wafer or the application wafer to provide power. The semiconductor package may include another wafer that is bonded to the memory wafer to provide power supply routing to all of the plurality of modules. The semiconductor package may include an application wafer that features an application programming interface. The semiconductor package may further include a plurality of vias coupled to the memory wafer and the application wafer. The semiconductor package may include a plurality of vias positioned on the peripheral of the memory wafer. The semiconductor package may include a plurality of vias positioned at the center of the memory wafer.

[0024] The semiconductor package may include a heat sink material that is interposed between the application wafer and the memory wafer. The semiconductor package may feature an application wafer that includes a Network-on-Chip (NoC) with routers interconnecting a plurality of processing elements. The semiconductor package may include a memory wafer configured with a single write-grouping that features a single write port. This configuration ensures that the write port operates as the sole write interface for the entire memory wafer.

[0025] The semiconductor package may be configured such that the plurality of modules is designed for a write endurance of at least once a day for 10 years. The semiconductor package may be configured such that the plurality of modules is designed for a write endurance of at least once a week for 10 years. The semiconductor package may include a memory wafer that is formed from a complete slice of an ingot. The semiconductor package may include a memory wafer that is round. The semiconductor package may feature a memory wafer that has a diameter of 12 inches and includes an aggregate memory capacity of the plurality of modules configured to provide at least 10 trillion bytes of storage. The semiconductor package may comprise an application wafer and a memory wafer that are interconnected using bump-less bonding technology.

[0026] The semiconductor package may include a memory wafer that incorporates memory repair mechanisms configured for in-situ repair of faulty memory modules. The semiconductor package may be configured such that each module of the plurality of modules includes fault-tolerance routing. This configuration enables any module within the memory wafer to route data around a failed module.

[0027] The semiconductor package may include functionality such that when a module within the semiconductor package is determined to be faulty, a corresponding processing element is also determined to be faulty. The semiconductor package may further comprise an interconnection between the memory wafer and the application wafer. This interconnection is configured to utilize through-silicon vias (TSVs), which are designed to form interconnected buses between the memory wafer and the application wafer. The semiconductor package may include an AI inference fabric disposed on the application wafer.

[0028] The semiconductor package may include a plurality of modules configured to support deep neural network operations through large-scale dataflow mapping. The semiconductor package includes each module within the plurality of modules being configured to receive configuration data corresponding to a spatial dataflow. This configuration data is used to optimize the execution of an AI inference task.

[0029] In some embodiments, a semiconductor package may feature a first read port disposed on the surface of the memory wafer. The first module within this configuration comprises a memory array that includes a plurality of memory unit cells arranged within the first module. Additionally, the module may be equipped with read peripheral circuitry that is configured to read data stored within the memory array via the first read port. The semiconductor package may feature an arrangement where the footprint of the read peripheral circuitry on the surface overlaps with the footprint of the memory array on the surface.

[0030] In some embodiments, a semiconductor package may include a first read port on the surface of the memory wafer, where the footprint of the first read port overlaps the footprint of the memory array on the same surface. The semiconductor package may feature a configuration where the footprint on the surface of the read peripheral circuitry is coextensive with the footprint on the surface of the memory array. The semiconductor package may include a memory wafer where the footprint on the surface of the first read port is coextensive with the footprint on the surface of the memory array. The memory array, the read peripheral circuitry, and the first read port are vertically disposed to minimize the module footprint on the surface of the memory wafer. This vertical arrangement ensures that the memory array, the read peripheral circuitry, and the first read port are stacked in a vertical configuration on the surface of the memory wafer. Consequently, the surface area occupied by each module is minimized.

[0031] The semiconductor package may also include a write peripheral circuitry connected to the write port on the surface of the memory wafer. This write peripheral circuitry defines a first footprint on the surface, which does not overlap with the footprint of the memory array and does not overlap with the footprint of the first read port. Additionally, the write port defines a second footprint on the surface, which does not overlap with the footprint of the memory array and does not overlap with the footprint of the first read port. The semiconductor package may include a write peripheral circuitry connected to the write port on the surface of the memory wafer, wherein the first footprint of the write peripheral circuitry overlaps with the second footprint of the write port. This configuration ensures that the two footprints share a common area on the surface of the memory wafer.

[0032] The semiconductor package may further include a plurality of wafer pairs. In this configuration, the application wafer and the memory wafer are combined to form a first wafer pair within the plurality of wafer pairs. The semiconductor package may include a plurality of through-silicon vias (TSVs) that connect the first wafer pair to a second wafer pair among the plurality of wafer pairs. The semiconductor package may feature a plurality of through-silicon vias (TSVs) disposed on the peripheral of the first wafer pair. These TSVs may be strategically positioned around the edges of the wafer pair to facilitate vertical interconnections between the stacked wafers, enabling efficient signal and power transfer within the multi-layered configuration. The semiconductor package may include a plurality of through-silicon vias (TSVs) positioned in a central portion of the first wafer pair. The semiconductor package may include a heatsink material disposed between the first wafer pair and the second wafer pair to facilitate thermal management. The semiconductor package may further include a redistribution layer disposed between the first and second wafer pairs. The

semiconductor package may include a redistribution layer that supplies power to one of the first and second wafer pairs via a through-silicon via (TSV). The semiconductor package may further include another redistribution layer disposed on a second side of the first wafer pair.

[0033] In some embodiments, a method of forming a semiconductor package may involve forming a memory wafer having a plurality of modules groups, where each modules group includes multiple modules such as a first module and a second module. The method also includes configuring the modules group with a write port that is capable of writing data to the entire group, a first read port designed specifically for reading data from the first module, and a second read port designed specifically for reading data from the second module. This procedural step can be followed by bonding an application wafer to the memory wafer, thereby enabling operative communication between the application wafer and the plurality of modules groups on the memory wafer.

[0034] The method of forming a semiconductor package may involve forming the memory wafer and the application wafer such that the outer perimeter of the memory wafer and the outer perimeter of the application wafer are coextensive relative to each other. The method of forming a semiconductor package may include forming the memory wafer and the application wafer such that the application wafer includes at least one of the following processing elements: an FPGA core, a microcontroller, a GPU processing element, an ASIC, a CPU core, a neural network accelerator, a digital signal processor (DSP), a reconfigurable processing unit, a machine learning inference engine, a tensor processing unit (TPU), a system-on-chip (SoC), a network processing unit (NPU), an embedded processor, a security module, and a programmable logic device.

[0035] The method of forming a semiconductor package may involve configuring the write port as a shared write port, which is set up to write to both the first module and the second module within the modules group. The method of forming a semiconductor package may involve configuring the shared write port to write to an address space. Specifically, the shared write port may be set up to write to the first module through a first portion of the address space and to the second module through a second portion of the address space. The method of forming a semiconductor package may involve configuring the write port as a shared write port, which is specifically set up to write solely to the modules group. This ensures that the shared write port directs its write operations only within the confines of the designated modules group and does not extend its write functionality to other parts of the semiconductor package. The method of forming a semiconductor package may involve forming at least one write-group within the plurality of modules groups. This write-group is designed to include a respective read port for

each write-group and a common write port that services the entire write-group, enabling efficient data writing and reading operations within the designated modules. The method of forming a semiconductor package may include creating a first write group and a second write group within at least one write-group. The first write group encompasses approximately the first half of the plurality of modules groups, while the second write group includes approximately the second half of the plurality of modules groups. The semiconductor package may be configured such that the first write group is identified as the modules group. The method of forming a semiconductor package may involve creating the modules group as a write group and configuring the write port as a common write port for the entire write group. This configuration allows the write port to serve as the sole interface for writing data to all the modules within the modules group. The semiconductor package may include a write group that encompasses all the modules of the memory wafer. This configuration allows the write group to span across the entirety of the memory wafer, ensuring that all modules are included within the write group. The method of forming a semiconductor package may further involve cutting the memory wafer into a rectangle. The semiconductor package may involve a method of forming a memory wafer that includes cutting the memory wafer into a square. The method of forming a semiconductor package may involve configuring each module within the plurality of modules as a microvault.

[0036] In some embodiments, a method of forming a semiconductor package may involve configuring each module within the plurality of modules to include an independent read port. This configuration allows for concurrent reading via a respective independent read port for any of the modules. Specifically, the first read port is designed to read data from the first module while simultaneously, another read operation can be conducted from the second module using the second read port.

[0037] The method of forming a semiconductor package may include configuring the modules group to have a single write address space. This configuration is designed to enable writing data to the plurality of modules within the modules group using this single write address space. The method of forming a semiconductor package may include forming the first read port to have a first read address space and the second read port to have a second read address space. The method of forming a semiconductor package may further involve configuring the first read address space such that it numerically overlaps with the second read address space.

[0038] According to some embodiments, the method of forming a semiconductor package may also involve configuring the first read address space to be coextensive with the second read address space. The method of forming a semiconductor package may involve

configuring the first read address space to be contiguous with the second read address space. The method of forming the semiconductor package may include forming a plurality of processing modules in the application wafer.

[0039] The method of forming a semiconductor package may include providing a redistribution layer on either the memory wafer or the application wafer to supply power. The method of forming a semiconductor package may involve bonding another wafer to the memory wafer to provide power supply routing to all of the plurality of modules. This additional wafer is integrated in such a way to ensure effective power delivery across the various modules located within the memory wafer. The method of forming a semiconductor package may involve forming an application programming interface (API) within the application wafer. The method of forming the semiconductor package may involve coupling a plurality of vias to the memory wafer and the application wafer. The method of forming a semiconductor package involves coupling a plurality of vias to the memory wafer and the application wafer. Additionally, this process further includes disposing the plurality of vias on the peripheral of the memory wafer. This arrangement may strategically position the vias around the edges of the memory wafer to facilitate the interconnection between the memory wafer and the application wafer.

[0040] The method of forming a semiconductor package may involve coupling a plurality of vias to the memory wafer and the application wafer. Specifically, this embodiment further includes disposing the plurality of vias at the center of the memory wafer. The method of forming a semiconductor package may involve interposing a heat sink material between the application wafer and the memory wafer.

[0041] The method of forming a semiconductor package may include forming a Network-on-Chip (NoC) within the application wafer. The NoC is designed with routers that interconnect a plurality of processing elements, facilitating efficient communication and data transfer between these elements on the application wafer.

[0042] According to some embodiments, the method of forming a semiconductor package may involve configuring the memory wafer to have a single write-grouping with a single write port. This configuration ensures that the single write port operates as the sole write interface for the entire memory wafer. The method of forming a semiconductor package may include configuring the plurality of modules to ensure a write endurance of at least once a day for a period of 10 years. The method of forming a semiconductor package may involve configuring the plurality of modules to endure writes at a frequency of at least once a week for a duration of 10 years. The method of forming a semiconductor package may involve forming

the memory wafer from a complete slice of an ingot. The method of forming a semiconductor package may include forming the memory wafer as a round wafer.

[0043] The method may include configuring the memory wafer to be a 12-inch wafer. This 12-inch memory wafer may be designed to incorporate an aggregate memory from the plurality of modules, providing a combined storage capacity of at least 10 trillion bytes.

[0044] In some embodiments, a method of forming a semiconductor package may further involve interconnecting the application wafer and the memory wafer using bump-less bonding technology. This method ensures an efficient and reliable connection between the two wafers.

[0045] The method of forming a semiconductor package may include incorporating memory repair mechanisms within the memory wafer. These mechanisms are configured to enable in-situ repair of faulty memory modules, ensuring continued functionality and reliability of the memory components within the semiconductor package. The method of forming a semiconductor package may involve configuring each module within the plurality of modules to include fault-tolerance routing. This fault-tolerance routing is designed to enable any module within the memory wafer to route data around a failed module. This ensures continued data flow and communication integrity within the memory wafer, allowing the remaining operational modules to bypass the affected area efficiently. The method may further involve determining a corresponding processing element to be faulty when it is determined that a module within the semiconductor package has failed.

[0046] The method of forming a semiconductor package may involve interconnecting the memory wafer and the application wafer using through-silicon vias (TSVs). These TSVs are configured to form interconnected buses between the memory wafer and the application wafer. The method of forming a semiconductor package may involve forming an AI inference fabric on the application wafer. The method of forming a semiconductor package may involve configuring the plurality of modules to support deep neural network operations through large-scale dataflow mapping. The method of forming a semiconductor package involves configuring each module within the plurality of modules to receive configuration data corresponding to a spatial dataflow. This configuration data is used to optimize the execution of an AI inference task.

[0047] A method of forming a semiconductor package involves positioning the first read port on the surface of the memory wafer. This process includes creating a memory array composed of several memory unit cells organized within the first module. Additionally, read peripheral circuitry is formed to facilitate reading the stored data within the memory array

through the first read port. The method may further involves forming the footprint on the surface of the read peripheral circuitry to overlap with the footprint on the surface of the memory array. The method may further include forming a footprint on the surface of the first read port such that it overlaps with a footprint on the surface of the memory array. The method of forming a semiconductor package may involve creating a footprint on the surface of the read peripheral circuitry such that it is coextensive with the footprint on the surface of the memory array. The method may further involve forming a footprint on the surface of the first read port so that it is coextensive with a footprint on the surface of the memory array.

[0048] According to some embodiments, the method involves forming the memory array, the read peripheral circuitry, and the first read port in a vertical arrangement on the surface of the memory wafer. This vertical configuration is designed to reduce the footprint of the module on the surface of the memory wafer, effectively minimizing the space occupied by each module. The method of forming a semiconductor package may include creating a write peripheral circuitry connected to the write port on the surface of the memory wafer. This write peripheral circuitry defines a first footprint on the surface, and this first footprint does not overlap with the footprint of the memory array nor with the footprint of the first read port. Additionally, the write port defines a second footprint on the surface, and this second footprint does not overlap with the footprint of the memory array nor with the footprint of the first read port.

[0049] In some embodiments, a method of forming a semiconductor package may involve creating a write peripheral circuitry connected to the write port on the surface of the memory wafer. This write peripheral circuitry is configured to define a first footprint on the surface, which does not overlap with the footprint of the memory array and does not overlap with the footprint of the first read port. Additionally, the write port defines a second footprint on the surface, which does not overlap with the footprint of the memory array and does not overlap with the footprint of the first read port. The method further includes forming the first footprint to overlap the second footprint. The method of forming a semiconductor package may involve forming a plurality of wafer pairs, where the application wafer and the memory wafer combine to form a first wafer pair within the plurality of wafer pairs.

[0050] In some embodiments, a method of forming a semiconductor package may include connecting the first wafer pair to a second wafer pair among the plurality of wafer pairs using a plurality of through-silicon vias (TSVs). The method of forming a semiconductor package may further involve disposing the plurality of through-silicon vias (TSVs) on a peripheral of the first wafer pair. The method of forming a semiconductor package may involve

positioning the plurality of through-silicon vias (TSVs) on a central portion of the first wafer pair. The method of forming a semiconductor package may involve positioning a heatsink material between the first wafer pair and the second wafer pair. The method of forming a semiconductor package may further involve disposing a redistribution layer between the first and second wafer pairs. The method of forming a semiconductor package may further comprise supplying power to one of the first and second wafer pairs via a through-silicon via (TSV) using the redistribution layer. The method of forming a semiconductor package may further involve disposing an additional redistribution layer on a second side of the first wafer pair.

[0051] In some embodiments, a method of using a semiconductor package may involve writing data to a modules group using a write port, where the modules group comprises a plurality of modules, including a first module and a second module. The method further includes reading data from the first module using a first read port and reading data from the second module using a second read port. Additionally, the method encompasses operatively communicating between a memory wafer that contains the modules group and an application wafer that is bonded to the memory wafer.

[0052] In some embodiments, a method of using a semiconductor package may involve writing data to a modules group using a write port, where the modules group comprises a plurality of modules, including a first module and a second module. The method further includes reading data from the first module using a first read port and reading data from the second module using a second read port. Additionally, the method encompasses operatively communicating between a memory wafer that contains the modules group and an application wafer that is bonded to the memory wafer. The method also involves aligning the outer perimeter of the memory wafer to be coextensive with the outer perimeter of the application wafer. The method of using a semiconductor package may involve writing data to the first module and the second module through the use of a shared write port. This approach facilitates the writing process by utilizing a common interface to manage data input to both modules within the modules group.

[0053] According to some embodiments, the method may involve configuring the shared write port to write to an address space. Specifically, the shared write port may be arranged to write to the first module through a first portion of the address space and to the second module through a second portion of the address space. The method may involve configuring a shared write port specifically to write solely to the modules group. The method of using a semiconductor package may involve configuring a plurality of modules groups. Each of these modules groups may be designed to have its own respective read port and a common

write port, ensuring efficient data operations within the modules groups. The method of using a semiconductor package may involve configuring the modules groups into a first write group and a second write group.

[0054] The first write group may comprise approximately the first half of the plurality of modules groups, while the second write group comprises approximately the second half of the plurality of modules groups. The method may involve configuring the first write group to be the modules group. The semiconductor package may involve configuring the modules group as a write group and utilizing a common write port for the write group. The method of forming a semiconductor package may involve configuring all the modules of the memory wafer into a single write group. This configuration allows the write port to serve as the common write interface for all the modules within the memory wafer, thereby enabling unified and coordinated data writing operations across the entire memory wafer.

[0055] The method may involve configuring each module within the plurality of modules to include an independent read port. This configuration allows for concurrent reading via respective independent read ports, ensuring that data can be read simultaneously from the modules without contention. This means that each module can be accessed directly and independently for reading operations, facilitating efficient and parallel data retrieval from the memory wafer. The method of using a semiconductor package may involve configuring the modules group to possess a single write address space. This configuration allows data to be written to the plurality of modules within the modules group using this unified write address space. The method of forming a semiconductor package may involve configuring the first read port to have a first read address space and configuring the second read port to have a second read address space. The method of forming a semiconductor package involves configuring the first read address space to numerically overlap with the second read address space.

[0056] The method may also involve configuring the first read address space to be coextensive with the second read address space. The method of configuring a semiconductor package may involve setting the first read address space to be contiguous with the second read address space. This configuration ensures that the read address spaces for the first and second read ports are adjacent to each other in the memory address space. The method for using a semiconductor package may involve configuring a plurality of processing modules on the application wafer.

[0057] The method may involve supplying power to either the memory wafer or the application wafer using a redistribution layer. The method of using a semiconductor package may involve routing power to all of the plurality of modules. The method of using a

semiconductor package may further include querying an application programming interface (API) on the application wafer. The method of using a semiconductor package may involve communicating between a plurality of processing elements on the application wafer via a Network-on-Chip (NoC) with routers. The method of using a semiconductor package may involve configuring the memory wafer to feature a single write-grouping that includes a single write port. This configuration ensures that the write port operates as the sole write interface for the entire memory wafer, thereby facilitating data writing operations across the memory modules groupings. The method of using a semiconductor package may encompass configuring the plurality of modules to support write endurance of at least once a day for a duration of 10 years. This configuration ensures that each module can endure the specified write frequency over the designated time period, maintaining reliable performance throughout the package's operational lifespan. The semiconductor package may involve configuring the plurality of modules for a write endurance of at least once a week for 10 years. The method of forming a semiconductor package may involve facilitating communication between the application wafer and the memory wafer. The method of using a semiconductor package may involve detecting a faulty memory module within the modules group. This detection process ensures that any defective modules are identified during the operation of the memory wafer and the application wafer.

[0058] According to some embodiments, the method of using a semiconductor package may involve routing around a faulted module within the modules group. This process ensures that data can be redirected to avoid the faulty module, maintaining the integrity and flow of data within the system. The method involves determining a corresponding processing element to be faulty when a module within the modules group is determined to have failed. This process ensures that any defective modules are identified alongside their associated processing elements, maintaining the overall integrity and functionality of the semiconductor package.

[0059] In some embodiments, a method of using a semiconductor package may involve establishing communication between the memory wafer and the application wafer through the use of through-silicon vias (TSVs). These TSVs are configured to form an interconnected bus system, facilitating the transfer of data and signals between the two wafers. Conforming to certain embodiments, a method of using a semiconductor package may involve configuring each module within the plurality of modules to receive configuration data corresponding to a spatial dataflow. This configuration data is designed to optimize the execution of an AI inference task.

[0060] BRIEF DESCRIPTION OF THE DRAWINGS

[0061] These and other aspects will become more apparent from the following detailed description of the various embodiments of the present disclosure with reference to the drawings wherein:

[0062] Fig. 1 is a block diagram of an integrated circuit that may be part of a semiconductor device such as a chiplet in accordance with an embodiment of the present disclosure;

[0063] Fig. 2 shows a perspective view of an assembly having the integrated circuit of Fig. 1 implemented on a semiconductor device that is electrically connected to another device to form the assembly in accordance with an embodiment of the present disclosure;

[0064] Fig. 3 shows a block diagram illustrating the memory address space of the integrated circuit of Fig. 1 in accordance with an embodiment of the present disclosure;

[0065] Fig. 4 shows a block diagram illustrating the memory address space with the signal interfaces of the integrated circuit of Fig. 1 in accordance with an embodiment of the present disclosure;

[0066] Fig. 5 shows an illustration of an integrated circuit that may be part of a semiconductor device such as a chiplet in accordance with an embodiment of the present disclosure;

[0067] Fig. 6 shows a perspective of an assembly having the integrated circuit of Fig. 1 implemented on a semiconductor device that is electrically connected to a system-on-a-chip in accordance with an embodiment of the present disclosure;

[0068] Fig. 7 shows a perspective view of an assembly having a semiconductor device with an array of processing elements and a second semiconductor device having an array of microvaults;

[0069] Fig. 8 shows an assembly of a semiconductor devices including several memory types in accordance with an embodiment of the present disclosure;

[0070] Fig. 9 shows an assembly of semiconductor devices including a semiconductor device with a system-on-chip and another semiconductor with microvaults disposed on top in accordance with an embodiment of the present disclosure;

[0071] Fig. 10 shows a semiconductor assembly incorporating a daisy-chained configuration of microvaults operatively connected to a multiplexer and managed by a counter for coordinated data selection and retrieval, in accordance with an embodiment of the present disclosure.

[0072] Fig. 11 shows a semiconductor assembly incorporating a daisy-chained configuration of microvaults in multiple semiconductor devices that are operatively connected

to multiplexers and managed by counters for coordinated data selection and retrieval, in accordance with an embodiment of the present disclosure;

[0073] Fig. 12 illustrates a three-dimensional (3D) memory column configured as a 3D-NOR or 3D-AND structure, featuring a series of ferroelectric field-effect transistors (FeFETs) with interconnected drain terminals linked to a common select line and individual gate terminals connected to respective read/write enable lines, all coupled to a common bit line, in accordance with an embodiment of the present disclosure;

[0074] Fig. 13 depicts a three-dimensional (3D) memory column configured as a 3D-NAND structure, consisting of a vertical stack of ferroelectric field-effect transistors (FeFETs), in accordance with an embodiment of the present disclosure;

[0075] Fig. 14 depicts a three-dimensional (3D) memory column configured as a 3D-NAND with an integrated pass gate, in accordance with an embodiment of the present disclosure;

[0076] Fig. 15 illustrates a three-dimensional (3D) memory column 1500, which may be configured as either a 3D-NOR or a 3D-AND structure with independent Read/Write enable capabilities, in accordance with an embodiment of the present disclosure;

[0077] Fig. 16 shows a cross-sectional view of a 3D memory structure configured as a single-port 3D NAND, in accordance with an embodiment of the present disclosure;

[0078] Fig. 17 shows a cross-sectional view of a 3D memory structure that is a dual-port 3D NAND arrangement, in accordance with an embodiment of the present disclosure;

[0079] Fig. 18 illustrates a 3D memory structure that can be configured as a 3D NOR Vertical Transistor memory array, in accordance with an embodiment of the present disclosure;

[0080] Fig. 19 shows a planar FeFET in accordance with an embodiment of the present disclosure;

[0081] Fig. 20 shows electrical characteristics of an embodiment of a FeFET in accordance with an embodiment of the present disclosure; and

[0082] Fig. 21 shows a semiconductor package using a memory wafer and an application wafer in accordance with an embodiment of the present disclosure.

[0083] **DETAILED DESCRIPTION**

[0084] Fig. 1 shows a block diagram of an integrated circuit 100 that may be packaged as a bondable chiplet (e.g., face-to-face chiplet bondable) in accordance with an embodiment of the present disclosure. The integrated circuit (IC) 100 includes a modules group 106 consisting of modules 108, 110, 112, and 114. The IC 100 also features a shared write port 102, configured to write to the modules group 106 using a write peripheral 104. Additionally, it

includes read peripherals 116, 118, 120, and 122 and read ports 124, 126, 128, and 130, configured to read from the modules 108, 110, 112, 114.

[0085] The write port 102 may be configured to provide a single write address space for all of the modules group 106 where each of the modules 108, 110, 112, 114 has a dedicated read port 124, 126, 128, 130 respectively. The integrated circuit 100 may be packaged as part of a chiplet configured to be electrically connected to another integrated circuit device (e.g., another chiplet, or IC package, with or without electrical contacts, electrical bumps, etc.). The chiplet may be electrically connected to another device including, for example, by bonding, soldering, wafer-to-wafer bonding, face-to-face chiplet bonding, chiplet-to-wafer bonding, chiplet-to-interposer bonding, and/or may be connected together with an interposer or other interfacing technology. None, one, or more of interposers may be used or other interfacing technologies that are common to heterogeneous 3D system-in-package solutions may be utilized in electrically connecting a chiplet to another device.

[0086] Each read port (124, 126, 128, 130) in the chiplet may feature electrical contacts on a side of the chiplet or on multiple sides of the chiplet. The read ports 124, 126, 128, 130 may use multi-cycle pipelined circuitry. Upon bonding to another device (e.g., wafer, chiplet, chip, SOC, package, FPGA, etc.), the electrical contacts may line up in a manner that provides dedicated access to specific modules of the modules 108, 110, 112, 114. For instance, a processing/computing element may have exclusive access to module 108 via the read port 124, which may contain the neural network weights in a register file. Similarly, a different processing/computing element may have exclusive read access to module 110 via the read port 126, which includes a different register file. In this specific embodiment, this arrangement of the electrical contacts ensures that each computing/processing element has the dedicated access it needs to carry out its specific computation efficiently thereby providing a compact, modular, and scalable system that allows different processing elements to maintain dedicated access to specific modules 108, 110, 112, 114. Without dedicated access, different processing elements might have to queue up to use the same resource which would slow down overall processing speed. By providing dedicated access, the proposed chiplet ensures that each processing element can operate at its maximum capability without interference from other computing elements in this specific embodiment.

[0087] The write peripheral 104 is a peripheral circuitry responsible for processing and writing data into the memory cells found within the modules 108, 110, 112, 114. The write peripheral 104 may include dedicated contacts so that a chip electrically connected (e.g., bonded) to a chiplet of the integrated circuit, such that the write port 102 is accessible via a

shared write logic system that involves utilizing a shift register-based, different voltage design, preferably high voltage design, that has a shared write address and data components. This shared write logic system is designed to be accessed via a bonded chip, another bonded chiplet, and/or via other circuitry in the same package as the integrated circuit 100. A shift register could allow the system to move data through a series of stages, with each subsequent stage receiving the data from the previous stage. By utilizing a shift register, the system can increase the data throughput while maintaining a low rate of data transfers. The shared write address space refers to the location where data is written in the chiplet.

[0088] In another embodiment, an interlock 132 may disable the read ports 124, 126, 128, 130 while data is being written to the modules group 106 via the write port 102. Likewise, the interlock 132 may disable the write port 102 when read operations are being carried out on the read ports 124, 126, 128, 130. The written data can later be accessed concurrently by all processing elements that need to read the data via a respective one of the read ports 124, 126, 128, 130. This ensures that all processing elements have the most commonly used data available to them without regard to other reads being concurrently carried out by other processing elements.

[0089] The write peripheral 104 circuit includes a write driver. This unit receives the data to be written and converts it into suitable signals that can change the state of the memory cells. Depending on the type of memory technology used, these signals could involve voltage levels, current pulses, or other types of energy. The shared write logic system may be high voltage due to the specific voltage requirements of the chiplet. The write driver must provide enough power to reliably change the state of the memory cells, but it must also operate within suitable parameters to avoid causing damage or unnecessary wear.

[0090] The write peripheral 104 circuit may also feature a data buffer or write buffer. This component temporarily stores the data to be written, allowing the write operation to be performed at an optimal pace. By balancing the speed of incoming data with the speed at which the memory cells can be written, the write buffer helps prevent data loss and optimizes system performance.

[0091] The write peripheral 104 may also include, in some embodiments, a write control unit that orchestrates the sequence of operations in the write process. It generates control signals to activate the write driver at the appropriate times, controls the flow of data from the write buffer, and coordinates the timing of the write operations. By synchronizing these various activities, the write control unit ensures efficient and reliable write operations.

[0092] The write peripheral 104 may also include data encoding mechanisms to improve reliability and data integrity. For example, before the data is written to the memory cells, these mechanisms encode it in a way that allows potential errors to be detected, and in some cases, corrected when the data is later read. This can be helpful in systems where data integrity has a higher priority, such as in servers or scientific research devices.

[0093] The write peripheral 104 may also include a timing unit that serves as the system's heartbeat, supplying clock signals that synchronize the operation of the system's various components. In some systems, it may include components like oscillators, clock generators, or phase-locked loops. The timing unit may ensure that all operations occur at the suitable time relative to each other.

[0094] The IC 100 may be implemented as a face-to-face bonded chiplet, with modules 108, 110, 112, and 114 formed from a non-volatile memory. In some specific embodiments, the IC 100 may also feature a dynamic allocation circuitry to allocate memory blocks to the modules group 106 based on the usage of the modules group 106 (e.g., each module 108 may include dynamic allocation circuitry for dynamically allocating a range of read locations for a respective processing element).

[0095] The IC 100 features a plurality of clocks, with each clock of the plurality of clocks feeding a respective module of the plurality of modules, providing each respective module with decoupled timing relative to the other modules of the plurality of modules. The modules group 106 may be arranged in any topology known to one of ordinary skill in the relevant art. Bit-cell density can be up to 10 times more dense than embedded SRAM cells in the modules group 106.

[0096] The IC 100 may be formed on a chiplet that includes a first side and a second side, with the second side configured for bonding to a second semiconductor device. The IC 100 may include a high voltage write logic adjacent to the first side of the chiplet. A decoder circuitry, a driver circuitry, and a register circuitry may be formed on the silicon substrate portion of the chiplet, while the modules group 106 is formed on a second layer portion of the chiplet. The second semiconductor device may comprise a plurality of processing elements. Each processing element includes a respective interface to communicate with a respective module of the plurality of modules on the modules group 106 when the second semiconductor device is bonded to the chiplet.

[0097] The silicon substrate traditionally serves as the initial stage of IC fabrication, focusing on the creation of active components, particularly transistors. Techniques like diffusion, ion implantation, oxidation, and material deposition are employed to fashion the

intricate structures of transistors. These processes operate at small scales. The application of photolithography, etching, and implantation techniques enables the definition of transistor structures with precision. The silicon substrate's significance lies in its ability to establish the fundamental building blocks necessary for signal processing, amplification, and control within the IC. This layer is sometimes called Front-End-Of-The-Line ("FEOL").

[0098] Next in the manufacturing process, a second layer may be added that traditionally takes on the role of interconnect fabrication, facilitating the electrical connections between various IC components. This phase traditionally focused on the creation of passive components, including interconnects, vias, and metal-insulator-metal (MIM) capacitors. The second layer processes typically differ from the processes used on the silicon substrate in terms of precision and scale. The interconnects are formed by depositing and patterning metal layers, typically aluminum or copper, to construct the wiring network. Dielectric layers, such as silicon dioxide or low-k dielectrics, are introduced to insulate the interconnects and prevent signal interference between different wiring layers. The second layer's traditional function is to establish the necessary interconnections that enable the routing and distribution of electrical signals throughout the IC. However, as described herein, circuitry may be utilized within this second layer (sometimes referred to as Back-End-Of-The-Line ("BEOL")).

[0099] Alternate embodiments of the IC 100 may be implemented as a stacked die, a monolithic design, TSVs, or silicon through vias. In a stacked die design, several dies may be stacked on top of each other, with each die performing different functions, such as memory and processing. The stacked die may communicate through wire bonds, microbumps, or bump-less bonds. In a monolithic design, the various functions and modules of the IC 100 may be integrated onto a single die, forming a more compact and power-efficient design.

[00100] Additionally, the IC 100 may include one or more interlocks 132 to prevent conflicts in reading and writing data. The modules group 106 may be formed from a variety of non-volatile or semi-volatile (e.g., very long refresh periods) memory technologies, such as Static Random-Access Memory (SRAM), Ferroelectric Field Effect Transistor (FeFET), Ferroelectric Random Access Memory (FeRAM), Resistive Random Access Memory (ReRAM), Spin-Orbit Torque (SOT) Memory, Spin Transfer Torque (STT) Memory, charge trap, floating gate memories, and/or Schottky diodes.

[00101] The modules group 106 may utilize a Static Random-Access Memory (SRAM) Topology. The SRAM topology may employ a cross-coupled flip-flop structure (e.g., latching flip-flops), ensuring the stored data remains intact as long as power is supplied. Thus, in some

specific embodiments, the modules group 106 may utilize heterogeneous types of memory including volatile and non-volatile memory types.

[00102] The modules group 106 may utilize a Flash Memory Topology. The Flash memory is a non-volatile memory technology used in applications where data persistence is needed, such as solid-state drives (SSDs) and USB flash drives. The flash memory topology disclosed herein features a matrix of memory cells, each consisting of a floating-gate transistor or charge trap device. The modules group 106 may also use wear-leveling techniques to prolong the lifespan of the memory cells.

[00103] The modules group 106 may utilize a Ferroelectric Random-Access Memory (FeRAM) Topology. The FeRAM topology utilizes a ferroelectric material capable of retaining polarization states. One such memory topology may, in specific embodiments, utilize a FeFET to retain state information and program the ferroelectric material. These ferroelectric materials may be used to retain state information and act as a memory bit cell.

[00104] The modules group 106 may utilize a Phase Change Memory (PCM) Topology, which is a non-volatile memory technology that utilizes reversible phase changes in materials to store data. The PCM topology may include any phase change material, for example a chalcogenide alloy or a chalcogenide glass housed within a memory cell.

[00105] The modules group 106 may utilize a Resistive Random-Access Memory (ReRAM) Topology, which is a non-volatile memory technology based on resistive switching phenomena. The ReRAM topology may utilize a thin-film material that exhibits reversible changes in resistance upon the application of electrical stimuli.

[00106] The modules group 106 may utilize a Spin-Orbit Torque (SOT) Magnetic Random-Access Memory Topology. SOT-MRAM is a type of non-volatile memory that utilizes spin-orbit torque to switch the magnetic state of a storage element. The SOT-MRAM topology may incorporate a magnetic tunnel junction (MTJ) structure and leverages the spin-orbit coupling effect to write and read data. The magnetic tunnel junction may have a dielectric layer between a magnetic fixed layer and a magnetic free layer. Writing may be done by switching magnetization of the free magnetic layer by injecting an in-plane current in an adjacent SOT layer. Reading may be done by putting current into the magnetic tunnel junction. The SOT-MRAM can optimize the spin-orbit materials by using current-driven switching schemes while minimizing write energy consumption, in some specific embodiments.

[00107] The modules group 106 may utilize a Spin Transfer Torque (STT) Magnetic Random-Access Memory Topology. The STT-MRAM is another type of non-volatile memory that relies on spin transfer torque to manipulate the magnetic state of a storage element. The

STT-MRAM topology can use a magnetic tunnel junction (MTJ) structure, where the magnetization orientation determines the stored data. Additionally, the orientation of a magnetic layer in a magnetic tunnel junction or spin valve can be changed using a spin-polarized current, for example.

[00108] The IC 100 may include a single write peripheral 104 with a dedicated clock, or each module 108, 110, 112, 114 may have its own dedicated write peripheral utilizing a shared clock (not shown in Fig. 1). Additionally, the modules group 106 may be organized into separate partitions, each with a dedicated read peripheral 116, 118, 120, 122 having an independent clock.

[00109] Another possible embodiment of the IC 100 includes an interface (e.g., the same, different, higher or lower voltage) to enable data transfer external to the packaging of the IC 100. The IC 100 may also include an integrated microcontroller unit (MCU) or a digital signal processor (DSP) for processing data within the IC in yet additional specific embodiments.

[00110] Fig. 2 shows a perspective view an assembly 200 of the integrated circuit 212 of Fig. 1 implemented on a chiplet 230 that is bonded to a second device 226 in accordance with an embodiment of the present disclosure. The integrated circuit 212 is the circuitry within the chiplet 230. The second device 226 may be a chiplet, semiconductor wafer, semiconductor package, encased circuitry, etc. For example, the second device 226 may be an AI accelerator such that each processing unit has read access to one module (or a predetermined set) of the modules group 236. In yet another embodiment, the second device 226 may be a network controller where there is an offload circuit to read the data from each of the modules to processing incoming/outgoing packets, etc. The assembly 200 includes a modules group 236 having a plurality of modules, including a first module 232 and a second module 234. Fig. 2 shows several modules, however, for clarity, only modules 232, 234 have reference numbers. The integrated circuit 212 further comprises a shared write port 222. The shared write port 222 interfaces into the write peripheral 202.

[00111] Although the second device 226 may use the shared write port 222 via an address & data bus with a clock and a enable signal to write data to any modules within the modules group 236, other ways of writing data may be considered. For example, serial connections, parallel connections, various buses, or ports, may be used, such as a DDR (Double Data Rate) Interface, a SRAM (Static Random-Access Memory) Interface, a NAND Flash Memory Interface, a NOR Flash Memory Interface, a HBM (High Bandwidth Memory) Interface, a GDDR (Graphics Double Data Rate) Interface, a NVMe (Non-Volatile Memory

Express) Interface, SPI, IC2, etc. Each of the modules has a read port with a read address 218 (to send an address to a module 234) and read data 214 (which is the data read from the module 232).

[00112] The modules group 236 is formed on a chiplet 230 having two sides including a surface 228 that can be bonded to and complement a second device 226. The chiplet 230 may be formed by forming circuitry on a silicon substrate 204 and then by adding a second layer 206. In other embodiments, these layers may be reversed and/or other layers may be added, removed, etc. The read address 218 and read data 220 are used for reading the module 232.

[00113] Although the second device 226 may use an address & data bus with a clock and a enable signal to read data from the module 232, other ways of reading data may be considered. For example, serial connections, parallel connections, various buses, or ports, may be used, such as a DDR (Double Data Rate) Interface, a SRAM (Static Random-Access Memory) Interface, a NAND Flash Memory Interface, a NOR Flash Memory Interface, a HBM (High Bandwidth Memory) Interface, a GDDR (Graphics Double Data Rate) Interface, a NVMe (Non-Volatile Memory Express) Interface, SPI, IC2, etc.

[00114] All of the read ports (e.g., 218 and 222) are configured to be inactive when a write operation is applied to the shared write port 222. The read ports may also be configured to process reads concurrently with each other. The shared write port 222 is configured to write to an address space, where the shared write port 222 is configured to write to the first module 232 via a first portion of the address space and write to the second module 234 via a second portion of the address space. Each module of the plurality of modules 236 includes an independent read port for concurrent reading via a respective independent read port of any of the plurality of modules.

[00115] Each read port for a respective module may include contacts for circuitry found within the second device 226 to interface via metallic contacts. Thus, there may be metallic contacts on the top layer 208 that are configured to interface with metallic contacts on the surface 228 of the chiplet 230 such that the metallic contacts allow for a read space that is coextensive with a read space of a module of the modules 236. The read spaces of the modules group 236 may all be coextensive with each other (as is described with reference to Figs. 3 and 4).

[00116] In one embodiment, the read peripheral for the first module 232 is implemented on a silicon substrate 204 (sometimes referred to as a Front-end-of-the-line). The second layer 206 (sometimes call the Back-end-of-the-line) may be built next in the manufacturing process

on top of the silicon substrate 204 (and any circuitry) and may contain the respective memory bit cells. In an alternative embodiment, the read peripheral for the first module 232 is implemented in the second layer 206 and is disposed between the modules group 236 and the surface 228 of the chiplet 230.

[00117] The modules group 236 may be configured to process write commands only during reset. The write commands may be “slow write” commands. That is, the modules group 236 may have very low write speeds relative to its read speed. The write logic may be frozen (or disabled) when the modules group 236 are used for reading data. In some specific embodiments, the integrated circuit 212 provides functionality to allocate memory blocks to the modules group 236 based on the usage of the modules group 236. In other embodiments, the memory addresses are fixed along with the allocation. The integrated circuit 212 may be implemented as a face-to-face bonded chiplet 230. The face-to face bonding may be bump-less wafer bonding.

[00118] The modules group 236 can have a single write peripheral 202. In other embodiments, each module of the modules group 236 may have a dedicated write peripheral that utilizes a shared clock. In yet other embodiments, the modules group 236 may also be organized into separate partitions each with partition having a dedicated read peripheral, where each dedicated read peripheral has an independent clock. The partitions may be one, two, or more modules of the modules group 236.

[00119] The write peripheral 202 circuitry's overall architecture may include a series of different components, including write driver, address decoders, sense amplifiers, data input latches, data bus, etc. and/or some combination thereof. Write drivers or write buffers, may be tasked with transferring data onto the memory cell. They may enhance the input signal to achieve a level appropriate for the memory cell. Address decoders may be used to interpret the memory address that is fed as an input where the data needs to be written. By activating the specific row and column of the memory array linked to that address, they may be used to select the target memory cell. Sense amplifiers may be used to identify and boost the signal from the memory cells during reading operations, also participate in refreshing the memory cell post data write in write operations. The write operation is instigated by a write enable signal. When a write command is initiated, this signal propels the write drivers and decoders into the writing process. Data input latches may be used as temporary storage units, retaining the data set to be written into the memory until the write operation is implemented. A data bus with a transmission route, can be used to facilitate the movement of data from the data input latches to the memory cells.

[00120] A write operation to the modules group may be performed through a priority arbitration circuit that facilitates the modules to be accessed in a predetermined order, and the shared write port 222 may be configured to write to a virtual address space that is mapped onto a physical memory space. The integrated circuit 212 may include a high voltage write logic used within the write peripheral 202, and the second semiconductor device 226 may comprise a plurality of processing elements, whereby each processing element includes a respective interface to communicate with a respective module of the modules group 236. Furthermore, the chiplet 230 may include an interface to the shared write port 222 on the second side to thereby interface with a complementary interface on the second semiconductor device 226.

[00121] The integrated circuit 212 may also include a power gating circuitry that selectively powers down a module of the modules 236 when not in use. Additionally, the integrated circuit 212 may have a write peripheral 202 of the modules group 236 connected to a dedicated I/O pad to enable data transfer external to the package of the integrated circuit.

[00122] The integrated circuit 212 may utilize multiple modules of the modules group 234 grouped together. These modules may be synchronized with one another in specific embodiments. In some cases, all the modules are synchronized, while in other instances, only specific modules are to be synchronized. For instance, the circuit on a second device 226 may need to synchronize with a specific module when reading data from one of the modules in the modules group 236.

[00123] To synchronize the modules, the integrated circuit 212 may use various timing technologies. In some cases, a plurality of clocks may feed each respective module of the modules group 236, thereby allowing each module to have decoupled timing relative to the other modules in the group. This decoupling ensures that any delay in one module will not affect the functioning of other modules. It is worth noting that the clocks used may or may not need to be synchronized. In some cases, a common clock can be used to synchronize the modules. In yet other embodiments, the clock signal or signals may be provided by the second device 226.

[00124] In alternative embodiments, other synchronization techniques can be used, such as phase comparison of the clock signals or a phase-locked loop (PLL) synchronization method. Another embodiment for synchronizing the modules in the IC could use delay-locked loop (DLL) synchronization. In this method, a delay element is added to the clock signal path, and the output is compared to the input clock signal. The feedback loop adjusts the delay element until the output of the DLL matches the input, resulting in synchronization of the clock signals.

[00125] In another embodiment, the integrated circuit 212 could use a combination of different synchronization techniques to achieve synchronization between the modules. For example, some modules may use PLL synchronization while others use clock delay lines or DLL synchronization, depending on their specific requirements. Additionally, the integrated circuit 212 can also use redundant synchronization techniques to ensure reliability and redundancy in case one method fails. For example, the integrated circuit 212 could use both PLL synchronization and DLL synchronization simultaneously, so that if one method fails, the other can still maintain synchronization.

[00126] Fig. 3 shows a block diagram 300 illustrating the memory address space of the integrated circuit of Fig. 1 in accordance with an embodiment of the present disclosure. The memory address space includes a write address space 316 and read data address spaces 310, 312, 314.

[00127] The write address space 316 consists of various units where data, e.g., weights, and/or instructions can be stored. These units are referred to as memory addresses. The modules group 302 includes multiple memory modules 304, 306, 308. The write address space 316 may be distributed among the memory modules 304, 306, 308 such that the write address space 316 spans from 0 to $N*M-1$. As shown in Fig. 3, the modules group 302 has N memory modules 304, 306, 308, where N is a positive integer, and each module has a memory size of M . The total number of unique write memory addresses in the write address space will be $N*M$, which can be referenced by an integer from 0 to $N*M-1$.

[00128] Starting at 0, memory addresses of the write address space 316 are ordered sequentially up to $N*M-1$. In other words, the first address is 0 and the final address is $N*M-1$, encompassing a total of $N*M$ addresses. This ordering can be linear (each address increases by one) or some other specified pattern depending.

[00129] The write memory addressing can be implemented in a variety of ways based on the system architecture. One method used in a specific embodiment is to use the base and limit registers. The base register holds the smallest legal physical write memory address, and the limit register specifies the size of the range. Therefore, to generate a logical address, you would add the base to the relative address. In other embodiments, a memory addressing scheme may be used where the base used is set to be 0. Yet additional write addressing techniques will be appreciated by one of ordinary skill in the relevant art.

[00130] For any device that writes to the modules group 302, each memory module can possess a unique set of write memory addresses such all memory addresses within the modules group 302 is unique with respect to writing data. e.g., the first module starting at 0 and the last

one ending at $N*M-1$. This allocation, in some embodiments, may be dependent on the memory management system of the device writing data to the modules 304, 306, 308, which could range from simple fixed partitioning schemes to more complex dynamic partitioning models.

[00131] For instance, in a straightforward linear model where each module (304, 306, or 308) has an equal size of M addresses, the first module 304 would possess write addresses 0 to $M-1$, the second module would have write addresses M to $2M-1$, the third module would have write addresses $2M$ to $3*M-1$, and so forth. The N th module 308, therefore, would possess write addresses from $(N-1)*M$ to $N*M-1$.

[00132] It is contemplated that one of ordinary skill in the relevant art may use other implementations of write memory addresses from 0 to $N*M-1$ that depends on various factors such as the hardware architecture, operating system, memory management schemes, and the nature of the programs being run on the system, etc.

[00133] The modules group 302 has different read data address spaces 310, 312, 314. These read address spaces 310, 312, 314 may have overlapping addresses spaces, may have contiguous address spaces, or may have coextensive address spaces. The read address spaces 310, 312, 314 may be independent relative to each other. The system includes three independent read address spaces, labeled as read address spaces 310, 312, and 314. Each of these read address spaces is distinct from the others, meaning that reads can be performed in each space without affecting the others.

[00134] The read address spaces 310, 312, 314 may be defined as contiguous blocks of memory addresses, each with its own starting address and ending address. In modules group 302, each read address space 310, 312, 314 may have a range of addresses that corresponds to values from 0 to $M-1$, where M is a maximum value determined by the size of the modules 304, 306, 308 being used.

[00135] In one embodiment, allowing one processing unit to interface with each read address space 310, 312, 314, the concurrent reads may be implemented as described herein. The independence of the read address spaces 310, 312, 314 ensures that each processing unit can access its desired data without causing any interference or conflict with other processing units.

[00136] Fig. 4 shows a block diagram illustrating the memory address space with the signal interfaces of the integrated circuit of Fig. 1 in accordance with an embodiment of the present disclosure. The signals used in Fig. 4 may be used with any embodiment described herein. However, one of ordinary skill in the relevant art will appreciate that different signaling schemes may be used.

[00137] The modules group 402 includes modules 404, 406, 408 that share a common write peripheral 411. The write peripheral 411 includes a write address bus that includes the address of the data being written, a write data bus that includes the data, a write clock cause the writes to occur (e.g., either on a leading or trailing edge of the clock signal, etc.). The writes only occur if the write enable signal indicates a write should occur. Any logic may be used, e.g., high voltage may correspond to 1 and a low voltage may correspond to 0, or vice versa. In some embodiments, the write peripheral 411 may be on the chiplet 230 and in other embodiments, the write peripheral 411 is on the second device 226.

[00138] The modules group 402 has modules 404, 406, 408 where each has a respective read peripheral 410, 412, 414. Each of the read peripheral 410, 412, 414 has a read address bus to send an address for reading, a read data bus to receive the data, a read clock which is the clock used to control the timing of the output of the digital data, and an output enable that is a precondition to outputting data. Any logic may be used, e.g., high voltage may correspond to 1 and a low voltage may correspond to 0, or vice versa. In yet additional embodiments, multi-bit or analog data storage may be used. In some embodiments, one or more of the read peripherals 410, 412, 414 may be on the chiplet 230 and in other embodiments, one or more of the read peripherals 410, 412, 414 are on the second device 226.

[00139] Fig. 5 shows an illustration of an integrated circuit 500 that may be part of a semiconductor device such as a chiplet in accordance with an embodiment of the present disclosure. The integrated circuit 500 may be disposed on a semiconductor device, such as a chiplet, that has a silicon substrate 506 and a second layer portion 508. Within the integrated circuit 500, there may be an array section forming the module 502 where a three-dimensional column array of memory bit cells 522 has the necessary components to store memory in non-voltage, semi-volatile memory, or a memory format as described herein.

[00140] The integrated circuit 500 may include a modules group having a plurality of modules including a first module and a second module, etc. even though only a single module 502 is shown. The memory bit cells 522 are written to by the shared write port 512, 516, which includes both a write address bus line 512 and a write data bus 516. These buses run through the second layer 508 and can be connected to a second semiconductor device via an interposer. The second device has electrical contacts that complement those on the surface 518, allowing it to be electrically coupled to the write address and data buses. The memory bit cells 522 can be read from via the read port 524, 526, which includes a read address bus line 524 and a read data bus 526. Both of these buses can also run through the second layer 508 to the

second semiconductor device coupled to the surface 518, which also has complementary electrical contacts to allow it to be electrically coupled to the read address and data buses.

[00141] Various kinds of memory technologies may be used for the memory bit cells 522, such as a vertical connectivity fabric structure formed from non-volatile memory unit cells arranged in a three-dimensional column array 522. The memory bit cells 522 may utilize one or more of a cross-point, 3D NANDs, 3D NORs, 3D ANDs, and/or a stacked planar layer.

[00142] In some embodiments, the integrated circuit 500 is electrically connected to a second semiconductor device (not shown in Fig. 5) comprising another integrated circuit, which may be a system-on-chip or a Field-Programmable-Gate-Array. In some embodiments, the memory bit cells 522 may be formed from various non-volatile memory types, such as FeFET, FeRAM, ReRAM, SOT, or STT. Additionally, alternatively, or optionally, the memory bit cells may be formed from non-volatile memory unit cells having 2-terminal devices, 3-terminal devices, or 4-terminal devices.

[00143] For example, the memory unit bit cells 522 may be formed from ferroelectric materials, such as a ferroelectric tunnel junction, a diode, a capacitor, a single-gate transistor, or a dual-gate transistor. Alternatively, the memory unit bit cells 522 may be formed from memristive materials, such as at least one ReRAM, or magnetic materials, such as at least one spin-orbit-torque device or at least one spin-transfer-torque device. Moreover, the non-volatile memory unit cells 522 may also be formed from phase-change materials or anti-ferroelectric materials.

[00144] In some alternative embodiments, the non-volatile memory unit cells 522 can be formed from other types of materials, such as phase change materials, anti-ferroelectric materials, or multi-bit PCM materials. The non-volatile unit cells can be formed utilizing different structures, such as resistive random-access memory (RRAM) technology, magnetic random-access memory (MRAM) technology, or ferroelectric random-access memory (FRAM) technology.

[00145] Moreover, in some implementations, 3D NAND technology may be utilized to form the memory unit bit cells 522. For example, the memory unit bit cells 522 may be formed from stacked memory layers where each layer includes a plurality of memory cells that can be accessed using shared bit lines. In such a case, the read port 524, 526 may be coupled to the bit lines, and the write port 512, 516 may be coupled to the word lines that control the access to each layer.

[00146] In another embodiment, the 3D connectivity fabric structure can be built with stacked layers of either NAND gates, NOR gates, or AND gates, and in some cases, different

types of logic gates may be combined to optimize the structure's functionality. In addition, the 3D connectivity fabric structure may be formed utilizing through-silicon-via (TSV) technology, which allows the vertical interconnection of the different layers of the structure.

[00147] Additionally, the non-volatile memory unit cells may include 2-terminal devices, such as a capacitive or a memristive device with or without an additional selector device such as a diode in series, 3-terminal devices, such as a floating-gate transistor, a transistor with an access gate, or 4-terminal devices, such as a transistor with two access gates. The type and configuration of the non-volatile memory unit cells 522 may depend on the specific application requirements, including the speed, power consumption, and reliability of the circuit. The memory unit cell may include or be a single ferroelectric transistor or 6T SRAM cell. The memory unit cell may be a combination of many different devices, including, but not limited to, one or more of a transistor, a memristor, a capacitor, etc.

[00148] In some embodiments, a ferroelectric material can be utilized to form the non-volatile memory unit cells 522. The ferroelectric material may be implemented as any kind of device, including, but not limited to, a thin-film device, such as a ferroelectric tunnel junction, a capacitor, a single-gate transistor, or dual-gate transistors, etc.

[00149] In another embodiment, the non-volatile memory unit cells 522 may be formed from a memristive material, such as a Metal Oxide Memristor (MOM), Conductive-Bridging RAM (CBRAM), or valence change memory (VCM), each of which provides different benefits regarding power consumption, speed, endurance, etc.

[00150] Moreover, in some embodiments, the non-volatile memory unit cells 522 may be formed from a magnetic material, such as spin-orbit-torque (SOT) devices, spin-transfer-torque (STT) devices, or perpendicular magnetic tunnel junctions (p-MTJ).

[00151] In one embodiment, the modules group may include many modules where each of which can be accessed through dedicated read ports 524, 526 with a dedicated read peripheral 520 while sharing the same write port 512, 516 and shared write peripheral 510. The shared write port 512, 516 can be configured to selectively write to one or more of the plurality of modules within the modules group including the memory bit cells 522. Each of the modules may have the same or different sizes, and different module sizes may be configured to optimize the utilization of the memory array with different operating scenarios, etc.

[00152] Furthermore, the integrated circuit 500 may be formed utilizing different manufacturing processes and techniques, which include but not limited to, a CMOS or Bipolar-CMOS-DMOS (BCD) process, a silicon-on-insulator (SOI) process, a FinFET process, a silicon germanium (SiGe) process, a gallium arsenide (GaAs) process, etc.

[00153] In some embodiments, the where a three-dimensional column array of memory bit cells 522 forms is configured as a microvault. Additionally or alternatively, each of the microvaults will have a dedicated read peripheral 520 and a write peripheral 510 that is also a dedicated write peripheral rather than a shared write peripheral. That is, in some embodiments, each microvault includes a dedicated write connection and a dedicated read connection, predetermined number of microvaults (e.g., 2 or 4), may have a dedicated write connection and a dedicated read connection, with or without dedicated respective peripheries, etc.

[00154] Fig. 6 shows a perspective of an assembly 600 having the integrated circuit of Fig. 1 implemented on a semiconductor device, such as the chiplet 230, that is electrically connected to a system-on-a-chip ("SOC") 610 in accordance with an embodiment of the present disclosure. The semiconductor device, in this embodiment, is the chiplet 230 that is electrically connected to a system-on-a-chip ("SOC") 610.

[00155] Referring to Fig. 6, the SOC 610 includes a silicon substrate 602 on which a plurality of processing elements is formed, including a processing element 606. The processing elements can communicate with each other through a Network-on-Chip ("NOC") 604, which is a communication fabric that directs data transfer between the processing elements. The communication fabric can take various forms, including buses, switches, NOCs, etc. The NOC 604 in the SOC 610 directs data traffic between the various nodes (e.g., the processing element 606) and links, which provide the communication paths between the nodes.

[00156] The plurality of processing elements including the processing element 606 processing elements can be any suitable type of processors capable of executing instructions, including microprocessors, graphics processing units (GPUs), digital signal processors (DSPs), or application-specific integrated circuits (ASICs).

[00157] Additionally, the SOC 610 may comprise various modules, such as module 232, which are grouped together to provide memory functionality to the assembly 600 as described here. The modules in modules group 236 can be coupled to a respective processing element provide it readable memory. In some embodiments, the coupling between the module (e.g., module 232) and the processing element (e.g., 606) can be achieved through interconnects on the silicon substrate 602.

[00158] After the circuitry is formed on the silicon substrate 602, a second layer 608 can be disposed on top of the substrate. The second layer 608 can be any suitable material, such as an insulating material, a metal, a dielectric, or interconnect layer, and it may be bonded to the

chiplet 230. The bonding can be done using any suitable technique, including but not limited to, adhesives, soldering, or welding, etc.

[00159] In general, the assembly 600 provides a means of integrating the chiplet 230, which can include the integrated circuit of Fig. 1, with the SOC 610. Integrating the chiplet 230 provides various advantages, such as enhanced functionality, higher performance, and lower power consumption. Moreover, the integration of the chiplet 230 with the SOC 610 can be accomplished in various ways, depending on the particular application and design objectives of the system.

[00160] The assembly 600 can incorporate various variations and modifications, depending on the specific requirements of the system. For example, the processing elements formed on the silicon substrate 602 can vary in their number, type, and arrangement. Similarly, the modules in modules group 236 can vary in their number, type, and function.

[00161] Furthermore, the second layer 608 can be modified to include additional functionality. For instance, the second layer 608 can include passive components, such as resistors, capacitors, and inductors, or active components, such as transistors or diodes. Incorporating these components in the second layer 608 can further enhance the functionality and performance of the system.

[00162] In another variation, the assembly 600 can incorporate a heterogeneous integration approach, where the chiplet 230 is fabricated using a different technology than that used for the SOC 610. This approach allows for the optimal use of different fabrication technologies for different parts of the system, resulting in improved performance and reduced power consumption.

[00163] Fig. 7 shows a perspective view of an assembly 700 having a semiconductor device 707 with an array of processing elements 706 (on a grid of processing element 706a,a to 706n,n, where the first subscript is the columns and the second subscript is the row), and a second semiconductor device 709 having an array of microvaults 708. The array of microvaults 708 are on a grid of microvaults 708a,a to 708n,n, where the first subscript is the columns and the second subscript is the row). These subscripts may line up such that a respective subscript of a processing element 706 corresponds to a respective subscript of a microvault 708. The microvaults 708 are a type of module described herein where it is positioned in a vertical direction, e.g., above a respective processing element 706. The semiconductor device 707 may be a chiplet. Also, the semiconductor device 709 may also be a chiplet. The chiplets 707,709 may be bonded together. The different layers 710 (e.g., 710a through 710d can correspond)

may be allocated for separate AI models (e.g., parameters in a neural network, such as a CNN or transformer model).

[00164] The assembly 700 is the overarching structure that houses the various components shown in Fig. 7. It provides mechanical support and integration for the other elements, allowing them to function as a unified system.

[00165] The assembly 700 includes two semiconductor devices - the semiconductor device 707 and the semiconductor device 709. The semiconductor device 707 contains an array of processing elements labeled 706a,a to 706n,n. Similarly, the semiconductor device 709 contains an array of microvaults labeled 708a,a to 708n,n.

[00166] The subscripts a,a to n,n indicate that the processing elements 706 and microvaults 708 are arranged in a grid pattern, with the first subscript referring to the column and the second subscript referring to the row. This grid arrangement allows each processing element 706 to have a corresponding microvault 708 positioned vertically above it. For example, processing element 706a,a has microvault 708a,a above it, processing element 706b,b has microvault 708b,b above it, and so on. The alignment of the grid allows tight integration between the processing and storage components.

[00167] In some embodiments, the semiconductor devices 707 and 709 are potentially separate chiplets that are integrated using packaging techniques into the unified assembly 700. The chiplet form factor allows greater flexibility and customization in assembling the system. This arrangement is such that each processing element 706 can access its respective microvault 708 located above it to retrieve relevant data, such as weights for neural networks or AI models. This may provide high bandwidth and low latency access to the data needed for efficient processing.

[00168] Input data enters the system via the input DRAM memories 702. This data flows into the processing elements 706, where it is operated on locally using weights or parameters from the vertically integrated microvaults 708. The processing results output via the output DRAM memories 704. The input DRAM memories 702 consist of multiple individual DRAM modules labeled 702a, 702b, and 702c. The DRAM memories 702 can be any type of dynamic random access memory, including but not limited to DDR SDRAM, LPDDR SDRAM, GDDR SDRAM, and HBM. The DRAM memories provide high-bandwidth data input capabilities to feed data, such as inference inputs or training data, into the processing pipeline.

[00169] In some embodiments, each individual DRAM module 702a, 702b, and 702c has a dedicated interface and data path to each processing element 706. For example, DRAM module 702a may feed data only to processing element 706a,a, while DRAM module 702b

feeds data only to processing element 706b,b. This provides modular scalability, as additional DRAM modules can be added to feed more processing elements.

[00170] The number of input DRAM memories 702 and individual modules 702a-702c may vary depending on the application requirements. For instance, there could be 4, 8, 16, or more input DRAM modules. The capacity of each module can range from gigabytes to terabytes depending on factors such as access speed, power, and cost budget.

[00171] High-speed interfaces like DDR5, GDDR6, or HBM3 may be used to maximize data transfer bandwidth between the input DRAM memories 702 and the processing elements 706 across the semiconductor device 707. Shared data buses, crossbar switches, or on-chip networks may interconnect groups of DRAM modules 702 and processing elements 706.

[00172] In some implementations, the input DRAM modules 702 may be stacked or arranged in a multi-dimensional configuration to increase overall memory capacity and bandwidth while reducing latency and power consumption. Specialized memory controllers and schedulers may manage parallel data access across multiple input DRAM modules 702.

[00173] The input DRAM memories 702 supply the high-bandwidth data needs of the parallel processing elements 706, enabling fast and efficient data-intensive computations such as neural network inferencing. Each processing element 706 can directly access the required input data from its dedicated DRAM module 702 without contending with other processors for data access.

[00174] In alternative embodiments, instead of, or in addition to, the DRAM memories 702, 706, adjacent accelerator chiplets may be in communication with the semiconductor device 707. That is, there may be a grid-like arrangement of assemblies 700 in communication with each other to perform AI inference and/or AI training (e.g., transformer inference, CNN inference, ANN inference, etc.). In some embodiments, there may be clusters of semiconductor devices 707 that share a bank or portion of DRAM memories 702 and/or 706. In some embodiments, the input DRAM memories 702 and the output DRAM memories 704 may be combined into the same DRAM memory.

[00175] The assembly 700 includes a semiconductor device 707 that comprises an array of processing elements labeled from 706a,a to 706n,n. Each processing element in the array may be configured to execute specialized computations and data processing operations. For instance, in some embodiments, the processing elements could be optimized for artificial intelligence workloads like neural network inference. In other cases, the processing elements may focus more on general-purpose capabilities. Ultimately, the capabilities of each processing

element depend on its specific microarchitecture which can be tailored for certain applications if desired.

[00176] The processing elements 706 can access nearby memory storage to retrieve data that feeds into their computations. This memory may be physically separate from the processing element arrays 706 as is the case with the microvaults 708 shown in Figure 7. The processing elements 706 and microvaults 708 are aligned so that each microvault is positioned directly above its corresponding processing element in a vertical configuration. This tight coupling provides fast data transfer speeds between each vault-element pair.

[00177] In terms of physical implementation, the array of processing elements 706 resides within the semiconductor device 707. The semiconductor device 707 could potentially be manufactured as a standalone chiplet using advanced packaging techniques. This modular chiplet can then be integrated with other components like the microvault chiplet 709 through high-density interconnections. Some options include bumpless hybrid bonding, interposers, or even monolithic 3D integration. Ultimately, combining chiplets allows creating powerful heterogenous systems with optimized dies.

[00178] The specific number, design, and interconnect scheme of processing elements 706 present can vary between implementations of assembly 700. For instance, simpler systems may need only a 2x2 grid of elements whereas a sophisticated AI accelerator could feature a 32x32 array. The processing elements 706 themselves can also have different memory access routes across the assemblies. Point-to-point links, crossbar switches, or shared buses are possible connection structures. Such architectural decisions depend on the performance and area constraints trying to be met.

[00179] The second semiconductor device 709 is a separate device from the first semiconductor device 707. Like the first semiconductor device 707, the second semiconductor device 709 may also be implemented as a chiplet. The second semiconductor device 709 includes an array of microvaults 708, arranged on a grid spanning from microvault 708a,a to 708n,n.

[00180] As mentioned, the microvaults 708 on the second semiconductor device 709 are positioned vertically above the processing elements 706 on the first semiconductor device 707. Each microvault 708 lines up with and corresponds to the processing element 706 underneath it, based on the subscripts identifying their position in the grid. For example, microvault 708a,a is vertically aligned with and corresponds to processing element 706a,a. This allows each processing element 706 to access the microvault 708 above it.

[00181] The microvaults 708 act as a memory structure, storing things like AI model weights that can be accessed by the processing elements 706 underneath during operations like neural network inference. The microvaults 708 may be optimized for very fast read times but slower write times. This allows the processing elements 706 to rapidly access the weights and data needed for their computations, while less frequently updated data can still be written at a slower pace.

[00182] In some embodiments, the second semiconductor device 709 containing the array of microvaults 708 is directly bonded to the first semiconductor device 707 with the processing elements 706. This bonding aligns each microvault 708 with its corresponding processing element 706 underneath. Electrically conductive interconnects between the devices allow each processing element 706 to communicate directly upwards with its respective overlying microvault 708. This provides a compact, modular, and efficient system architecture.

[00183] The microvaults 708 may contain multiple memory layers, labeled 710a to 710d, each storing weights or data for a different AI model. For example, layer 710a contains the weights for model A, layer 710b contains the weights for model B, and so on. Stacking these layers vertically contributes to the high density and fast access times of the microvault design.

[00184] The microvaults 708 can be implemented using various memory technologies, including but not limited to SRAM, FeFET, ReRAM, SOT, and STT, optimized for fast readout times to supply data to the processing elements 706 with minimal latency. Specific embodiments may configure the microvaults 708 to have much faster read speeds compared to their write speeds. The microvaults 708 may be implemented using any FeFET or memory structure described herein.

[00185] In some embodiments, each microvault 708 may have a capacity between 4 kilobytes to 128 kilobytes for storing parameters for machine learning models or other data. The bit density per layer may exceed 0.4 gigabits per square millimeter. The microvaults' 708 compact size, between less than 100 micrometers on each side in one embodiment or 12 micrometers by 12 micrometers in another, allows high-density integration of the memory modules.

[00186] The array-based arrangement of the microvaults 708 may enable concurrent parallel data access by the processing elements 706, supporting high-throughput data processing by the assembly 700. The one-to-one alignment of the microvaults 708 and processing elements 706 also ensures that each processing element 706 has dedicated access to its required data without contention.

[00187] The microvaults 708 share the semiconductor device interface provided by the second semiconductor device 709, which facilitates writing data to the microvaults 708 from the input DRAM memories 702. Reading data from the microvaults 708 to the processing elements 706 and output DRAM memories 704 is handled through dedicated pathways between each vertically aligned microvault 708 and processing element pair.

[00188] The microvault memory layers 710 refer to multiple layers of microvault memories stacked vertically within the second semiconductor device 709. As illustrated in Fig. 7, there are four separate microvault memory layers labeled as 710a, 710b, 710c, and 710d. Each layer contains an array of microvaults, such as the array of microvaults 708 shown in the diagram.

[00189] In one embodiment, the microvaults 708 may utilize a stacked 3D NAND architecture built from multiple layers of NAND memory arrays using charge trap flash technology. Each microvault 708 may contain a dedicated set of wordline drivers on the bottom layer to facilitate access to the 3D NAND cell arrays above, spaced by alternating dielectric layers. The 3D NAND implementation may be used to maximize density and throughput by leveraging vertical scaling.

[00190] In an alternative embodiment, the microvaults 708 employ a 3D NOR architecture constructed from multiple tiers of NOR flash memory arrays. Each plane features NOR strings with a source line and bit line architecture, stacked on top of each other using vias. The 3D NOR arrangement optimizes random read access times to stored data.

[00191] The microvaults 708 may also adopt a hybrid configuration with different types of volatile and/or non-volatile memory, such as combining FeRAM and ReRAM cells, organized into vertical sub-arrays. This heterogeneous 3D integration allows optimizing for speed, endurance, and retention within the same vault structure.

[00192] In certain embodiments, the microvaults 708 integrate processing logic like analog computing directly into the memory array stack itself. This processing-in-memory approach places basic computational operators within the memory peripheral or bit cells, enabling highly parallel and efficient in-situ data processing.

[00193] Some implementations may utilize 2.5D or 3D stacking to integrate the microvaults 708 with other components like logic, CPUs, GPUs or application-specific accelerators. This tight packaging integration via techniques like high-bandwidth memory cube architectures reduces data transfer latency and power consumption.

[00194] The microvaults 708 may also employ a virtualized architecture, with an external memory controller handling translation between the physical array organization and

dynamically allocated virtual memory domains. These virtual domains mapped onto the physical array effectively creates separate virtual vaults with flexible capacities tailored to application needs.

[00195] In certain embodiments, the microvaults 708 are designed as Computational RAM (CRAM) with integrated processing capabilities within the bit cell peripheral to enable highly parallel in-memory computing architectures. Gateless transistor structures integrated into the CRAM arrays facilitate efficient execution of bulk bitwise operations.

[00196] Some implementations arrange the microvaults 708 into modular Memory Processing Unit (MPU) structures containing dedicated processing logic tailored for workloads like AI inferencing. The MPU architecture couples vault arrays to vector processors via high-speed interfaces like HBM2 enabling low-latency data transfers.

[00197] The microvaults 708 may also implement content-addressable capabilities by integrating comparison logic into the memory peripheral. This facilitates searching or accessing data based on content rather than explicit addresses, enabling powerful pattern matching capabilities.

[00198] Certain embodiments may stack multiple microvault dies on top of base logic dies featuring things like GPUs or AI accelerators. This creates dense, high-bandwidth heterogeneous systems optimized for data-centric workloads while minimizing data movement.

[00199] The microvault memory layers 710 may be fabricated utilizing three-dimensional integrated circuit manufacturing processes to stack multiple dies or wafers containing microvault 708 arrays on top of each other. Through-silicon vias (TSVs) or other vertical interconnect technologies can be employed to enable communication between the layers.

[00200] In some embodiments, each microvault memory layer 710 corresponds to a different artificial intelligence (AI) model or application. For example, layer 710a could store the weights and parameters for AI model A, layer 710b could store the weights and parameters for AI model B, and so on. This allows multiple AI models to be stored efficiently within the same microvault memory 708 structure.

[00201] The microvaults 708 may possess capacities ranging from 4 kilobytes to 128 kilobytes in some embodiments. In other cases, the capacity could be between 4 kilobytes to 16 kilobytes. Each microvault could have lateral dimensions less than 100 micrometers by less than 100 micrometers, while extending vertically to incorporate potentially over 200 memory cell layers in some implementations.

[00202] In some embodiments, the bit density of per square millimeter per layer within the microvault memory layers 710 may facilitate high-capacity storage with a small footprint. The layers may utilize non-volatile memory technologies, such as FeFET, STT-MRAM, or ReRAM, to retain data when power is removed.

[00203] In operation, the processing elements 706 may access weights or parameters from the microvault memory layers 710 to perform neural network inferencing or other machine learning computations.

[00204] The inference results of the various AIs may be sent to the output DRAM memories 704, which comprise individual DRAM memory modules labeled 704a, 704b, and 704c. The output DRAM memories 704 are positioned adjacent to the array of microvaults 708 and the second semiconductor device 709. The output DRAM memories 704 may serve as temporary data storage that can buffer output data retrieved from the microvaults 708 before it is transmitted externally.

[00205] Each DRAM memory module 704a, 704b, and 704c may have similar or different storage capacities, depending on the design requirements. For example, in one embodiment, each module contains 16 megabits of storage. The DRAM storage cells utilize a capacitor to retain data bits in the form of electrical charges. Due to charge leakage, the DRAM memories require periodic refresh cycles to maintain the stored data integrity. To enable concurrent reads and writes across multiple modules, each DRAM module 704a, 704b, and 704c can have dedicated internal control circuitry and I/O ports.

[00206] The data outputs from the individual microvaults 708 may get aggregated and buffered in the output DRAM memories 704 before being transmitted to external components via peripheral circuitry. Buffering the data allows the transmission rate to be regulated to match the requirements of the external interfaces. It also enables data processing operations like formatting, encoding, or encryption to be performed by the second semiconductor device 709 prior to output.

[00207] In some embodiments, the output DRAM modules 704a, 704b, and 704c are designed to provide high-density, low-cost temporary data storage to support the high-bandwidth parallel reads from the array of microvaults 708. Optimizing these performance parameters allows efficient extraction of data from the microvaults to feed the computational workflows hosted on external chips or devices. Specific implementations may utilize various types of DRAM, including asynchronous DRAM, synchronous DRAM, graphics DRAM, and low-power DRAM tailored to the application. Overall, the output DRAM memories 704 facilitate seamless data movement from the integrated microvaults to external execution pipelines.

[00208] Fig. 8 shows an assembly 800 of a semiconductor devices 802, 804, 806, 808, 810, 812, 814 including several memory types in accordance with an embodiment of the present disclosure. Specifically, Fig. 8 shows an exemplary assembly 800 of semiconductor devices 802, 804, 806, 808, 810, 812, 814 configured to provide a hierarchical memory structure. The assembly 800 is modular and scalable, allowing for various combinations and numbers of semiconductor devices, which may be implemented as chiplets in certain embodiments, to be stacked to meet specific performance and density requirements.

[00209] The assembly 800 includes an application semiconductor 802, that may be a plurality of processing elements as described herein. On top of the semiconductor device 802 is semiconductor device 814 that includes an array of microvaults. One top of the semiconductor device 814 is semiconductor device 812, which may also include an array of microvaults. On top of the semiconductor device 812, are semiconductor devices 810, 808, which may be SRAM vaulted dies. One top of the semiconductor device 808, is semiconductor devices 806, 804 may be DRAM vaulted dies. These vaults 816 may be arranged in a grid-like fashion such that 816a,a to 816n,n subscripts the vaults. Each of these vaults may include a respective microvault from semiconductor devices 804, 812, respective SRAM vaults from semiconductor devices 810, 808, and respective DRAM vaults from semiconductor devices 806, 804.

[00210] At the base of assembly 800 lies the semiconductor device 802, which comprises a plurality of processing elements. These processing elements execute computational tasks and facilitating data flow within the system.

[00211] Directly above semiconductor device 802 is semiconductor device 814, which includes an array of microvaults. These microvaults utilize Field-Effect Transistors (FeFETs) known for their non-volatile characteristics and suitability for high-density memory applications. The FeFET-based microvaults may be designed to enable high-speed read operations essential for rapid data retrieval during processing tasks, such as AI inferencing, while supporting slower write operations that are more tolerant to latency. Stacked on top of semiconductor device 814 is semiconductor device 812, which similarly includes an array of microvaults. The presence of multiple layers of microvaults in semiconductor devices 814 and 812 exemplifies the scalable nature of the assembly, where additional memory capacities and functionalities can be integrated through additional layers.

[00212] Further contributing to the memory hierarchy, semiconductor devices 810 and 808, positioned above semiconductor device 812, are depicted as SRAM vaulted dies. SRAM

provides fast access memory that can serve as a cache or buffer to the slower, but denser, FeFET microvault memory layers beneath.

[00213] At the top of the assembly 800 and hence the memory structure are semiconductor devices 806 and 804, illustrated as DRAM vaulted dies. DRAM is typically used for main memory due to its relatively high speed and low cost per bit compared to SRAM, offering a balance between performance and economy.

[00214] The arrangement of vaults 816 in a grid-like fashion, subscripted from 816a,a to 816n,n, indicates that each processing element at a given location (a,b) in semiconductor device 802 has dedicated access to the corresponding vertically aligned vaults of microvaults and memory cells in the layers above. This vertical stacking and alignment ensure that data and control signals can be directly routed between processing elements and their respective memory stacks, facilitated by interconnect technologies such as through-silicon vias (TSVs) and micro-bumps, which are used in the assembly's 800 3D integrated circuit architecture.

[00215] The modular and scalable design of assembly 800 allows for various combinations of semiconductor devices or chiplets to be integrated into more extensive systems. The flexibility in the number and combination of stacks provides the adaptability to tailor the assembly to the requirements of different applications and performance demands. Each vault within the vaults 816 within the assembly 800 presents a multi-die structure that contributes to the overall capacity and performance of the system.

[00216] Fig. 9 shows an assembly of semiconductor devices including a semiconductor device with a system-on-chip 914 and another semiconductor 906 with microvaults disposed on top in accordance with an embodiment of the present disclosure. The assembly 900 integrates a semiconductor device 906 and a semiconductor device 914, which may be implemented as separate chiplets bonded together. The semiconductor device 914 includes various components to facilitate reading data from the microvaults on the semiconductor device 906, such as a read address register input interconnect 924, read data register 922, and read data register output interconnect 950. These components pass the read address to the microvaults on semiconductor device 906 and return the read data back to the semiconductor device 914.

[00217] Specifically, the read address enters via interconnect 924 into the read address register 926. The output of this register 926 connects through interconnects and bumpless bonds to another read address register 938 on the semiconductor device 906, which then addresses the target microvault 936. The microvault 936 outputs read data via interconnect 940 to a read data register 942, which passes the data back through bumpless bonds 910, 918 to

read data register 922 on semiconductor device 914. This data can then be accessed externally via the read data register output interconnect 950. Additionally, the semiconductor device 914 and 906 have interconnected Through-Silicon Vias 916 and 944 to allow communication with devices potentially stacked above semiconductor device 906.

[00218] The semiconductor device 906 features various memory structures to provide data storage capabilities. This includes a microvault 936, which offers high-density, low-latency data storage, along with other peripheral memory components like read data register 942 and read address register 938 to facilitate data reads. The microvault 936 resides on the BEOL portion of the chiplet, allowing dense 3D integration of memory layers. In some implementations, the microvault utilizes non-volatile memory technologies like FeFET or STT-MRAM for data retention without power.

[00219] The semiconductor device 914 comprises processing elements and data routing circuitry to retrieve and manipulate data stored in semiconductor device 906. Components like read address register 926 and read data register 922 handle sending read addresses and receiving data from the microvault 936 respectively. The device 914 also includes interconnects 924, 950 and Through-Silicon Via 916 to communicate externally.

[00220] The two devices 906 and 914 integrate via fine-pitch interconnects like bumpless hybrid bonds 908, 910, 918, 920, 930 and 932. This allows direct data transfer pathways between processing components in device 914 and memory structures in device 906. Alignment during bonding ensures dedicated access - for instance, read data register output interconnect 950 on 914 links directly to read data register 922 to receive requested data.

[00221] The pathways facilitating data flow during reads can be summarized as follows: A read address enters through interconnect 924 into read address register 926 on device 914. This gets communicated via interconnects and bumpless bonds to read address register 938 on device 906, which then addresses microvault 936. Requested data gets passed via interconnect 940 to read data register 942, then transfers through bonds back to read data register 922 on 914, where it becomes available externally via interconnect 950.

[00222] The assembly 900 exemplifies a modular, high-density architecture optimized for data-centric applications like AI inferencing. Tight integration of processing and storage dies via advanced packaging techniques allows localized data access with minimal latency and power. Scalability is also enabled by incorporating multiple chiplets, in this case devices 906 and 914. The assembly 900 illustrates a potential configuration suited for space-constrained, high-performance computing systems.

[00223] The Through-Silicon Via (TSV) 916 is an electrical connection that passes vertically through the semiconductor device 914. Its purpose is to provide a pathway for signals to travel between the top and to a processing element within the semiconductor device 914. This allows the device to be stacked and interconnected with other components in a vertical configuration. The TSV 916, along with other TSVs on the device, facilitates high-density 3D integration and heterogeneous stacking of multiple devices like chiplets.

[00224] The TSV 916 interacts with several other components within the system. On the top side of semiconductor device 914, it connects to interconnect 912, which couples it to bumpless bonds 918. These bonds interface with complementary bumpless bonds 910 on the bottom side of semiconductor device 906 when the two devices are stacked. This allows signals to travel from device 914 to device 906 through the TSV 916. The route continues as signals go through interconnect 902 to TSV 944 on device 906. TSV 944 provides a vertical signal pathway to the top surface of device 906 where additional devices could be stacked. In the reverse direction, signals can travel from TSV 944 down through device 906, back up TSV 916, and down into device 914. So the TSV 916 provides bidirectional vertical communication across device boundaries.

[00225] There are a few possible variations for the TSV 916 implementation. First, multiple TSVs arranged in an array could be used instead of a single via to increase throughput and redundancy. Second, the dimensions and materials of the TSV could be optimized- for example, smaller TSV diameters using denser materials like tungsten could be advantageous. Additionally, the interface circuitry driving signals into the TSV, like interconnects 912 and 902, could employ variable line drivers to support different voltage levels or signal integrity enhancements. Further embodiments may include integrated monitoring circuitry within TSV 916 to track metrics like temperature and link utilization. And alternative signaling schemes besides electrical signals could be employed in future cases. For instance, integrated silicon photonics utilizing modulated light to convey data through the TSVs could enable very high bandwidth and low latency connectivity. There are multiple avenues to further develop the capabilities of TSV-based vertical links like TSV 916 within these complex 3D integrated architectures.

[00226] There are several variations and alternatives for the interconnect 912 implementation. For example, different conductive materials such as copper or aluminum may be utilized to fabricate the pathways forming interconnect 912 and optimize for conductivity or thermal dissipation. Additionally, interconnect 912 may feature redundant signal paths or self-repair capabilities using spare interconnect lines to improve reliability and resilience. The

bumpless bonds 918 and 910 connecting devices 906 and 914 could also be replaced with other high-density bonding approaches like hybrid bonding or Through-Silicon Vias. Furthermore, alternate signaling schemes besides simple digital logic could be employed on interconnect 912, such as analog signaling or multi-level digital waveforms to enhance data transmission capabilities. The routing and dimensions of interconnect 912 can also be adapted according to bandwidth requirements or circuit layout considerations. Overall, many structural and functional alternatives exist for crafting interconnect 912 to meet application needs.

[00227] The bumpless bonds 918 are electrical connections located on the semiconductor device 914 between an interconnect 912 and bumpless bonds 910 of the semiconductor device 906. The bumpless bonds 918 provide an electrical pathway for signals to travel between the semiconductor device 914 and any additional semiconductor devices, such as the semiconductor device 906, stacked on top of the assembly 900. The signals communicated over the bumpless bonds 918 can include data signals, control signals, address signals, or any other signals needed to coordinate operations between the multiple semiconductor devices.

[00228] There are several possible variations for the bumpless bonds 918. The number of individual bond sites can range from just a few to hundreds, depending on signal bandwidth requirements. The bonding method can utilize techniques like direct bonding, plasma-activated bonding, adhesive bonding, or compression bonding. Hybrid bonding approaches are also possible, combining direct wafer bonds with intermediate metal bonds. The size and pitch of each bond site can vary and may use pitches under 10 micrometers to enable high-density connections. Redundant bonds can provide backup pathways. Shielding structures may surround bonds for noise immunity. Overall, many embodiments of bumpless bonds 918 are possible to meet cost, reliability, and performance needs.

[00229] The bumpless bonds 910 provide an interface for communicating signals between the semiconductor device 906 and semiconductor device 914. Specifically, the bumpless bonds 910 of the semiconductor device 906 are electrically coupled to the complementary bumpless bonds 918 of the semiconductor device 914. This allows signals like read/write data and addresses to be transmitted between the two devices. The bumpless nature of the bonds allows for a low-profile, high-density interconnection.

[00230] The bumpless bonds 910 interact with other components in the system to facilitate data transfer operations. For writes, data enters the semiconductor device 914 via the Through-Silicon Via 916, passes through interconnect 912 and bumpless bonds 918 before reaching bumpless bonds 910 of device 906. For reads, addresses flow from the read address

register 926 of device 914 through interconnects 928, 930 and bumpless bonds 932 into the read address register 938 on device 906. Read data then returns through bumpless bonds 908 and 920 back to device 914. So the bumpless bonds 910 provide key data and address routing between the devices.

[00231] Possible variations of the bumpless bonds 910 include using different bond densities, materials, or electrical contact configurations to optimize performance. The bonds can use alloying or doping techniques to improve conductivity. Additionally, the routing of signals can be changed, for example by using separate ports for input and output instead of shared ports. More bumpless bonds can be added to increase bandwidth between devices. Shielding may be added around the bonds to reduce interference. Overall, many modifications to the bumpless bonds 910 are possible within the scope of electrically interconnecting multiple devices.

[00232] The Through-Silicon Via (TSV) 944 is an electrical connection that passes vertically through the semiconductor device 906 from the top surface to the bottom surface. Its purpose is to facilitate communication of signals and data between the semiconductor device 906 and any additional semiconductor devices potentially stacked on top of it in a 3D integrated circuit configuration. The TSV 944 enables high-density interconnections between multiple stacked semiconductor layers, providing an efficient means for data routing and signaling.

[00233] The TSV 944 interfaces with surrounding circuitry within the semiconductor device 906, allowing signals to be transmitted upwards or downwards depending on the system configuration. On one end, the TSV 944 couples to the read data register 942 via interconnect 946. The read data register 942 can use the TSV 944 path to transfer read data from the microvault 936 to external semiconductor devices. This enables efficient data offloading from the on-chip memory. On the other end, the TSV 944 continues through to the top surface of semiconductor device 906, where it may interface with complementary contacts or interconnects on the bonded semiconductor above it. This facilitates the vertical transfer of signals and data along the assembly 900.

[00234] There can be many variations in the specific implementation of the TSV 944. Its dimensions can range from a few microns to tens of microns to match pitch requirements. The TSV 944 can be tapered, straight, or have non-uniform cross-sections. It may utilize different conductive materials as liners and fills, including metals like copper, tungsten or alloys. Insulating liners made of materials like silicon dioxide can separate the conductive fill from the substrate. The contacts and interconnects coupling into the TSV 944 can also have diverse layouts. Multiple TSVs can be placed adjacent to each other in a high-density array

configuration if desired. Overall, many architectural optimizations in the design and fabrication process of the TSV 944 are possible within the scope of the present disclosure.

[00235] Fig. 10 shows a semiconductor assembly 1000 incorporating a daisy-chained configuration of microvaults 1036, 1058, operatively connected to a multiplexer 1060 and managed by a counter 1062 for coordinated data selection and retrieval, in accordance with an embodiment of the present disclosure. This assembly 1000 is designed to carry out data processing tasks, potentially for applications such as artificial intelligence (AI) and machine learning, where high-speed data access and processing are utilized.

[00236] The assembly 1000 comprises two primary semiconductor devices: semiconductor device 1006 and semiconductor device 1014. Semiconductor device 1014 is depicted as containing several interfaces and registers for data communication, including a read address register input interconnect 1024. This interconnect 1024 facilitates the delivery of read addresses to a read address register 1026, which temporarily holds these addresses before they are transmitted to corresponding microvaults 1036, 1058 in semiconductor device 1006 for data retrieval operations.

[00237] In semiconductor device 1014, interconnect 1028 serves as a pathway for read addresses from the read address register 1026 to transition to bumpless bonds 1030. Bumpless bonds 1030 and 1032 represent high-density, low-profile electrical connections between semiconductor device 1014 and semiconductor device 1006, ensuring the transmission of read addresses with minimal signal loss and physical space requirements.

[00238] Through interconnect 1034, the received addresses reach the read address register 1038 in semiconductor device 1006, which then directs the microvault 1036 to output the requested read data. Microvault 1036, an memory storage unit, may encompass a variety of memory technologies, such as FeFETs and/or 3D-NAND structures as described herein, to facilitate the storage and rapid retrieval of data.

[00239] The multiplexer 1060 selects the appropriate data stream from multiple microvault 1036, 1058 outputs. Controlled by a counter 1062, which may operate according to a predefined sequence or be driven by external control signals, the multiplexer 1060 arbitrates between the outputs of microvault 1036 and another microvault, denoted as microvault 1058. Microvault 1058, similar in function and potential memory technology to microvault 1036, provides an additional source of data for the multiplexer 1060 to select from.

[00240] Once the desired data is selected by the multiplexer 1060, it is temporarily stored in a read data register 1042, also located within semiconductor device 1006. This register 1042 acts as a buffer, holding the data for subsequent processing or transmission. The read data

is then routed via interconnect 1004 to bumpless bonds 1008, which facilitate the transfer of data to the semiconductor device 1014.

[00241] Bumpless bonds 1010, 1018 facilitate the continued data's journey through the assembly 1000, ensuring data transfer from semiconductor device 1006 to semiconductor device 1014. Once the read data arrives at semiconductor device 1014, it is channeled via interconnect 1048 to a read data register, specifically read data register 1022, where it can be accessed by external systems, such as an application-specific integrated circuit (ASIC) or a system-on-chip (SoC), via the read data register output interconnect 1050.

[00242] Additionally, the assembly 1000 encompasses Through-Silicon Vias (TSVs) 1016 and 1044, providing vertical electrical connections through the semiconductor devices 1014 and 1006, respectively. These TSVs enable the stacking of additional semiconductor devices or chipllets atop the assembly 1000, thus allowing for vertical expansion of the system's capabilities. Interconnects 1012 and 1046 serve as horizontal pathways for signals to travel to and from the TSVs 1016 and 1044, respectively.

[00243] Although the description presents a specific configuration, the assembly 1000 may be subject to various modifications and alternative embodiments. For example, the number and arrangement of microvaults, the specific types of memory technologies employed within the microvaults, and the configuration of interconnects and bonding areas may be tailored to meet the requirements of different applications.

[00244] In some embodiments, the semiconductor devices 1006 and 1014 may be designed to accommodate additional functionality, such as thermal management layers for heat dissipation, hardware-based encryption modules for data security, or power management circuits to optimize energy consumption. The detailed structure of Fig. 10, therefore, serves as a foundation upon which a variety of sophisticated semiconductor systems can be constructed, each tailored to the specific needs of its intended application.

[00245] In the configuration of assembly 1000 as depicted in Fig. 10, the microvaults, exemplified by microvaults 1036 and 1058, present a daisy-chaining configuration that allows for an expandable and flexible memory architecture within the semiconductor device 1006. This daisy-chaining is facilitated through a series of interconnected pathways and controlled by the multiplexer 1060 in coordination with the counter 1062.

[00246] Each microvault, such as 1036 and 1058, is designed to hold and provide rapid access to data, which may be in the form of stored charge, magnetic states, ferroelectric material states, or other physical embodiments of binary information. The microvaults are interconnected such that the output of one microvault can be routed to the input of another,

creating a chain of memory elements. This is achieved through a series of interconnects, such as interconnect 1034 for microvault 1036 and interconnect 1056 for microvault 1058, which serve as conduits for the read data signals emanating from the microvaults.

[00247] The multiplexer 1060 manages the flow of data from this daisy chain of microvaults. It is designed with multiple inputs, each connected to the output of a microvault via respective interconnects. In the example provided, interconnect 1040 carries the read data from microvault 1036, and interconnect 1056 carries the read data from microvault 1058 to the multiplexer 1060. The multiplexer 1060 is capable of selecting which input to connect to its output at any given time, thus controlling which microvault's data is forwarded to the read data register 1042.

[00248] The counter 1062 orchestrates the operation of the multiplexer 1060. It may be a binary counter or any form of sequential logic circuit that produces a series of output states in response to a clock signal. The counter 1062 progresses through its states with each tick of the clock, which may be provided by an external clock source or generated internally within the semiconductor device 1006. As the counter 1062 advances, it outputs a control signal that instructs the multiplexer 1060 on which input to select.

[00249] For instance, at the first clock pulse, the counter 1062 may instruct the multiplexer 1060 to connect the output from microvault 1036 to the read data register 1042. At the next clock pulse, the counter may switch the connection to microvault 1058's output, and so on, cycling through the available microvaults in a predefined order. The counter's sequence and timing can be configured based on the desired data access patterns and the specific requirements of the processing tasks at hand.

[00250] This clock-driven coordination allows for an efficient and organized retrieval of data from a potentially large array of microvaults. It ensures that each microvault has an equal opportunity to present its data for processing, and it simplifies the control scheme by reducing it to a predictable, rhythmic progression of states. This is particularly advantageous in systems where a large volume of data must be processed in parallel, as it provides a systematic method for accessing and utilizing the stored information.

[00251] It should be noted that while Fig. 10 illustrates only two microvaults, the described daisy-chaining mechanism can be extended to accommodate any arbitrary number of microvaults. Additional microvaults can be added to the chain, with each new microvault connected to the multiplexer via an additional input line. The multiplexer 1060 and counter 1062 would be scaled accordingly to manage the increased number of inputs, maintaining the same clock-driven, sequential data retrieval process across the expanded memory architecture.

[00252] This daisy-chaining of microvaults, in conjunction with the multiplexing and counter-driven control system, exemplifies a modular and scalable approach to memory design in semiconductor devices. It allows for the customization of memory arrays to match the capacity and performance needs of a wide range of applications, from embedded systems to large-scale data centers, providing a versatile solution for modern computing challenges.

[00253] Fig. 11 shows a semiconductor assembly 1100 incorporating a daisy-chained configuration of microvaults in multiple semiconductor devices 1106, 1116 that are operatively connected to multiplexers and managed by counters for coordinated data selection and retrieval, in accordance with an embodiment of the present disclosure.

[00254] The present detailed description relates to Fig. 11 of the accompanying drawings, which illustrates an embodiment of an assembly 1100 as part of an integrated circuit. The assembly 1100 may be seen as a hierarchical structure that includes a bottom semiconductor device 1122, a middle semiconductor device 1116, and a top semiconductor device 1106 (each of these may be chiplets). Each semiconductor device is configured to interface with the others through a series of bumpless bonds, such as bumpless bonds 1128, 1129, 1158, 1159, 1160, 1161, 1162, and 1163, which facilitate electrical connectivity without the added profile of traditional bonding methods, thus enabling a compact and dense stacking of semiconductor layers.

[00255] The bottom semiconductor device 1122 includes a read address register 1126, which may be configured to store and communicate read addresses to microvaults located across the assembly 1100. The read address register 1126 communicates via an interconnect 1174, which serves as a conduit for signals directed to bumpless bonds 1128. These bonds in turn engage with bumpless bonds 1129 of the middle semiconductor device 1116, thus transferring the read addresses into the middle semiconductor device. The bottom semiconductor device 1122 also comprises a read data register 1124, which may serve as a repository for read data received. Read data is received through an interconnect 1176 that connects to bumpless bonds 1158, which are in communication with bumpless bonds 1159 of the middle semiconductor device 1116.

[00256] The middle semiconductor device 1116 serves as an intermediary layer within the assembly 1100, housing microvaults such as microvault 1132 and microvault 1118, each of which may be designed to store and rapidly provide access to data. These microvaults are linked to other components within the device via interconnects, such as interconnect 1130 and interconnect 1134, which guide the flow of read addresses and read data, respectively. The middle semiconductor device 1116 also features a read address register 1146, which receives

read addresses from interconnect 1130, and a read data register 1154, which collects read data from TSV 1152.

[00257] The multiplexer 1114, within the middle semiconductor device 1116 selects between various data streams. This multiplexer is controlled by a phase counter 1110, which determines the sequence of data selection based on the input received from the top semiconductor device 1106 via the TSV 1152. The read data register 1112 serves as a holding area for the selected data stream from the multiplexer 1114.

[00258] The top semiconductor device 1106 features a phase counter 1102 and a read data register 1104, which are used for coordinating data flow within the assembly 1100. The phase counter 1102, in conjunction with multiplexer 1150, dictates the output of read data from microvaults 1138 and 1164 based on the selected phase. The read data register 1104 captures the output from the multiplexer 1150, which is then relayed through interconnect 1108 to bumpless bonds 1162, facilitating communication with the middle semiconductor device 1116.

[00259] The assembly 1100 illustrates the integration of multiple microvaults across different semiconductor devices. For example, the microvault 1132 on the middle semiconductor device 1116 may receive read addresses from read address register 1146 via an interconnect 1134, path "1". Similarly, microvault 1118 may receive read addresses from the same register via an interconnect 1156, path "2". Microvaults 1138 and 1164 on the top semiconductor device 1106 receive read addresses via interconnect 1140 and paths "3" and "4", respectively, from read address register 1142, which is in communication with the middle semiconductor device 1116 via TSV 1136 and bumpless bonds 1160 and 1161.

[00260] Each microvault, such as 1132, 1118, 1138, and 1164, can potentially output read data to the multiplexer 1114 or 1150, where the data is then selected based on the configuration of the respective phase counter, 1110 or 1102. The selected data is temporarily stored in a read data register, either 1112 or 1104, before being transmitted down the assembly 1100 through the respective bumpless bonds and interconnects, ultimately reaching the read data register 1124 of the bottom semiconductor device 1122. This arrangement allows for a synchronized read-out of data from all microvaults, which may be essential in applications requiring parallel processing and high-speed data access.

[00261] The TSVs, such as 1136 and 1152, provide vertical connectivity across the semiconductor devices, enabling the integration of additional layers or functionalities atop the existing assembly 1100. These TSVs are coupled to various interconnects and bumpless bonds that establish the necessary pathways for signal transmission both within and between the semiconductor devices.

[00262] In some embodiments, the microvaults within the assembly 1100 may include various memory technologies, such as 3D-NAND or 3D-NOR structures, and are arranged to facilitate parallel processing and efficient data retrieval. Each microvault may include additional features, such as thermal management layers for heat dissipation, hardware-based encryption modules for data security, or power management circuits to optimize energy consumption.

[00263] Datapath 1 within the assembly 1100 exemplifies a route through which read addresses and corresponding read data are transmitted across the assembly, specifically directing operations from the read address register 1126 located on the bottom semiconductor device 1122 to the read data register 1124 within the same device.

[00264] The process begins with the read address register 1126 holding a specific read address. This address is sent through interconnect 1174, which acts as a channel for the signal. The read address is then transmitted to bumpless bonds 1128, which are meticulously designed to create a reliable electrical connection without the physical protrusion associated with traditional bonding methods. These bonds ensure a low-profile interface that preserves the compactness of the semiconductor stack.

[00265] The signal continues from bumpless bonds 1128 to engage with bumpless bonds 1129 of the middle semiconductor device 1116. The read address is carried forward by interconnect 1130, which delivers the address to the read address register 1146 of the middle semiconductor device. Read address register 1146, in turn, propagates the read address through interconnect 1134, designated as path "1" guiding the signal to the microvault 1132.

[00266] Upon receiving the read address, microvault 1132 accesses the requested data. This data is then outputted through interconnect 1170 and directed to the multiplexer 1114. In some embodiments, multiplexer 1114 functions as a selective switch that chooses between data streams based on the configuration determined by phase counter 1110. This phase counter may be designed to cycle through a sequence that dictates the timing and selection of data streams, ensuring that each microvault is read in a coordinated manner.

[00267] The selected data from multiplexer 1114 is then captured by read data register 1112, which holds the data momentarily. The data is subsequently sent via interconnect 1120, which carries the signal to bumpless bonds 1159. These bonds are part of a sophisticated electrical interconnection system that, along with bumpless bonds 1158 on the bottom semiconductor device 1122, enables vertical and horizontal integration within the semiconductor stack.

[00268] The signal, now in the form of read data, traverses from bumpless bonds 1159 to bumpless bonds 1158 and is finally introduced into interconnect 1176. This interconnect completes the connection to read data register 1124, which is configured to receive and hold the read data. The read data register 1124 may be equipped to retain the data for subsequent processing or external communication.

[00269] Datapath 2 within the assembly 1100 delineates a route specifically designed for the transmission of read addresses from the read address register 1126 on the bottom semiconductor device 1122 to the microvault 1118 located on the middle semiconductor device 1116 and the subsequent transfer of read data back to the read data register 1124 on the bottom device.

[00270] The journey commences at the read address register 1126, where a read address is held in preparation for dispatch. This register is an part of the semiconductor device's control mechanism, orchestrating the retrieval of data by issuing specific addresses to the memory units. From the read address register 1126, the read address is sent through the interconnect 1174, which provides a secure and reliable pathway for electrical signals within the integrated circuit.

[00271] The read address continues from the interconnect 1174 to bumpless bonds 1128, which offer a seamless and low-profile connection to the middle semiconductor device 1116 via the corresponding bumpless bonds 1129. These bonds maintain the signal's integrity during inter-layer communication and are designed to accommodate the requirements of modern semiconductor architectures.

[00272] The signal is then channeled via interconnect 1130 to the read address register 1146 within the middle semiconductor device 1116. The read address register 1146 acts as a secondary store and hold register. From there, the read address is directed down interconnect 1156, labeled as path "2," which terminates at the microvault 1118.

[00273] Upon receipt of the read address, microvault 1118 accesses the corresponding data. This data retrieval process is facilitated by the microvault's internal architecture, which may comprise an array of memory cells optimized for rapid access and data stability. The read data is outputted from microvault 1118 and travels through interconnect 1172, which leads to the multiplexer 1114.

[00274] Multiplexer 1114 determines which data stream to forward based on input from the phase counter 1110. The phase counter 1110 operates in synchronization with the system clock or an external control signal, cycling through various states to control the selection process of the multiplexer 1114 in a precise and predictable manner.

[00275] The output of the multiplexer 1114, now carrying the selected read data, is conveyed to read data register 1112. This register temporarily stores the read data, acting as a buffer. The read data is then dispatched via interconnect 1120 towards bumpless bonds 1159.

[00276] Bumpless bonds 1159 form the interface with bumpless bonds 1158 on the bottom semiconductor device 1122, where the signal is transmitted downward through the assembly. The read data then traverses the interconnect 1176 to reach its final destination, the read data register 1124. The read data register 1124 captures the read data, holding it in readiness for further processing or transmission to external circuits.

[00277] Datapath 3 within the assembly 1100 is another communication route that illustrates the data transfer sequence from the read address register 1126 on the bottom semiconductor device 1122, through various components, ultimately to the microvault 1138 on the top semiconductor device 1106, and then back to the read data register 1124 on the bottom device.

[00278] The sequence initiates at the read address register 1126, which serves as the origin point for read addresses. The register 1126 securely holds the address before it is dispatched through the interconnect 1174. Interconnect 1174 acts as a dedicated channel, ensuring that the read address is conveyed with precision to the bumpless bonds 1128. These bumpless bonds 1128 facilitate a streamlined connection to the bumpless bonds 1129 of the middle semiconductor device 1116, preserving the integrity and compactness of the signal pathway.

[00279] Upon reaching the middle semiconductor device 1116, the read address is relayed through interconnect 1130 to the read address register 1146. The read address register 1146 acts as a juncture that further propagates the address signal through the Through-Silicon Via (TSV) 1136. The TSV 1136 is a vertical interconnect that pierces through the semiconductor substrate, providing a direct link from the middle semiconductor device 1116 to the top semiconductor device 1106, thus exemplifying the 3D integration capabilities of semiconductor design.

[00280] The read address ascends from TSV 1136 and emerges onto bumpless bonds 1160 on the middle semiconductor device 1116. The bumpless bonds 1160 are connected to the bumpless bonds 1161 of the top semiconductor device 1106. The address signal is conducted through interconnect 1140 to the read address register 1142 on the top device 1106.

[00281] The read address register 1142, upon receiving the read address, directs the signal along interconnect 1144. This path, denoted as "3," leads the address to the microvault 1138. Microvault 1138, designed for data storage, retrieves the requested information in

response to the read address. The read data is outputted through interconnect 1166, which feeds the data into the multiplexer 1150.

[00282] The multiplexer 1150 in the top semiconductor device 1106 is governed by the phase counter 1102, which dictates the selection of the data stream to be channeled to the read data register 1104. The chosen data stream is temporarily housed in the read data register 1104, where it awaits downstream transmission.

[00283] The read data departs from the read data register 1104 via interconnect 1108, which connects to bumpless bonds 1162. These bonds 1162 engage with the corresponding bumpless bonds 1163 on the middle semiconductor device 1116, transferring the read data to TSV 1152.

[00284] TSV 1152 operates as a vertical conduit, allowing the read data to traverse down into the middle semiconductor device 1116, where it is received by the read data register 1154. The read data register 1154 holds the read data momentarily before it is directed to the multiplexer 1114 through interconnect 1156.

[00285] The multiplexer 1114 in the middle semiconductor device 1116, coordinated by the phase counter 1110, selects the appropriate data for output. The read data is then channeled to the read data register 1112, where it is briefly stored. Following this, the read data travels via interconnect 1120 to bumpless bonds 1159.

[00286] The bumpless bonds 1159 form an interface with bumpless bonds 1158 on the bottom semiconductor device 1122. The read data signal is then carried through interconnect 1176, culminating its journey at the read data register 1124 on the bottom device.

[00287] Datapath 4 within the assembly 1100 is a path that establishes the flow of read addresses from the read address register 1126 on the bottom semiconductor device 1122 to the microvault 1164 located on the top semiconductor device 1106, and subsequently facilitates the movement of read data back down to the read data register 1124 on the bottom device.

[00288] This datapath begins at the read address register 1126, which is responsible for holding and issuing the read addresses necessary for data retrieval from the microvaults. The read address is sent from the register 1126 through interconnect 1174, a pathway that maintains signal integrity and facilitates electrical communication.

[00289] From interconnect 1174, the read address is directed to bumpless bonds 1128. These bonding areas create an interconnection between the bottom semiconductor device 1122 and the middle semiconductor device 1116 through bumpless bonds 1129. The design of these bumpless bonds facilitates data transmission.

[00290] Once the read address reaches the middle semiconductor device 1116, it is carried forward by interconnect 1130 to the read address register 1146. This register acts as an intermediary, preparing the address for its vertical ascent through the device stack. The address is then transmitted via the Through-Silicon Via (TSV) 1136 that facilitates vertical integration by providing a direct electrical link through the semiconductor substrate.

[00291] After ascending through TSV 1136, the read address emerges onto bumpless bonds 1160, which are aligned to connect with bumpless bonds 1161 on the top semiconductor device 1106. The read address then proceeds along interconnect 1140 to the read address register 1142 located on the top device.

[00292] The read address register 1142 serves to forward the read address to its final destination, the microvault 1164, through interconnect 1148, labeled as path "4". Microvault 1164, upon receiving the read address, retrieves the requested data, which is then outputted through interconnect 1168. This data is directed to the multiplexer 1150, which is under the control of phase counter 1102.

[00293] The phase counter 1102 determines which data stream is selected by the multiplexer 1150, which then sends the read data to the read data register 1104. The read data register 1104 acts as a temporary repository, holding the data until it can be sent downwards through the device stack.

[00294] The data leaves the read data register 1104 and travels via interconnect 1108 to bumpless bonds 1162. These bonds maintain a connection to bumpless bonds 1163 on the middle semiconductor device 1116. The read data is then transferred to TSV 1152, which carries the data vertically down to the read data register 1154 on the middle semiconductor device 1116.

[00295] The read data register 1154 temporarily holds the read data before it is fed into the multiplexer 1114 via interconnect 1156. The multiplexer 1114, coordinated by the phase counter 1110, channels the appropriate data stream to the read data register 1112. This register serves as a staging area for the read data, which is then sent through interconnect 1120 to bumpless bonds 1159.

[00296] Bumpless bonds 1159 interface with bumpless bonds 1158 on the bottom semiconductor device 1122 to provide a downward transmission of the read data. Finally, the signal is routed through interconnect 1176 and arrives at read data register 1124, where the data is made available for subsequent processing or external communication.

[00297] Fig. 12 illustrates a three-dimensional (3D) memory column 1200 configured as a 3D-NOR or 3D-AND structure, featuring a series of ferroelectric field-effect transistors

(FeFETs) 1202 with interconnected drain terminals 1204 linked to a common select line 1212 and individual gate terminals 1206 connected to respective read/write enable lines 1214 (e.g., 1214a for FeFET 1202a, all coupled to a common bit line, in accordance with an embodiment of the present disclosure.

[00298] Thus, Fig. 12 depicts a three-dimensional (3D) memory column, designated as element 1200, which can be configured in various embodiments as either a 3D-NOR or 3D-AND structure, providing flexibility in the application and use of the integrated circuit. This memory column is an assembly of multiple ferroelectric field-effect transistors (FeFETs), collectively referred to as fefets 1202, where each FeFET is indicated by elements such as 1202a, 1202b, 1202c, and 1202d, among others potentially present in the array.

[00299] Within each FeFET, such as 1202a, there is a drain terminal 1204a. This drain terminal is part of the memory cell's output path and is connected to a common select line 1212. In some embodiments, the common select line 1212 serves as a control mechanism that enables the selection of a particular FeFET for data read or write operations.

[00300] The gate terminal of each FeFET, exemplified by 1206a for FeFET 1202a, is individually connected to a respective read/write enable line, such as 1214a. This enables control of the FeFET's state, allowing it to be in a conductive (on) state for reading or writing data or in a non-conductive (off) state to prevent data flow. The presence of individual read/write lines for each FeFET may allow for precise control and operation of each memory cell.

[00301] Moreover, each FeFET, such as 1202a, comprises a source terminal, such as 1208a, which is coupled to a common bit line 1210. The bit line 1210 provides a conduit for data being written to or read from the FeFETs. In some embodiments, this bit line can be shared across multiple memory columns, which can facilitate parallel processing and increased data throughput.

[00302] In accordance with various embodiments of the present disclosure, the 3D memory column 1200 may incorporate additional elements and configurations to enhance performance and functionality. For instance, the 3D memory column 1200 may include insulating materials, conductive pathways, and other structural components not explicitly shown in Fig. 12 but which are inherent to the implementation of such 3D memory structures. The FeFETs 1202 may also exhibit variations in terms of material composition, structural dimensions, and electrical properties, contributing to a range of performance characteristics suitable for different applications.

[00303] Furthermore, the memory column 1200 may be incorporated into larger memory arrays, forming part of a memory module or system. These arrays can be arranged in various configurations, such as rows and columns, to create a matrix that efficiently addresses the demands of high-density data storage. The memory column 1200 can also be interfaced with other circuit elements and control logic, which may govern the operation of the memory array, including data management protocols, error correction algorithms, and power optimization strategies.

[00304] In some embodiments, the memory column 1200 may be fabricated using advanced semiconductor manufacturing techniques, such as photolithography, etching, deposition, and planarization processes. The choice of materials for the FeFETs, including the ferroelectric material, the semiconductor channel, and the conductive elements, can be selected based on desired electrical characteristics, such as charge retention, switching speed, and energy efficiency.

[00305] The 3D memory column 1200, as illustrated in Fig. 12, may include FeFETs, such as 1202, fabricated from a variety of materials that provide the necessary electrical and physical properties to achieve the desired functionality. For instance, in some embodiments, the channel layer of each FeFET in the FeFETs 1202 could be constructed from materials such as Indium Gallium Zinc Oxide (IGZO) or other Amorphous Oxide Semiconductors (AOS) like Zinc Tin Oxide or Indium Tungsten Oxide (IWO). These materials are selected for their electronic properties, such as carrier mobility and stability.

[00306] The ferroelectric material rigidly coupled to the channel layer in each FeFET may comprise hafnium zirconium oxide (HfZrO₂) or other transition metal oxides, perovskites, etc. These ferroelectric materials are chosen for their ability to maintain a polarization state when an electric field is applied, which is used for the non-volatile memory characteristics of the FeFETs. The thickness, crystalline structure, and stoichiometry of the ferroelectric layer can be controlled to achieve the desired coercive voltage, remanent polarization, and other electrical parameters for reliable data storage and retrieval.

[00307] The drain 1204 and source 1208 terminals of the FeFETs 1202 are connected to the common select line 1212 and common bit line 1210, respectively. These common lines may be formed from conductive materials such as tungsten, titanium nitride, or other metals and metal alloys that provide low-resistance pathways for electrical signals. The configuration of these terminals and their respective common lines ensures that the FeFETs can be accessed and controlled effectively during operation.

[00308] Each gate terminal, such as the gate 1206 of the FeFET 1202a, is connected to its respective read/write enable line, such as 1214a. The gate terminals help control the state of the FeFET, and the materials chosen for these terminals may include various conductive materials that can provide a reliable electrical interface with the ferroelectric material. The read/write enable 1214 lines are designed to deliver suitable voltage levels to the gates 1206 of the FeFETs 1202 for switching between states.

[00309] The memory column 1200 as a whole is designed to support a range of operating parameters. In some embodiments, these parameters may include, but are not limited to, an off-state current of less than 10^{-8} amps per centimeter cubed, an on-state current greater than 10^{-7} amps per centimeter cubed, and a channel mobility that is maintained despite the presence of the ferroelectric layer. The channel layer's thickness can be less than 30 nm to ensure high device density, while the ferroelectric layer's characteristics, such as coercive voltage and remanent polarization, are optimized to provide the necessary memory functionality.

[00310] In some embodiments, the FeFETs 1202 may include additional materials or dopants to enhance their electrical properties. For instance, dopants such as gallium (Ga), indium (In), or zinc (Zn) may be introduced into the channel layer to modulate the carrier concentration or to adjust the threshold voltage of the FeFETs. Similarly, the ferroelectric layer may include dopants like lanthanum (La) or niobium (Nb) to adjust its ferroelectric properties.

[00311] In other embodiments, the 3D memory column 1200 may be integrated with additional semiconductor devices and structures to form complex memory systems. These systems can provide storage capabilities and support various memory architectures.

[00312] Fig. 13 depicts a three-dimensional (3D) memory column 1300 configured as a 3D-NAND structure, consisting of a vertical stack of ferroelectric field-effect transistors (FeFETs) 1302, each with source 1304 and drain 1308 terminals. The source 1304 of each FeFET, such as 1304a for FeFET 1302a, is coupled to the start of a bit line 1310 or connected to the drain of the preceding FeFET, exemplified by source 1304b of FeFET 1302b coupled to drain 1308a of FeFET 1302a. Each FeFET includes a gate 1306, such as 1306a for FeFET 1302a, connected to a respective read/write enable line, illustrated by 1314a for FeFET 1302a, in accordance with an embodiment of the present disclosure.

[00313] The 3D memory column 1300 is composed of a series of vertically stacked Field-Effect Transistors (FeFETs), identified collectively as FeFETs 1302. These transistors, which include FeFETs 1302a, 1302b, 1302c, 1302d, and so on, are characterized by their incorporation of ferroelectric materials within their gate structure. Each FeFET in the series is of the memory column contributes to the memory storage capabilities of the device.

[00314] In the depicted embodiment, each FeFET, such as FeFET 1302a, includes a source terminal 1304, for instance, source 1304a, which is coupled to a bit line 1310. The bit line 1310 serves as a conduit for electrical signals that are used to read from and write to the memory cell associated with FeFET 1302a. In scenarios where FeFET 1302a is not the bottom-most transistor in the column, its source 1304b may be connected to the drain 1308a of the immediately preceding FeFET, such as FeFET 1302a, facilitating a serial connection that defines the vertical NAND architecture.

[00315] Each FeFET within the FeFETs 1302 is further equipped with a gate terminal 1306, exemplified by gate 1306a for FeFET 1302a. This gate terminal 1306 is coupled to a respective read/write enable line, exemplified by 1314a for FeFET 1302a. The read/write enable line 1314a is responsible for controlling the state of the FeFET, allowing it to either conduct or prevent the flow of current through the device, thereby enabling the writing or reading of data.

[00316] Moreover, each FeFET of the FeFETs 1302 also includes a drain terminal 1308, such as drain 1308a for FeFET 1302a, which is typically connected to the source of the subsequent FeFET in the vertical stack. This arrangement ensures that the charge stored in the ferroelectric material of the gate can modulate the current flowing from the source to the drain, allowing for the storage and retrieval of data.

[00317] The memory column 1300 within Fig. 13 is indicative of a memory architecture that can be utilized in various applications, from portable electronics to enterprise-level data storage systems. In some embodiments, the ferroelectric material used in the FeFETs may include various compositions, such as hafnium oxide, zirconium oxide, or any combination thereof, which can be doped with elements such as lanthanum or yttrium to adjust the ferroelectric properties as required.

[00318] In some variations, the 3D memory column 1300 can incorporate additional features that enhance performance, reliability, or manufacturability. For instance, the FeFETs 1302 may include protective layers to shield the ferroelectric material from environmental factors or process-induced damage. The column 1300 may also be integrated with other circuit elements, such as capacitors or diodes, to facilitate operations like charge pumping or to provide additional functionality within the memory array.

[00319] The 3D memory column 1300 may be fabricated from a variety of materials that confer specific electrical properties to enhance device performance. In some embodiments, the channel layer of each FeFET may be formed from materials such as Indium Gallium Zinc

Oxide (IGZO). Other materials for the channel layer could include Amorphous Oxide Semiconductors (AOS) like Zinc Tin Oxide or Aluminum Zinc Oxide.

[00320] The ferroelectric layer within the FeFETs 1302 may comprise materials such as Hafnium Zirconium Oxide (HfZrO₂). The ferroelectric layer's thickness and material composition can be controlled through methods like Atomic Layer Deposition (ALD) to achieve the desired coercive voltages, remanent polarizations, and endurance characteristics. In some implementations, the coercive voltage of the ferroelectric layer may be tuned to be between -3 Volts to +3 Volts, facilitating low-voltage operation of the memory devices.

[00321] The source and drain terminals of the FeFETs 1302 may be composed of conductive materials such as Tungsten or Titanium Nitride. These materials may also be selected to optimize the contact resistance with the channel layer, reducing overall power consumption and improving the Ion/Ioff ratio of the device.

[00322] Additionally, the FeFETs 1302 may be engineered to exhibit specific electrical parameters. For instance, the channel layer's thickness may be less than 30 nm in some embodiments. In some embodiments, the channel layer may demonstrate a carrier concentration of 10^{17} to 10^{20} per centimeter-cubed, which can be adjusted through doping with elements such as Gallium, Indium, or Zinc to modulate the electrical properties.

[00323] The memory cells formed by the FeFETs 1302 within the 3D memory column 1300 may also target operational parameters such as read and write latencies, endurance, and energy consumption. For example, read and write operations may be executed with energies less than 10 picojoules and within timeframes less than 20 nanoseconds, contributing to the low power and high-speed attributes of the memory column.

[00324] Furthermore, the 3D-NAND configuration of the memory column 1300 may be designed to achieve a high off-state resistance to on-state resistance ratio (R_{off}/R_{on}), which is critical for distinguishing between different data states and ensuring reliable data retention. This ratio may be about 10^3 or greater, which helps to maintain a high signal-to-noise ratio during memory operations.

[00325] The FeFETs 1302 in the memory column 1300 may also be designed to sustain a high degree of reliability, with endurance ratings greater than or equal to 10^{11} cycles, ensuring the longevity and durability of the memory device. This endurance is complemented by the ferroelectric layer's ability to maintain data retention for at least 1 minute at room temperature, which is 25°C.

[00326] Fig. 14 depicts a three-dimensional (3D) memory column configured as a 3D-NAND with an integrated pass gate, in accordance with an embodiment of the present

disclosure. This figure illustrates a series of ferroelectric field-effect transistors (FeFETs) 1402, each including source 1404 and drain 1408 terminals, gated by respective gate terminals 1406 and coupled to read/write enable lines 1414. The FeFETs are interconnected, forming a vertical memory structure with pass gates 1418 linked to a pass gate line 1416.

[00327] Thus, Fig. 14 illustrates a three-dimensional (3D) memory column, designated as element 1400, which can be configured as a 3D-NAND structure with an integrated pass gate. This configuration enables enhanced control over individual memory cells within the 3D structure, potentially improving read/write operations and facilitating efficient memory management.

[00328] In detail, the 3D memory column 1400 comprises multiple ferroelectric field-effect transistors (FeFETs), collectively referred to as FeFETs 1402. Each FeFET within the 1402 series, such as 1402a, 1402b, 1402c, 1402d, etc., is a constituent memory cell of the 3D memory column 1400. These FeFETs are utilized for their ability to retain data in a non-volatile manner due to the ferroelectric properties of their gate material, which allows for data retention without continuous power supply for a time.

[00329] Each FeFET in the series 1402 includes a source, exemplified by source 1404a for FeFET 1402a. The source 1404 for each FeFET is either coupled to a bit line, illustrated as bit line 1410 for FeFET 1402a, or is connected to the drain of the preceding FeFET in the series. For instance, source 1404b of FeFET 1402b is electrically coupled to drain 1408a of FeFET 1402a. This serial connection forms the basis for the daisy-chain configuration, which is used in NAND architectures, allowing for sequential access to the array of FeFETs.

[00330] Furthermore, each FeFET within the FeFETs 1402 is equipped with a gate terminal, such as gate 1406a for FeFET 1402a. The gates of the FeFETs are connected to their respective read/write enable lines, which are depicted as element 1414 in the figure. For example, gate 1406a of FeFET 1402a is influenced by read/write enable line 1414a. These enable lines control the application of appropriate voltages for the reading and writing of data.

[00331] Additionally, each FeFET in the FeFETs 1402 series includes a drain, such as drain 1408a for FeFET 1402a. This drain is connected to the source of the subsequent FeFET in the series, thus establishing the continuity for the columnar structure of the 3D memory stack.

[00332] In some embodiments, each FeFET of the FeFETs 1402 incorporates a pass gate, for example, pass gate 1418a, which is connected to a pass gate line, represented by 1416 in the figure. The pass gate line 1416 is a conductive pathway that provides electrical signals to control the pass gates 1418 of the FeFETs. The inclusion of pass gates in the FeFETs may

allow for improved isolation between memory cells during operation, thereby reducing interference and potentially enhancing the reliability of data storage and retrieval.

[00333] The 3D memory column 1400, as depicted in Fig. 14 also encompasses a diverse range of materials and parameters that could be utilized to optimize its performance in various embodiments. Each FeFET 1402 within the column could be fabricated using a variety of semiconductor materials. For instance, the channel layer of the FeFETs could be formed from materials such as Indium Gallium Zinc Oxide (IGZO).

[00334] The ferroelectric layer, which is a defining characteristic of the FeFETs, may be composed of materials like Hafnium Zirconium Oxide (HfZrO₂) or other perovskite materials, which are known for their remanent polarization. This property determines the data retention capabilities of the FeFETs. The coercive voltage of this layer, which affects the energy required to switch the polarization state, is another critical parameter that can be adjusted according to the requirements of the specific application, with a range in one embodiment being between -3 Volts to 3 Volts.

[00335] The source and drain terminals of the FeFETs, which include elements 1404 and 1408, respectively, could be composed of conductive materials such as Tungsten or Titanium Nitride. These materials provide pathways for current flow, which is for switching. The read/write enable lines 1414, which control the gates 1406 of the FeFETs, could also be fabricated from similar materials, ensuring consistent electrical characteristics throughout the device.

[00336] In terms of the physical parameters, the channel layer thickness may be less than 30 nm in thickness. The electron mobility within these channel layers may be maintained at a predetermined level even when the layer is less than 30nm in thickness

[00337] The pass gates 1418 may be manufactured using low-resistance materials to enable quick switching times, which is beneficial when the memory column is accessed frequently during operation.

[00338] Fig. 15 illustrates a three-dimensional (3D) memory column 1500, which may be configured as either a 3D-NOR or a 3D-AND structure with independent Read/Write enable capabilities, in accordance with an embodiment of the present disclosure. This memory column encompasses a series of vertically aligned FeFETs 1502, such as FeFETs 1502a, 1502b, 1502c, 1502d, and so forth, each integrated with a source 1504 (e.g., source 1504a for FeFET 1502a) linked to a respective read enable line 1520 (e.g., read enable line 1520a for FeFET 1502a), and a gate 1506 (e.g., gate 1506a for FeFET 1502a) connected to a corresponding write enable line 1522 (e.g., write enable line 1522a for FeFET 1502a). All FeFETs within the column share

a common bit line 1510 connected to their drains 1508, enabling the column to perform coordinated memory operations.

[00339] Fig. 15 presents a detailed depiction of a three-dimensional (3D) memory column 1500, which can be configured as a 3D-NOR or 3D-AND structure with independent Read/Write enable functionalities. This memory column is an assembly of Field-Effect Transistors with ferroelectric gate layers, commonly referred to as FeFETs 1502, which are individually identified, for example, as 1502a, 1502b, 1502c, 1502d, etc., each representing a memory cell within the column.

[00340] In the illustrated embodiment, each FeFET 1502 includes a source 1504, such as source 1504a corresponding to FeFET 1502a. The source 1504 is designed to be electrically coupled to a respective read enable line 1520, such as read enable line 1520a which is dedicated to FeFET 1502a. The read enable line 1520 functions to selectively activate the FeFET 1502 for reading operations, allowing the readout of stored data from the memory cell.

[00341] Additionally, each FeFET 1502 is equipped with a gate 1506, exemplified by gate 1506a for FeFET 1502a. The gate 1506 is connected to a respective write enable line 1522, such as write enable line 1522a, which is specific to FeFET 1502a. The write enable line 1522 serves to selectively activate the FeFET 1502 for writing operations, enabling the storage of data within the memory cell.

[00342] Furthermore, each FeFET 1502 includes a drain 1508, for instance, drain 1508a affiliated with FeFET 1502a. The drain 1508 is connected to a common bit line 1510. The bit line 1510 acts as a conduit for transferring data to and from the memory cells during read and write operations. The commonality of the bit line 1510 across multiple FeFETs 1502 signifies that data from any activated memory cell can be routed through this shared path.

[00343] In some embodiments of the memory column 1500, the configuration of the FeFETs 1502 allows for a high density of memory cells vertically stacked within a compact footprint.

[00344] The ferroelectric material utilized in the gate 1506 of the FeFETs 1502 may comprise various compositions, such as hafnium oxide-based materials, which can be deposited using atomic layer deposition techniques. The ferroelectric property of the material allows for data retention, enabling the memory cells to maintain stored information even when power is not supplied.

[00345] The source 1504, gate 1506, and drain 1508 of each FeFET 1502 can be fabricated from materials that provide predetermined electrical performance. These materials may include metals such as tungsten or copper, or metal nitrides such as titanium nitride.

[00346] The read enable lines 1520 and write enable lines 1522 can be designed to minimize crosstalk and interference between adjacent lines, in some embodiments. In some specific embodiments, shielding layers or insulating materials may be included to further isolate the signal paths.

[00347] Furthermore, the described memory column 1500 may be integrated within a larger semiconductor device, such as a processor or a storage module. It may form part of a system-on-chip (SoC) or be included in a multi-chip module (MCM), contributing to a data storage and retrieval system.

[00348] The materials that constitute the FeFETs 1502 within the memory column 1500 are selected to provide specific electrical and physical properties to optimize the performance of the integrated circuit. For instance, the channel layer in each FeFET may be formed from advanced semiconductor materials, such as Indium Gallium Zinc Oxide (IGZO) or other Amorphous Oxide Semiconductors (AOS) like Zinc Tin Oxide or Cadmium Oxide. These materials are chosen for their excellent electron mobility characteristics and stability.

[00349] The ferroelectric layer, integral to the FeFETs 1502, may be fabricated from various ferroelectric materials that exhibit suitable polarization properties. Materials such as Hafnium Zirconium Oxide (HfZrO₂) or Lead Zirconate Titanate (PZT) could be utilized. These materials can be doped with elements such as Lanthanum, Yttrium, or other suitable dopants to modify their ferroelectric properties, such as coercive voltage, remanent polarization, and crystallization temperature. The ferroelectric layer's thickness and material composition may be adjusted to achieve desired memory characteristics, such as write endurance and retention time, while ensuring the layer remains compatible with the overall semiconductor manufacturing process, or other considerations, etc.

[00350] The source 1504, gate 1506, and drain 1508 terminals of the FeFETs 1502 may be composed of conductive materials like Tungsten, Titanium Nitride, Nickel, or Molybdenum. Connections to the read enable lines 1520 and the write enable lines 1522 may be facilitated through conductive vias or contacts.

[00351] The read enable lines 1520 and write enable lines 1522, along with the common bit line 1510, may be patterned using lithographic techniques to achieve the predetermined precision and alignment for proper functionality. These lines may be insulated from one another using dielectric materials like Silicon Dioxide (SiO₂), Silicon Nitride (Si₃N₄), or low-k dielectrics to reduce parasitic capacitance and crosstalk.

[00352] Each element within the memory column 1500 may consider factors such as line width, spacing, and aspect ratio to ensure manufacturability, functionality, and/or other

goals or characteristics. The materials and processes used in the construction of the memory column 1500 are chosen to ensure compatibility with standard semiconductor fabrication techniques, such as photolithography, etching, deposition, and annealing, while also enabling the integration of materials and structures.

[00353] The fabrication of the FeFETs 1502 within the memory column 1500 may involve deposition techniques such as atomic layer deposition (ALD), chemical vapor deposition (CVD), or physical vapor deposition (PVD) to create uniform and/or non-uniform layers.

[00354] Fig. 16 presents a cross-sectional view of a 3D memory structure, designated as 1600, configured as a single-port 3D NAND, in accordance with an embodiment of the present disclosure. The structure includes a first vertical structure 1608a and a second identical vertical structure 1608b, each comprising a dielectric column 1610a, 1610b, a channel column 1612a, 1612b disposed around the dielectric column, and a ferroelectric column 1614a, 1614b disposed around the channel column. A series of horizontal gate-electrode layers 1606a-c are disposed at predetermined distances from each other, adjacent to the ferroelectric column along the length of the vertical structures. The assembly further includes a drain select layer 1602 and a source select layer 1604, with respective end dielectric columns 1618a, 1618b, and 1616a, 1616b positioned at the interfaces with the vertical structures, illustrating a detailed and intricate design for high-density data storage.

[00355] Thus, fig. 16 provides a cross-sectional view of a three-dimensional (3D) memory structure, designated as 1600, which is configured as a single-port 3D NAND architecture. This structure incorporates a pair of vertical structures, 1608a and 1608b, which may be fabricated to be substantially identical, as indicated by their respective subscripts a and b, suggesting the potential for a modular and scalable memory array design.

[00356] Each vertical structure, exemplified by the first vertical structure 1608a, includes a dielectric column 1610a. The dielectric column may adopt various geometric forms—it can be cylindrical, substantially cylindrical, or feature curves. Additionally, it may present a tapered form, having different diameters at each end, implying a design that narrows towards the top. Both solid and hollow configurations of the dielectric column are contemplated within the scope of the disclosure, offering design flexibility for different electrical and structural requirements.

[00357] Surrounding the dielectric column 1610a is a channel column 1612a, which is the locus for charge carriers during device operation. The channel column is also described as

potentially cylindrical, substantially cylindrical, or feature curves and/or and may exhibit similar variations in diameter along its length as the dielectric column.

[00358] Enveloping the channel column 1612a is a ferroelectric column 1614a, which extends along the length of the channel column but may recede at the ends, thereby meaning the ferroelectric column 1614 does not extend the entire length of the channel column 1612a.

[00359] Intersecting with the vertical structures is a series of horizontal gate-electrode layers, 1606a-c, which are positioned at predetermined distances from one another. These layers play a role in controlling the operational states of the device by influencing the electric field within the ferroelectric column.

[00360] Atop the 3D memory structure 1600 sits a drain select layer 1602, parallel to the horizontal gate-electrode layers 1606. Where the drain select layer 1602 meets the vertical structures 1608a, 1608b, end dielectric columns, 1618a and 1618b, are discernible. These end dielectric columns 1618 interface with the channel column 1612 and the drain select layer 1602, contributing to the isolation and control of the charge carriers within the channel column. They may contact the ferroelectric layer 1614, as they envelop the channel column 1612 at different positions along its length.

[00361] Similarly, a source select layer 1604 is situated at the bottom of the structure 1600, again parallel to the horizontal gate-electrode layers 1606. Corresponding end dielectric columns, 1616a and 1616b, are present where the source select layer 1604 interfaces with the vertical structures, serving analogous functions to the end dielectric columns 1618 near the drain select layer 1602.

[00362] The horizontal gate-electrode 1606 layers could be constructed from a range of conductive materials, including metals and metal compounds, which may offer different work functions, conductivity, and compatibility with other materials in the structure. Similarly, the ferroelectric column 1614 might incorporate a variety of ferroelectric materials each with its unique polarization characteristics, coercive fields, and dielectric constants, affecting the device's memory retention and switching behaviors.

[00363] The channel column 1612 materials can be chosen based on their electronic properties, such as carrier mobility and bandgap, to achieve the desired levels of on-state and off-state current. The dielectric column 1610 provides the electrical insulation necessary to prevent leakage currents and ensure the proper functioning of the device.

[00364] The dielectric column, such as 1610a for the first vertical structure, may be constructed from materials that offer insulating properties to mitigate any potential leakage

currents. Choices for the dielectric material may be Hafnium Oxide (HfO₂) or Silicon Dioxide (SiO₂).

[00365] Surrounding the dielectric column, the channel column (1612a) has channel material can be selected from a wide range of semiconducting materials that offer predetermined carrier mobility. For example, Indium Gallium Zinc Oxide (IGZO) can be used for its electrical properties. The channel layer's thickness may vary, with some embodiments considering a thickness less than 30 nm. This thickness is chosen to achieve a predetermined electrical performance. The ferroelectric column, like 1614a, may include perovskite structures, such as Lead Zirconate Titanate (PZT).

[00366] The horizontal gate-electrode layers, represented by 1606a-c, are composed of conductive materials that facilitate the application of an electric field to the ferroelectric column, such as Tungsten or Titanium Nitride, which may be chosen for their electrical behavior. The selection of gate-electrode materials also takes into consideration factors such as work function, thermal stability, and ease of integration with the existing semiconductor manufacturing processes.

[00367] The drain and source select layers, 1602 and 1604 respectively, are incorporated to enable the addressing of individual memory cells within the array. The materials used for these layers are chosen for their conductive properties and compatibility with the channel and ferroelectric materials. The design of these layers may also incorporate considerations for reducing parasitic capacitance and ensuring swift data access.

[00368] The end dielectric columns, like 1618a and 1616a, provide electrical insulation at the ends of the channel column, where the ferroelectric material does not extend.

[00369] The disclosed embodiments within the 3D memory structure 1600 outline an assembly capable of providing data storage. The design allows for variations in structural dimensions, such as the diameter of the cylindrical columns, which can be uniform or tapered. Additionally, the option for solid or hollow configurations may be used.

[00370] Fig. 17 shows a 3D memory structure that is a dual-port 3D NAND arrangement in accordance with an embodiment of the present disclosure. The three-dimensional (3D) memory structure illustrated in Fig. 17, referred to as 3D memory structure 1700, exemplifies a dual-port 3D NAND arrangement to provide a memory functionality. This structure is characterized by two primary vertical formations, designated as the first vertical structure 1708a and the second vertical structure 1708b, which may be identical or near identical, as evidenced by the designating subscripts 'a' and 'b'.

[00371] The first vertical structure 1708a includes a hollow or solid, tapered pass-gate electrode column 1718a that is substantially cylindrical in shape. The pass-gate electrode column 1718a may be made of titanium nitride and may have a larger diameter at the bottom end compared to the top end.

[00372] Surrounding the pass-gate electrode column 1718a is a dielectric column 1710a that may be made of hafnium oxide. The dielectric column 1710a is also substantially cylindrical with a slightly tapered shape, having a marginally larger diameter at the top. The dielectric column 1710a provides electrical isolation between the pass-gate electrode and subsequent layers.

[00373] Disposed around the dielectric column 1710a is a cylindrical channel column 1712a that may be made of IGZO semiconductor material. The channel column 1712a features curves along its length and has a uniform diameter throughout. The thickness of the channel column may be less than 30 nm.

[00374] Enclosing the channel column 1712a is a PZT ferroelectric column 1714a that covers most of the length of the channel column 1712a but recedes at the ends, leaving a portion of the channel column 1712a uncovered. The ferroelectric column 1714a is substantially cylindrical and contains lead, zirconium and titanium as key elemental constituents.

[00375] The vertical structures 1708a and 1708b traverse through several horizontal gate-electrode layers 1706a, 1706b and 1706c that may be made of tungsten, which are positioned at fixed intervals to form an interconnected grid layout. These layers influence the electric field within the ferroelectric column 1714a during memory operations.

[00376] At the top of the memory structure 1700, a drain select layer 1702 (e.g., titanium nitride) runs parallel to the horizontal gate-electrode layers 1706. Where the drain select layer 1702 intersects the vertical structures 1708a and 1708b, end dielectric columns 1718a and 1718b are visible. These end columns (e.g., made of HfO₂) touch the ferroelectric column 1714a on one end and surround the uncovered portion of channel column 1712a, providing insulation.

[00377] Similarly at the bottom of the structure 1700, a source select layer 1704 (e.g., made of tungsten), also parallel to the electrode layers 1706, interfaces with the vertical structures. End dielectric columns 1716a and 1716b can be observed at these intersection points, enclosing the open ends of the channel columns 1712a and 1712b.

[00378] Within the hollow region of the pass-gate electrode columns 1718a and 1718b at the ends, a thin dielectric horizontal layers 1720a and 1720b may be placed near the bottom terminals (e.g., HfO₂). These layers seal off the bottom open ends of the vertical hollow voids.

[00379] Fig. 18 illustrates a 3D memory structure 1800 that can be configured as a 3D NOR Vertical Transistor memory array. The 3D memory structure 1800 comprises a first vertical structure 1808a and an identical (or substantially identical) second vertical structure 1808b arranged adjacent to one another.

[00380] The first vertical structure 1808a includes a vertical plug column 1802a that provides an electrical connection to the lower portions of the 3D memory structure. The vertical plug column 1802a may have a uniform diameter along its entire length or may have a larger diameter on its lower end than on its upper end. The plug column 1802a can be fabricated as a solid column or as a hollow column in various embodiments.

[00381] Disposed adjacent to the vertical plug column 1802a is a source electrode column 1804a and a drain electrode column 1816a. The source electrode column 1804a and drain electrode column 1816a provide electrical connections to the source and drain nodes of the vertical transistors formed along the vertical structure 1808a. The source electrode column 1804a and drain electrode column 1816a may be comprised of various conducting materials including, but not limited to, tungsten, titanium nitride, tantalum nitride, nickel, molybdenum, platinum, palladium, cobalt, gold, aluminum, copper, hafnium, hafnium nitride, iridium, iridium oxide, ruthenium, ruthenium oxide, silicides, graphene, carbon nanotubes, doped polysilicon, indium tin oxide, silver, aluminum-doped zinc oxide, gallium, gallium arsenide, indium gallium zinc oxide, metal alloys such as AlCu and TiW, and conducting polymers.

[00382] Surrounding the vertical plug column 1802a, source electrode column 1804a, and drain electrode column 1816a is a channel column 1812a that provides the semiconductor channel region for the vertical transistors along the first vertical structure 1808a. The channel column 1812a may be formed from materials including, but not limited to, indium gallium zinc oxide (IGZO), indium zinc oxide (IZO), zinc tin oxide (ZTO), aluminum zinc oxide (AZO), indium tungsten oxide (IWO), gallium zinc oxide (GZO), hafnium indium oxide (HIO), cadmium oxide (CdO), polysilicon, polygermanium, cadmium selenide (CdSe), copper indium gallium selenide (CIGS), crystalline silicon, crystalline germanium, gallium arsenide (GaAs), indium phosphide (InP), indium antimonide (InSb), silicon carbide (SiC), gallium nitride (GaN), zinc oxide (ZnO), pentacene, P3HT, polythiophene, PPV, graphene, carbon nanotubes (CNTs), methylammonium lead halides, cesium lead halides, lead sulfide (PbS), lead selenide (PbSe), cadmium selenide (CdSe), indium arsenide (InAs), and other semiconducting materials.

[00383] Surrounding the channel column 1812a is a ferroelectric column 1814a that provides the gate dielectric for the vertical transistors along the first vertical structure 1808a.

The ferroelectric column 1814a may be comprised of ferroelectric materials including, but not limited to, perovskite oxides, lead zirconate titanate (PZT), barium titanate (BaTiO_3), strontium titanate (SrTiO_3), bismuth ferrite (BiFeO_3), potassium niobate (KNbO_3), lithium niobate (LiNbO_3), lithium tantalate (LiTaO_3), sodium bismuth titanate ($\text{Na}_{0.5}\text{Bi}_{0.5}\text{TiO}_3$), bismuth titanate ($\text{Bi}_4\text{Ti}_3\text{O}_{12}$), bismuth zinc niobate ($\text{Bi}(\text{Zn}_{1/2}\text{Ti}_{1/2})\text{O}_3$), bismuth lanthanum titanate (BiLaTiO_3), bismuth nickel titanate (BiNiTiO_3), PMN-PT, PLZT, neodymium-doped bismuth titanate ($\text{Bi}_{4-x}\text{Nd}_x\text{Ti}_3\text{O}_{12}$), hafnium-based oxides like hafnium oxide (HfO_2) and doped hafnium oxide, tungsten bronze structure materials, barium strontium niobate (BSN), lead barium niobate (PBN), potassium tantalate niobate (KTN), bismuth titanate ($\text{Bi}_4\text{Ti}_3\text{O}_{12}$), strontium bismuth tantalate (SBT), calcium bismuth niobate (CBN), organic ferroelectrics like PVDF, TrFE and P(VDF-TrFE) copolymers, aurivillius phase oxides, rare earth manganites like YMnO_3 , lanthanum-modified PLZT, nickel manganese oxide (NiMnO_3), relaxor ferroelectrics like PMN, PST and PIN, multiferroic materials like TbMnO_3 , EuTiO_3 , SbSI, GeTe, SnTe, PZT thin films, SBT thin films, HfO_2 -based thin films, layered superlattices, and $\text{PbTiO}_3/\text{SrTiO}_3$.

[00384] The 3D memory structure 1800 further comprises multiple horizontal gate electrode layers 1806 including layers 1806a, 1806b, 1806c etc. The horizontal gate electrode layers 1806 are disposed at regular intervals along the vertical structures 1808 and provide the gate electrodes for the vertical transistors. The gate electrode layers 1806 may be formed from materials such as tungsten, titanium nitride, tantalum nitride, nickel, molybdenum, platinum, palladium, cobalt, gold, aluminum, copper, hafnium, hafnium nitride, iridium, iridium oxide, ruthenium, ruthenium oxide, silicides, graphene, carbon nanotubes, doped polysilicon, indium tin oxide, silver, aluminum-doped zinc oxide, gallium, gallium arsenide, indium gallium zinc oxide, metal alloys such as AlCu and TiW, and conducting polymers.

[00385] Each of the horizontal gate electrode layers 1806 may be surrounded by an oxide/nitride/oxide (ONO) stack 1810, such as 1810a surrounding gate electrode layer 1806a, to provide insulation between the gate electrodes.

[00386] The second vertical structure 1808b in the 3D memory structure 1800 is identically configured as the first vertical structure 1808a. The two vertical structures 1808a and 1808b are arranged horizontally adjacent to each other with a spacing that allows integration of the gate electrode layers 1806 and ONO stacks 1810. Together, the first and second vertical structures 1808a, 1808b along with the horizontal gate electrode layers 1806 can be configured as a 3D NOR memory architecture.

[00387] Fig. 19 illustrates an embodiment of a planar FeFET 1900. The FeFET 1900 comprises a substrate 1910 upon which various layers and components are formed. The substrate 1910 may be comprised of silicon or other suitable semiconductor materials. Disposed on top of the substrate 1910 is a layer of TiN 1912. The TiN layer 1912 may act as an electrode and can be deposited by sputtering or other suitable deposition techniques.

[00388] On top of the substrate 1910 and TiN layer 1912, a layer of HZO 1908 is disposed. HZO 1908 comprises hafnium, zirconium, and oxygen and can exhibit ferroelectric properties. The HZO 1908 may be deposited by ALD, CVD, PVD or other suitable deposition methods and can have a thickness in the range of 5-50 nm. Acting as a ferroelectric layer, the HZO 1908 enables the non-volatile storage of data in the FeFET 1900.

[00389] Deposited conformally on top of the HZO 1908 is a layer of IWO 1906. IWO 1906 comprises indium, tungsten, and oxygen. It can be deposited by sputtering or other suitable techniques and may have a thickness in various ranges. The IWO 1906 layer serves as a control oxide layer in the FeFET 1900.

[00390] On top of the IWO 1906 layer, a drain contact 1904 and source contact 1914 are formed. The drain contact 1904 and source contact 1914 may comprise metals such as copper, aluminum, or alloys thereof and can be deposited by PVD, CVD or other suitable methods. The drain contact 1904 and source contact 1914 allow electrical connection to the FeFET 1900. They may have thicknesses in the range of 50-500 nm.

[00391] In operation, a voltage applied to the drain 1904, source 1914, and TiN gate contact 1912 can control the ferroelectric polarization of the HZO 1908 layer. The polarization state can be used to store information in a non-volatile manner, enabling memory storage capabilities. The IWO 1906 layer helps improve the switching speed and endurance of the FeFET 1900. Overall, the layered structure shown in Fig. 19 enables a FeFET 1900 suitable for non-volatile memory applications.

[00392] Fig. 20 presents the transfer characteristics of a Ferroelectric FET (FeFET) device, illustrating the relationship between the gate voltage (V_{GS}) on the x-axis and the resulting drain current (I_D) on the y-axis. Fig. 20 may show the characteristics of a FeFET as disclosed herein. The x-axis spans from -1V to 1V, while the y-axis, on a logarithmic scale, displays current values from $10^{-12} \text{A}/\mu\text{m}$ to $10^{-4} \text{A}/\mu\text{m}$.

[00393] Two distinct curves represent the drain current behavior under clockwise (CW) and counterclockwise (MW) polarization of the FeFET. The blue curve (CW) starts at approximately $10^{-11} \text{A}/\mu\text{m}$ at -1V and exhibits a steep increase around -1 V, reaching just

above 10^{-5} A/ μ m at 1V. This demonstrates the rapid increase in drain current exhibited by the FeFET under forward bias in the clockwise polarization state.

[00394] Conversely, the red curve (MW) starts at approximately 10^{-11} A/ μ m at -1V and increases more gradually as it approaches 0V. At around 1V, it then follows closely with the blue curve past 1V. This demonstrates comparable drain current behavior under reverse bias conditions regardless of the polarization state.

[00395] Notably, the separation between the red and blue curves spans several orders of magnitude in the negative voltage range near -1V. This substantial difference in off-state current highlights the non-volatile memory effect achievable with the FeFET depending on its polarization direction. This large memory window is explicitly called out in the green box labeled “Large Memory Window” at the top left.

[00396] Additional key details provided include the FeFET device dimensions, with a width/length ratio of $1\mu\text{m}/50\text{nm}$ specified. The drain voltage is also fixed at 0.05V. Specific points along the curves are annotated, such as “MW @ $5\text{e-}7\text{A}/\mu\text{m} = 1\text{V}$ ” on the red MW curve denoting the 1V memory window at 5×10^{-7} A/ μ m drain current. Another point marked is “CW @ -0.5V = 1×10^{-6} ” on the blue CW curve, highlighting the clockwise current value of 1×10^{-6} A/ μ m at -0.5V gate voltage.

[00397] In summary, Fig. 20 depicts the bidirectional transfer characteristics of the FeFET device, highlighting the large memory window achievable through polarization switching and providing detailed voltage, current, and dimensional specifications to fully convey the measurement conditions and transistor performance. The paired curves effectively compare the clockwise and counterclockwise operation modes over the full gate voltage range.

[00398] Fig. 21 shows a semiconductor package using a memory wafer 2102 and an application wafer 2104 formed into a wafer pair 2106 in accordance with an embodiment of the present disclosure. The semiconductor package 2100 includes multiple wafer pairs 2106, 2108, 211 stacked vertically and bonded together.

[00399] The wafer pair 2106 comprises a memory wafer 2102 bonded face-to-face to an application wafer 2104. The memory wafer 2102 contains modules 2122 grouped into modules groups 2124. Each modules group 2124 includes multiple modules, such as first module 2126 and second module 2128.

[00400] Each module 2122 has an independent read port (not shown) to allow concurrent reading of data from the modules. The modules group 2124 shares a common write port with write peripheral to write data to the modules. The write peripheral contains circuitry such as drivers, decoders, buffers and controllers to facilitate writing data to the modules.

[00401] The application wafer 2104 contains processing elements (not shown) to operate on data read from the memory wafer 2102. The wafers 2102, 2104 may be connected through fine-pitch interconnects like hybrid bonds 2134 to enable high-bandwidth local data transfer between the processing logic and memory modules.

[00402] An interposer or interface may be integrated between wafer pair 2106 and the next wafer pair 2108 to facilitate communication of signals/power between the vertically stacked wafers. Conductive through-silicon vias TSVs 2120 provide vertical connectivity across the wafer stack.

[00403] A heat sink thermal material layer 2118 may be inserted between wafer pairs to dissipate heat. Additional wafer pairs like 2110 can be incorporated for increased functionality, storage capacity, etc. Redistribution layers may be included to route signals/power. The redistribution layers (RDL) on each wafer can be used for routing of electrical signals and/or power throughout the wafer stack. Each RDL may be composed of one or multiple metal layers, separated by dielectric materials (in some embodiments). The metal layers may be made of copper, aluminum, or other metals and/or alloys.

[00404] The RDL layer may include the integration of bypass capacitors on the RDL. These capacitors may be placed to stabilize the power supply by filtering out noise and providing a reservoir of charge to meet transient current demands. The inclusion of these capacitors directly on the RDL reduces the distance between the power source and the active circuits to thereby minimizing inductive and resistive losses.

[00405] The RDL may incorporate power supply Through-Silicon Vias (TSVs) on the back-side of the wafer. These TSVs are vertical interconnects that pass through the silicon wafer, providing a direct and low-resistance path for power supply connections from the top to the bottom of the wafer stack, for example. The back-side TSVs can be used in multi-wafer stacks to reduce the power delivery path length.

[00406] The design of the RDL may be laid out to minimize crosstalk and signal degradation to maintain high-speed data transfer rates. The signal routing on the RDL may be complemented by the use of shielding techniques, such as the incorporation of ground planes and the strategic placement of vias, to further enhance signal integrity, in some embodiments.

[00407] The RDL may be formed using processes such as photolithography for patterning, electroplating for metal deposition, and chemical-mechanical polishing (CMP), among others.

[00408] The memory wafer 2102 comprises multiple modules groups, including a modules group 2124. Each modules group, such as modules group 2122, contains multiple

modules, including a first module 2126 and a second module 2128. The exact number of modules in each modules group may vary in different embodiments.

[00409] The modules group 2124 may have a shared write port that allows data to be written to all modules within the group via a write peripheral. The write port utilizes a write address space that spans across the modules in modules group 2124. The shared write port may be facilitated by the write peripheral, which includes necessary circuitry like drivers, decoders, buffers, and controllers to ensure data is appropriately written to the modules.

[00410] Each module in the modules group 2124, such as the first module 2126 and the second module 2128, has its own dedicated read port. For example, the first module 2126 has a first read port that allows data to be read specifically from the first module 2126. Similarly, the second module 2128 has a second read port for reading data only from the second module 2128.

[00411] The dedicated read ports enable concurrent data reads across multiple modules. For instance, data can be read independently from the first module 2126 via the first read port concurrently with data being read from the second module 2128 via the second read port.

[00412] The memory wafer 2102 may contain multiple such modules groups 2124 with dedicated read ports per module and shared write ports per group to simultaneously enable writes to the group and parallel reads from individual modules. This architecture provides flexibility and performance.

[00413] The memory wafer 2102 is directly bonded to the application wafer 2104 to facilitate high-bandwidth communication between the processing elements on the application wafer 2104 and the data storage offered by the modules on the memory wafer 2102. The application wafer 2104 may contain various integrated circuit components to enable processing and communication with the modules on the memory wafer 2102.

[00414] In some embodiments, the application wafer 2104 includes processing elements such as CPUs, GPUs, DSPs, machine learning accelerators, and other logic circuits to manipulate data stored on and retrieved from the memory wafer 2102. These processing elements may be interconnected through an on-chip network or bus architecture to coordinate data transfers.

[00415] Additionally, the application wafer 2104 may feature an interface to an application programming interface (API) 2114 to facilitate interaction with the memory contents. This API 2114 may abstract the underlying memory implementation and provide a consistent interface through which applications can store and access data to package 2100, to the application wafer 2104, the memory module 2102, to one or more of the wafer pairs 216,

2106, 2110, or some combination thereof. The API 2114 may be a CXI or PCI interface, for example.

[00416] The application wafer 2104 may also comprise data routing circuitry such as the DRAM data input 2112 and DRAM data output 2116. These components manage the flow of data in and out of the wafer pair 2106, enabling communication with external systems. In some embodiments, the DRAM data input 2112 and/or the DRAM data output 2116 can also access the memory wafer 2102.

[00417] In some implementations, the application wafer 2104 contains voltage regulation and power delivery subsystems to provide stable power to the wafer pair 2106. It may also incorporate test, characterization, and built-in self-test (BIST) circuits to validate functionality.

[00418] Furthermore, the application wafer 2104 can have through-silicon vias (TSVs) 2120 and/or redistribution layers to assist with vertical system integration, thermal management, and intra-package signaling. Additional wafers may be stacked above wafer pair 2106 and interconnected using TSVs 2120.

[00419] The semiconductor package 2100 includes multiple wafer pairs, one of which is designated as wafer pair 2106. Wafer pair 2106 comprises a memory wafer 2102 bonded to an application wafer 2104.

[00420] The memory wafer 2102 includes modules, such as module 2122, arranged into modules groups like modules group 2124. The modules serve as data storage components, potentially utilizing non-volatile memory technologies. The modules group 2124 features dedicated data input/output capabilities. For example, as previously mentioned, in some embodiments a DRAM data input 2112 allows external input of data to the modules, while a DRAM data output 2116 enables data retrieval from the modules. In some embodiments, DRAM data input and output 2112, 2116 instead interfaces to the application wafer 2104.

[00421] The application wafer 2104 may contain various processing elements and routing circuitry to facilitate computation using data stored in the bonded memory wafer 2102. For instance, the application wafer 2104 could integrate components like CPUs, GPUs, AI accelerators, FPGAs, and other logic blocks that need to access large volumes of data. In some embodiments, the application programming interface 2114 provides the signaling and protocols enabling communication between the application wafer 2104 and memory wafer 2102 to facilitate data transmission across the wafer pair 2106.

[00422] Multiple wafer pairs 2106, 2108, 2110 including wafer pair 2106, can be integrated into a multi-layered configuration. Heat sink thermal material 2118 may be

incorporated between wafer pairs like 2106, 2108, and 2110 for thermal management. Through-silicon vias 2120 also enable high-density signaling across vertically stacked wafers. Together, these techniques facilitate wafer-scale heterogeneous integration to create computing systems combining processing and storage.

[00423] The wafer pair 2106 is architected to enable customized matching of compute engines (application wafer 2104) to large data repositories (memory wafer 2102), while optimizing metrics like throughput, latency, power efficiency, and reliability through advanced integration. Multiple wafer pairs can be interconnected across device layers to achieve increased capabilities. In some implementations, wafer pair 2106 may be designed based on a specific workload's computational requirements and data access patterns. Overall, wafer pair 2106 facilitates an efficient hardware mapping between processing logic and stored data.

[00424] The DRAM data input 2112 is configured to receive input data that can be written to the semiconductor package 2100. The data received by the DRAM data input 2112 may originate from various sources external to the package 2100, such as data buses, processing elements, sensor inputs, etc. that is coupled to a DRAM memory bank.

[00425] In some embodiments, the DRAM data input 2112 may be implemented as a high-speed interface port supporting standards such as DDR, GDDR, HBM, or custom protocols optimized for the target application. The specific signaling schemes, bus widths, clock rates and other physical layer attributes can be tailored to meet bandwidth, power, and latency requirements.

[00426] The DRAM data input 2112 can feature dedicated connections for each module or subsystem within the package 2100. For example, separate links may exist for the three modules 2102 shown connected to corresponding processing elements in the application wafer 2104. This enables concurrent data transfers to different memories without contention. The links may utilize various topologies such as point-to-point, daisy-chained, muxed or switched interconnects as per design requirements.

[00427] In some implementations, the connectivity between the DRAM data input 2112 and the internal DRAMs in package 2100 may involve intermediate buffers, queues, control logic or internal data networks like crossbar or mesh topologies to enable efficient routing, scheduling and quality-of-service minimums. This infrastructure transparently manages the data movement from input to destination DRAM modules.

[00428] The DRAM data input 2112 may also include circuitry for serializing/deserializing, encoding/decoding, encrypting/decrypting, and integrity checking to transform the incoming data streams appropriately before routing them towards the DRAM

storage elements. Features for handshake protocols, flow control, quality-of-service and traffic class differentiation may also be implemented.

[00429] In high reliability systems, the DRAM data input 2112 can incorporate redundancy and fail-over mechanisms. This provides resiliency against faults in the communication paths to any specific DRAM module. Multiple input ports with routing flexibility help facilitate memory access even in degraded conditions.

[00430] The application programming interface (API) 2114 provides an interface for the application wafer 2104 to communicate with and access the memory on the memory wafer 2102. The API 2114 abstracts the underlying memory implementation and provides a simplified way for the application to store and retrieve data. In some embodiments, the API 2114 can write data to the memory in the memory wafer 2102.

[00431] The API 2114 may provide functions such as OpenMemory, ReadMemory, WriteMemory, and CloseMemory. The OpenMemory function initializes access to a memory module within one of the modules groups 2124 on the memory wafer 2102. It takes as parameters information such as the module ID or address range to open. The ReadMemory and WriteMemory functions actually read and write data to the opened memory module. They take parameters such as a memory address, data buffer, and length. Finally, the CloseMemory function deinitializes access to the memory module when it is no longer needed.

[00432] The API 2114 may handle various interactions with the memory wafer 2102. This includes interfacing with any write ports and read ports to access the desired memory module. The API 2114 can utilize the appropriate ports, commands, and protocols to activate the target module, initiating row/column decoding, sense amplification, and data transfers. The modular memory architecture maximizes parallelism and throughput.

[00433] In some specific embodiments, the API 2114 allows the application wafer 2104 flexibility in reading, writing, and accessing groups of memory modules in the memory wafer 2102 without having to handle low-level responsibilities. In this specific embodiment, the application sees a high-level interface exposing the memory storage capability. The API 2114 may therefore serve as an intermediary ensuring efficient inter- or intra-wafer communication.

[00434] The DRAM data output 2116 refers to the data output connections to external DRAM (Dynamic Random Access Memory) components in communication with the semiconductor package 2100. The DRAM data outputs 2116 provide a means for reading out stored data from the memory arrays and transmitting that information externally for further processing.

[00435] In some embodiments, the DRAM data outputs 2116 may consist of a series of conductive traces, contact pads, solder bumps or pillars formed on the surface of the memory wafer 2102, the application wafer 2104, or the package 2100. These conductive elements can link to internal and/or external sense amplifiers, buffers and peripheral control logic.

[00436] The DRAM data outputs 2116 may also connect to complementary interfaces on the application wafer 2104, allowing data transfer between the memory and application dies in a face-to-face configuration. High-density microbumps or hybrid bonding may be employed to reduce parasitics and enable wide parallel data buses capable of high bandwidth operation.

[00437] In some implementations, the DRAM data outputs 2116 may have a predetermined pitch and size compatible with wafer-scale production and established interface standards. This may ensure interoperability with application dies from various sources and facilitates modular system design. Self-aligning bonding processes precision alignment between the fine-pitch outputs 2116 and application die contacts.

[00438] The Heat Sink Thermal Material layer 2118 is an optional layer that may be included between wafer pair 2106 and wafer pair 2108. The Heat Sink Thermal Material layer 2118 serves to facilitate heat dissipation and cooling between the stacked wafer pairs 2106 and 2108.

[00439] In some embodiments, the Heat Sink Thermal Material layer 2118 may be implemented using materials with high thermal conductivity such as aluminum, copper, graphite sheets, thermal adhesives, or thermal greases/compounds. The thickness and coverage area of the Heat Sink Thermal Material layer 2118 can be adjusted to provide the desired level of heat spreading and cooling capacity.

[00440] Additionally, the Heat Sink Thermal Material layer 2118 may contain microstructures, porous features, fins, pins, or channels to increase the surface area for more effective heat transfer. These enhancements to the geometry of the layer can improve lateral and vertical heat conduction.

[00441] In operation, the Heat Sink Thermal Material layer 2118 absorbs heat generated during the operation of wafer pair 2106 and spreads this thermal energy laterally to prevent localized hot spots. The heat is then conducted in the vertical direction towards wafer pair 2108 which can serve as a heat sink, facilitating cooling through greater heat dissipation across the larger surface area of the wafer pair 2108. This thermal management approach lowers the overall temperature rise within the semiconductor package 2100, enabling reliable operation of the stacked devices.

[00442] The inclusion of the Heat Sink Thermal Material layer 2118 is optional and may be advantageous in certain configurations of the semiconductor package 2100 where cooling demands are greater due to increased power density or device integration. By assisting in heat diffusion and cooling, the layer helps mitigate potential temperature build-up issues within the package.

[00443] The semiconductor package 2100 may include through silicon vias (TSVs) 2120 to facilitate vertical interconnections between the multiple wafer pairs 2106, 2108, 2110. The TSVs 2120 are electrical connections that pass vertically through a wafer, allowing signals and power to be transmitted between vertically stacked wafers.

[00444] In some embodiments, the TSVs 2120 are located around the peripheral of the wafer pairs 2106, 2108, 2110. The peripheral placement allows the central areas of the wafers to be dedicated to active circuits like the modules groups 2124. In other embodiments, the TSVs 2120 may be positioned towards the centers of the wafer pairs to optimize signal routing.

[00445] The TSVs 2120 enable high-density 3D integration of multiple dies or chiplets in the semiconductor package 2100. The TSVs 2120 may have straight, tapered, or variable diameters along their depths. They may be lined with insulating materials like silicon dioxide to prevent electrical leakage. The TSVs 2120 may then be filled with conductive materials such as copper or tungsten to carry signals and power between layers.

[00446] The TSVs 2120 can carry various signals between the wafer pairs 2106, 2108, 2110 including data, address, control, and clock signals. This allows different dies in separate wafer pairs to coordinate their operation. The TSVs also distribute power across the stack, eliminating, reducing, or supplementing the need for external power delivery, in some specific embodiments.

[00447] In the semiconductor package 2100, the TSVs 2120 enable communication between the memory wafer 2102 containing the modules groups 2124 and the application wafer 2104 which utilizes the memory. The TSVs 2120 provide power delivery and data transfer pathways between the stacked wafers. This allows high-bandwidth signaling with reduced latency between the computation and memory elements.

[00448] The number, density, dimensions, and materials of the TSVs 2120 can be optimized to meet thermal, electrical, and mechanical constraints. Additional circuitry may be embedded within the TSV structure for performance monitoring or redundancy.

[00449] The semiconductor package 2100 includes multiple wafer pairs 2106, 2108, 2110 stacked vertically. The wafer pair 2106 comprises a memory wafer 2102 bonded to an application wafer 2104.

[00450] The memory wafer 2102 contains multiple modules groups 2124 including modules such as module 2122. The modules 2122 are arranged in an array configuration across the memory wafer 2102. Each module 2122 functions as an independent data storage unit with dedicated read/write control circuitry and signal ports.

[00451] In some embodiments, the module 2122 is implemented as a microvault containing non-volatile memory cells optimized for high density and fast, parallel data access. The microvault design utilizes vertical 3D integration techniques to stack multiple layers of memory arrays, achieving high capacity within a compact footprint.

[00452] The module 2122 features independent read peripheral circuitry and a dedicated read port (not shown) to enable concurrent, contention-free reads by multiple processing elements on the application wafer 2104. The read peripheral circuitry facilitates the transfer of stored data through the read port to the bonded application wafer 2104.

[00453] Furthermore, the module 2122 contains write control circuitry coupled to a shared write port 2112 on the memory wafer 2102. The shared write port 2112 allows an external processor to write data to any module 2122 within the modules group 2124 by providing a common write interface. Appropriate write addressing schemes ensure data is routed to the correct module destination.

[00454] The module 2122 may exchange data signals with the application wafer 2104 through fine-pitch vertical connections like micro-bumps or through-silicon vias. Advanced integration techniques enable direct high-bandwidth links between each module 2122 and its corresponding processing element on the application wafer 2104.

[00455] Overall, the module 2122 provides an independently-accessible data storage unit optimized for fast parallel reads while also offering external write accessibility through shared write circuitry and signaling ports. Multiple similar modules can be integrated efficiently to construct dense, high-performance memory systems coupled to processing logic.

[00456] The semiconductor package 2100 includes a modules group 2124 on the memory wafer 2102. The modules group 2124 comprises a plurality of modules, including a first module 2126 and a second module 2128. The modules group 2124 has a shared write port (not shown) that is configured to write data to the plurality of modules in the group. This allows a common write logic system, potentially involving high voltage circuitry, to write data that can be shared across the modules.

[00457] Each module in the modules group 2124 also has an independent read port to enable concurrent reads from the modules. For example, the first module 2126 has a first read

port and the second module 2128 has a second read port. This allows each module to be read independently without contention.

[00458] The modules group 2124 may be organized into partitions, each partition having its own read peripheral circuitry and independent clocking. The number of modules in each partition can vary. In some cases, each partition may correspond to a single module.

[00459] The modules within the modules group 2124 can be implemented with various memory technologies, including but not limited to FeFET, ReRAM, SOT and STT. The modules can also have different capacities and speeds as suited for the application.

[00460] When the memory wafer 2102 is bonded to the application wafer 2104, the independent read ports of the modules in modules group 2124 may interface with processing elements on the application wafer 2104. This allows each processing element to have dedicated access to a module, enabling highly parallel and interference-free operation. The modules group 2124, through its shared write port and independent read ports, facilitates an efficient manycore architecture with a shared write space and independent read spaces. This allows data to be updated globally while being read locally without contention.

[00461] Various alternatives and modifications can be devised by those skilled in the art without departing from the disclosure. Accordingly, the present disclosure is intended to embrace all such alternatives, modifications, and variances. Additionally, while several embodiments of the present disclosure have been shown in the drawings and/or discussed herein, it is not intended that the disclosure be limited thereto, as it is intended that the disclosure be as broad in scope as the art will allow and that the specification be read likewise. Therefore, the above description should not be construed as limiting, but merely as exemplifications of particular embodiments. And those skilled in the art will envision other modifications within the scope and spirit of the claims appended hereto. Other elements, steps, methods, and techniques that are insubstantially different from those described above and/or in the appended claims are also intended to be within the scope of the disclosure.

[00462] The embodiments shown in the drawings are presented only to demonstrate certain examples of the disclosure. And the drawings described are only illustrative and are non-limiting. In the drawings, for illustrative purposes, the size of some of the elements may be exaggerated and not drawn to a particular scale. Additionally, elements shown within the drawings that have the same numbers may be identical elements or may be similar elements, depending on the context.

[00463] Where the term "comprising" is used in the present description and claims, it does not exclude other elements or steps. Where an indefinite or definite article is used when

referring to a singular noun, e.g., "a," "an," or "the," this includes a plural of that noun unless something otherwise is specifically stated. Hence, the term "comprising" should not be interpreted as being restricted to the items listed thereafter; it does not exclude other elements or steps, and so the scope of the expression "a device comprising items A and B" should not be limited to devices consisting only of components A and B. This expression signifies that, with respect to the present disclosure, the only relevant components of the device are A and B.

[00464] Furthermore, the terms "first," "second," "third," and the like, whether used in the description or in the claims, are provided for distinguishing between similar elements and not necessarily for describing a sequential or chronological order. It is to be understood that the terms so used are interchangeable under appropriate circumstances (unless clearly disclosed otherwise) and that the embodiments of the disclosure described herein are capable of operation in other sequences and/or arrangements than are described or illustrated herein.

[00465] Each of the characteristics and examples described herein, and combination thereof, may be said to be encompassed by the present disclosure. The present disclosure is thus drawn to the following non-limiting numbered aspects:

[00466] 1. A semiconductor package, comprising: a memory wafer having a plurality of modules groups including a modules group, the modules group having: a plurality of modules including a first module and a second module; a write port configured to write to the modules group; a first read port configured to read from the first module; and a second read port configured to read from the second module; and an application wafer bonded to the memory wafer, the application wafer in operative communication with the plurality of modules groups.

[00467] 2. The semiconductor package according to aspect 1, wherein an outer perimeter of the memory wafer and an outer parameter of the application wafer are coextensive relative to each other.

[00468] 3. The semiconductor package according to aspect 1, wherein the application wafer includes at least one of an FPGA core, a microcontroller, a GPU processing element, an ASIC, a CPU core, a neural network accelerator, a digital signal processor (DSP), a reconfigurable processing unit, a machine learning inference engine, a tensor processing unit (TPU), a system-on-chip (SoC), a network processing unit (NPU), an embedded processor, a security module, and a programmable logic device.

[00469] 4. The semiconductor package according to aspect 1, wherein the write port is a shared write port configured to write to the first and second modules.

- [00470] 5. The semiconductor package according to aspect 4, wherein the shared write port is configured to write to an address space, the shared write port is configured to write to the first module via a first portion of the address space and write to the second module via a second portion of the address space.
- [00471] 6. The semiconductor package according to aspect 1, wherein the write port is a shared write port configured to write to only the modules group.
- [00472] 7. The semiconductor package according to aspect 1, wherein the plurality of modules groups includes at least one write-group, wherein each write-group of the at least one write-group includes a respective read port and a common write port.
- [00473] 8. The semiconductor package according to aspect 7, wherein the at least one write-group includes a first write group and a second write group, wherein the first write group comprises about a first half of the plurality of modules groups and the second write group comprises about a second half of the plurality of modules groups.
- [00474] 9. The semiconductor package according to aspect 8, wherein the first write group is the modules group.
- [00475] 10. The semiconductor package according to aspect 1, wherein the modules group is a write group and the write port is a common write port for the write group.
- [00476] 11. The semiconductor package according to aspect 10, wherein the write group includes all modules of the memory wafer.
- [00477] 12. The semiconductor package according to aspect 1, wherein the memory wafer is cut into a rectangle.
- [00478] 13. The semiconductor package according to aspect 1, wherein the memory wafer is cut into a square.
- [00479] 14. The semiconductor package according to aspect 1, wherein each module of the plurality of modules is a microvault.
- [00480] 15. The semiconductor package according to aspect 1, wherein each module of the plurality of modules includes an independent read port for concurrent reading via a respective independent read port of any of the plurality of modules whereby the first read port is configured to read from the first module concurrently with another read of the second module using the second read port.
- [00481] 16. The semiconductor package according to aspect 1, wherein the modules group is configured to have a single write address space configured to write to the plurality of modules.

[00482] 17. The semiconductor package according to aspect 16, wherein the first read port is configured to have a first read address space and the second read port is configured to have a second read address space.

[00483] 18. The semiconductor package according to aspect 17, wherein the first read address space numerically overlaps with the second read address space.

[00484] 19. The semiconductor package according to aspect 17, wherein the first read address space is coextensive with the second read address space.

[00485] 20. The semiconductor package according to aspect 17, wherein the first read address space is contiguous with the second read address space.

[00486] 21. The semiconductor package according to aspect 1, wherein the application wafer includes a plurality of processing modules.

[00487] 22. The semiconductor package according to aspect 1, further comprising a redistribution layer on one of the memory wafer and the application wafer to provide power.

[00488] 23. The semiconductor package according to aspect 1, further comprising another wafer bonded to the memory wafer to provide power supply routing to all of the plurality of modules.

[00489] 24. The semiconductor package according to aspect 1, wherein the application wafer includes an application programming interface.

[00490] 25. The semiconductor package according to aspect 1, further comprising a plurality of vias coupled to the memory wafer and the application wafer.

[00491] 26. The semiconductor package according to aspect 25, wherein the plurality of vias are on a peripheral of the memory wafer.

[00492] 27. The semiconductor package according to aspect 25, wherein the plurality of vias is on a center of the memory wafer.

[00493] 28. The semiconductor package according to aspect 1, further comprising a heat sink material interposed between the application wafer and the memory wafer.

[00494] 29. The semiconductor package according to aspect 1, wherein the application wafer includes a Network-on-Chip with routers interconnecting a plurality of processing elements.

[00495] 30. The semiconductor package according to aspect 1, wherein the memory wafer has a single write-grouping with a single write port, the write port being the single write port.

[00496] 31. The semiconductor package according to aspect 1, wherein the plurality of modules is configured for a write endurance of at least once a day for 10 years.

- [00497] 32. The semiconductor package according to aspect 1, wherein the plurality of modules is configured for a write endurance of at least once a week for 10 years.
- [00498] 33. The semiconductor package according to aspect 1, wherein the memory wafer is formed from a complete slice of an ingot.
- [00499] 34. The semiconductor package according to aspect 1, wherein the memory wafer is round.
- [00500] 35. The semiconductor package according to aspect 1, wherein the memory wafer is a 12-inch wafer and includes an aggregate memory of the plurality of modules configured to provide at least 10 trillion bytes of storage.
- [00501] 36. The semiconductor package according to aspect 1, wherein the application wafer and the memory wafer are interconnected using bump-less bonding technology to thereby reduce thermal and electrical resistance.
- [00502] 37. The semiconductor package according to aspect 1, wherein the memory wafer includes memory repair mechanisms configured for in-situ repair of faulty memory modules.
- [00503] 38. The semiconductor package according to aspect 1, wherein each module of the plurality of modules is configured to include fault-tolerance routing to thereby enable a module of the plurality of modules to route data around a failed module within the memory wafer.
- [00504] 39. The semiconductor package according to aspect 38, wherein a corresponding processing element is determined if be faulty when the failed module is determined to be faulted.
- [00505] 40. The semiconductor package according to aspect 1, further comprising an interconnection between the memory wafer and the application wafer configured to utilize through-silicon vias (TSV) configured to form interconnected buses therebetween.
- [00506] 41. The semiconductor package according to aspect 1, further comprising an AI inference fabric disposed on the application wafer.
- [00507] 42. The semiconductor package according to aspect 1, wherein the plurality of modules are configured to support deep neural network operations through large scale dataflow mapping.
- [00508] 43. The semiconductor package according to aspect 1, wherein each module of the plurality of modules is configured to receive configuration data corresponding to a spatial dataflow to thereby optimize an execution of an AI inference task.

[00509] 44. The semiconductor package according to aspect 1, wherein the first read port is disposed on a surface of the memory wafer, the first module comprising: a memory array comprising a plurality of memory unit cells arranged within the first module; and a read peripheral circuitry configured to read data stored within the memory array via the first read port.

[00510] 45. The semiconductor package according to aspect 44, wherein a footprint on the surface of the read peripheral circuitry overlaps with a footprint on the surface of the memory array.

[00511] 46. The semiconductor package according to aspect 44, wherein a footprint on the surface of the first read port overlaps a footprint on the surface of the memory array.

[00512] 47. The semiconductor package according to aspect 44, wherein a footprint on the surface of the read peripheral circuitry is coextensive with a footprint on the surface of the memory array.

[00513] 48. The semiconductor package according to aspect 44, wherein a footprint on the surface of the first read port is coextensive with a footprint on the surface of the memory array.

[00514] 49. The semiconductor package according to aspect 44, wherein the memory array, the read peripheral circuitry, and the first read port are disposed vertically to minimize a module footprint on the surface of the memory wafer

[00515] 50. The semiconductor package according to aspect 44, further comprising a write peripheral circuitry connected to the write port on the surface of the memory wafer, wherein: the write peripheral circuitry defines a first footprint on the surface, the first footprint does not overlap with a footprint of the memory array and the first footprint does not overlap with a footprint of the first read port, and the write port defines a second footprint on the surface, the second footprint does not overlap with the footprint of the memory array and does not overlap with the footprint of the first read port.

[00516] 51. The semiconductor package according to aspect 50, wherein the first footprint overlaps the second footprint.

[00517] 52. The semiconductor package according to aspect 1, further comprising a plurality of wafer pairs, wherein the application wafer and the memory wafer form a first wafer pair of the plurality of wafer pairs.

[00518] 53. The semiconductor package according to aspect 52, further comprising a plurality of through-silicon vias connecting the first wafer pair to a second wafer pair of the plurality of wafer pairs.

- [00519] 54. The semiconductor package according to aspect 53, wherein the plurality of through-silicon vias are disposed on a peripheral of the first wafer pair.
- [00520] 55. The semiconductor package according to aspect 53, wherein the plurality of through-silicon vias are disposed on a central portion of the first wafer pair.
- [00521] 56. The semiconductor package according to aspect 53, further comprising a heatsink disposed between the first and second wafer pairs.
- [00522] 57. The semiconductor package according to aspect 53, further comprising a redistribution layer disposed between the first and second wafer pairs.
- [00523] 58. The semiconductor package according to aspect 57, wherein the redistribution layer supplies power to one of the first and second wafer pairs via a through-silicon via.
- [00524] 59. The semiconductor package according to aspect 53, further comprising another redistribution layer disposed on a second side of the first wafer pair.
- [00525] 60. A method of forming a semiconductor package, comprising: forming a memory wafer having a plurality of modules groups including a modules group, the modules group having: a plurality of modules including a first module and a second module; a write port configured to write to the modules group; a first read port configured to read from the first module; and a second read port configured to read from the second module; and bonding an application wafer to the memory wafer, the application wafer being in operative communication with the plurality of modules groups.
- [00526] 61. The method of aspect 60, further comprising forming the memory wafer and the application wafer such that an outer perimeter of the memory wafer and an outer perimeter of the application wafer are coextensive relative to each other.
- [00527] 62. The method of aspect 60, wherein the application wafer includes at least one of an FPGA core, a microcontroller, a GPU processing element, an ASIC, a CPU core, a neural network accelerator, a digital signal processor (DSP), a reconfigurable processing unit, a machine learning inference engine, a tensor processing unit (TPU), a system-on-chip (SoC), a network processing unit (NPU), an embedded processor, a security module, and a programmable logic device.
- [00528] 63. The method of aspect 60, wherein the write port is a shared write port configured to write to the first and second modules.
- [00529] 64. The method of aspect 63, further comprising forming the shared write port to write to an address space, wherein the shared write port is configured to write to the first

module via a first portion of the address space and write to the second module via a second portion of the address space.

[00530] 65. The method of aspect 60, wherein the write port is a shared write port configured to write to only the modules group.

[00531] 66. The method of aspect 60, further comprising forming at least one write-group in the plurality of modules groups, wherein each write-group includes a respective read port and a common write port.

[00532] 67. The method of aspect 66, further comprising forming a first write group and a second write group in the at least one write-group, wherein the first write group comprises about a first half of the plurality of modules groups and the second write group comprises about a second half of the plurality of modules groups.

[00533] 68. The method of aspect 67, wherein the first write group is the modules group.

[00534] 69. The method of aspect 60, further comprising forming the modules group as a write group and the write port as a common write port for the write group.

[00535] 70. The method of aspect 69, wherein the write group includes all modules of the memory wafer.

[00536] 71. The method of aspect 60, further comprising cutting the memory wafer into a rectangle.

[00537] 72. The method of aspect 60, further comprising cutting the memory wafer into a square.

[00538] 73. The method of aspect 60, wherein each module of the plurality of modules is a microvault.

[00539] 74. The method of aspect 60, further comprising forming each module of the plurality of modules to include an independent read port for concurrent reading via a respective independent read port of any of the plurality of modules, whereby the first read port is configured to read from the first module concurrently with another read of the second module using the second read port.

[00540] 75. The method of aspect 60, further comprising forming the modules group to have a single write address space configured to write to the plurality of modules.

[00541] 76. The method of aspect 75, further comprising forming the first read port to have a first read address space and the second read port to have a second read address space.

[00542] 77. The method of aspect 76, further comprising forming the first read address space to numerically overlap with the second read address space.

- [00543] 78. The method of aspect 76, further comprising forming the first read address space to be coextensive with the second read address space.
- [00544] 79. The method of aspect 76, further comprising forming the first read address space to be contiguous with the second read address space.
- [00545] 80. The method of aspect 60, further comprising forming a plurality of processing modules in the application wafer.
- [00546] 81. The method of aspect 60, further comprising providing a redistribution layer on one of the memory wafer and the application wafer to provide power.
- [00547] 82. The method of aspect 60, further comprising bonding another wafer to the memory wafer to provide power supply routing to all of the plurality of modules.
- [00548] 83. The method of aspect 60, further comprising forming an application programming interface in the application wafer.
- [00549] 84. The method of aspect 60, further comprising coupling a plurality of vias to the memory wafer and the application wafer.
- [00550] 85. The method of aspect 84, further comprising disposing the plurality of vias on a peripheral of the memory wafer.
- [00551] 86. The method of aspect 84, further comprising disposing the plurality of vias on a center of the memory wafer.
- [00552] 87. The method of aspect 60, further comprising interposing a heat sink material between the application wafer and the memory wafer.
- [00553] 88. The method of aspect 60, further comprising forming a Network-on-Chip with routers interconnecting a plurality of processing elements in the application wafer.
- [00554] 89. The method of aspect 60, further comprising forming the memory wafer to have a single write-grouping with a single write port, the write port being the single write port.
- [00555] 90. The method of aspect 60, further comprising configuring the plurality of modules for a write endurance of at least once a day for 10 years.
- [00556] 91. The method of aspect 60, further comprising con forming figuring the plurality of modules for a write endurance of at least once a week for 10 years.
- [00557] 92. The method of aspect 60, further comprising forming the memory wafer from a complete slice of an ingot.
- [00558] 93. The method of aspect 60, further comprising forming the memory wafer as a round wafer.

[00559] 94. The method of aspect 60, further comprising configuring the memory wafer as a 12-inch wafer and including an aggregate memory of the plurality of modules configured to provide at least 10 trillion bytes of storage.

[00560] 95. The method of aspect 60, further comprising interconnecting the application wafer and the memory wafer using bump-less bonding technology to thereby reduce thermal and electrical resistance.

[00561] 96. The method of aspect 60, further comprising forming memory repair mechanisms in the memory wafer configured for in-situ repair of faulty memory modules.

[00562] 97. The method of aspect 60, further comprising forming each module of the plurality of modules to include fault-tolerance routing to thereby enable a module of the plurality of modules to route data around a failed module within the memory wafer.

[00563] 98. The method of aspect 97, further comprising determining a corresponding processing element to be faulty when the failed module is determined to be faulted.

[00564] 99. The method of aspect 60, further comprising interconnecting the memory wafer and the application wafer using through-silicon vias (TSV) configured to form interconnected buses therebetween.

[00565] 100. The method of aspect 60, further comprising forming an AI inference fabric on the application wafer.

[00566] 101. The method of aspect 60, further comprising configuring the plurality of modules to support deep neural network operations through large scale dataflow mapping.

[00567] 102. The method of aspect 60, further comprising forming each module of the plurality of modules to receive configuration data corresponding to a spatial dataflow to thereby optimize an execution of an AI inference task.

[00568] 103. The method of aspect 60, wherein the first read port is disposed on a surface of the memory wafer, further comprising: forming a memory array comprising a plurality of memory unit cells arranged within the first module; and forming a read peripheral circuitry to read data stored within the memory array via the first read port.

[00569] 104. The method of aspect 103, further comprising forming a footprint on the surface of the read peripheral circuitry to overlap with a footprint on the surface of the memory array.

[00570] 105. The method of aspect 103, further comprising forming a footprint on the surface of the first read port to overlap with a footprint on the surface of the memory array.

[00571] 106. The method of aspect 103, further comprising forming a footprint on the surface of the read peripheral circuitry to be coextensive with a footprint on the surface of the memory array.

[00572] 107. The method of aspect 103, further comprising forming a footprint on the surface of the first read port to be coextensive with a footprint on the surface of the memory array.

[00573] 108. The method of aspect 103, further comprising forming the memory array, the read peripheral circuitry, and the first read port vertically to minimize a module footprint on the surface of the memory wafer.

[00574] 109. The method of aspect 103, further comprising forming a write peripheral circuitry connected to the write port on the surface of the memory wafer, wherein: the write peripheral circuitry defines a first footprint on the surface, the first footprint does not overlap with a footprint of the memory array and the first footprint does not overlap with a footprint of the first read port, and the write port defines a second footprint on the surface, the second footprint does not overlap with the footprint of the memory array and does not overlap with the footprint of the first read port.

[00575] 110. The method of aspect 109, further comprising forming the first footprint to overlap the second footprint.

[00576] 111. The method of aspect 60, further comprising forming a plurality of wafer pairs, wherein the application wafer and the memory wafer form a first wafer pair of the plurality of wafer pairs.

[00577] 112. The method of aspect 111, further comprising connecting the first wafer pair to a second wafer pair of the plurality of wafer pairs using a plurality of through-silicon vias.

[00578] 113. The method of aspect 112, further comprising disposing the plurality of through-silicon vias on a peripheral of the first wafer pair.

[00579] 114. The method of aspect 112, further comprising disposing the plurality of through-silicon vias on a central portion of the first wafer pair.

[00580] 115. The method of aspect 112, further comprising disposing a heatsink between the first and second wafer pairs.

[00581] 116. The method of aspect 112, further comprising disposing a redistribution layer between the first and second wafer pairs.

[00582] 117. The method of aspect 116, further comprising supplying power to one of the first and second wafer pairs via a through-silicon via using the redistribution layer.

[00583] 118. The method of aspect 112, further comprising disposing another redistribution layer on a second side of the first wafer pair.

[00584] 119. A method of using a semiconductor package, comprising: writing to a modules group using a write port, the modules group having a plurality of modules including a first module and a second module; reading from the first module using a first read port; reading from the second module using a second read port; and communicating operatively between a memory wafer having the modules group and an application wafer bonded to the memory wafer.

[00585] 120. The method of aspect 119, further comprising aligning an outer perimeter of the memory wafer to be coextensive with an outer parameter of the application wafer.

[00586] 121. The method of aspect 119, further comprising writing to the first and second modules using a shared write port.

[00587] 122. The method of aspect 121, further comprising configuring the shared write port to write to an address space, wherein the shared write port writes to the first module via a first portion of the address space and to the second module via a second portion of the address space.

[00588] 123. The method of aspect 119, further comprising configuring a shared write port to write to only the modules group.

[00589] 124. The method of aspect 119, further comprising configuring a plurality of modules groups, each write-group having a respective read port and a common write port.

[00590] 125. The method of aspect 124, further comprising configuring the modules groups into a first write group and a second write group, wherein the first write group comprises about a first half of the plurality of modules groups and the second write group comprises about a second half of the plurality of modules groups.

[00591] 126. The method of aspect 125, further comprising configuring the first write group as the modules group.

[00592] 127. The method of aspect 119, further comprising configuring the modules group as a write group and utilizing a common write port for the write group.

[00593] 128. The method of aspect 127, further comprising configuring all modules of the memory wafer in the write group.

[00594] 129. The method of aspect 119, further comprising configuring an independent read port for each module of the plurality of modules, enabling concurrent reading via respective independent read ports.

- [00595] 130. The method of aspect 119, further comprising configuring the modules group to have a single write address space for writing to the plurality of modules.
- [00596] 131. The method of aspect 130, further comprising configuring the first read port to have a first read address space and the second read port to have a second read address space.
- [00597] 132. The method of aspect 131, further comprising configuring the first read address space to numerically overlap with the second read address space.
- [00598] 133. The method of aspect 131, further comprising configuring the first read address space to be coextensive with the second read address space.
- [00599] 134. The method of aspect 131, further comprising configuring the first read address space to be contiguous with the second read address space.
- [00600] 135. The method of aspect 119, further comprising configuring a plurality of processing modules on the application wafer.
- [00601] 136. The method of aspect 119, further comprising supplying power to one of the memory wafer and the application wafer using a redistribution layer.
- [00602] 137. The method of aspect 119, further comprising routing power to all of the plurality of modules.
- [00603] 138. The method of aspect 119, further comprising querying an application programming interface on the application wafer.
- [00604] 139. The method of aspect 119, further comprising communicating between a plurality of processing elements on the application wafer via a Network-on-Chip with routers.
- [00605] 140. The method of aspect 119, further comprising configuring the memory wafer to have a single write-grouping with a single write port.
- [00606] 141. The method of aspect 119, further comprising configuring the plurality of modules for a write endurance of at least once a day for 10 years.
- [00607] 142. The method of aspect 119, further comprising configuring the plurality of modules for a write endurance of at least once a week for 10 years.
- [00608] 143. The method of aspect 119, further comprising communicating between the application wafer and the memory wafer.
- [00609] 144. The method of aspect 119, further comprising detecting a faulty memory module.
- [00610] 145. The method of aspect 119, further comprising routing around a faulted module within the modules group.

[00611] 146. The method of aspect 145, further comprising determining a corresponding processing element to be faulty when a module is determined to have failed.

[00612] 147. The method of aspect 119, further comprising communicating between the memory wafer and the application wafer using through-silicon vias (TSV) via an interconnected bus.

[00613] 148. The method of aspect 119, further comprising configuring each module of the plurality of modules to receive configuration data corresponding to a spatial dataflow to optimize the execution of an AI inference task.

What is Claimed is:

1. A semiconductor package, comprising:
 - a memory wafer having a plurality of modules groups including a modules group, the modules group having:
 - a plurality of modules including a first module and a second module;
 - a write port configured to write to the modules group;
 - a first read port configured to read from the first module; and
 - a second read port configured to read from the second module; and
 - an application wafer bonded to the memory wafer, the application wafer in operative communication with the plurality of modules groups.
2. The semiconductor package according to claim 1, wherein an outer perimeter of the memory wafer and an outer parameter of the application wafer are coextensive relative to each other.
3. The semiconductor package according to claim 1, wherein the application wafer includes at least one of an FPGA core, a microcontroller, a GPU processing element, an ASIC, a CPU core, a neural network accelerator, a digital signal processor (DSP), a reconfigurable processing unit, a machine learning inference engine, a tensor processing unit (TPU), a system-on-chip (SoC), a network processing unit (NPU), an embedded processor, a security module, and a programmable logic device.
4. The semiconductor package according to claim 1, wherein the write port is a shared write port configured to write to the first and second modules.
5. The semiconductor package according to claim 4, wherein the shared write port is configured to write to an address space, the shared write port is configured to write to the first module via a first portion of the address space and write to the second module via a second portion of the address space.
6. The semiconductor package according to claim 1, wherein the write port is a shared write port configured to write to only the modules group.
7. The semiconductor package according to claim 1, wherein the plurality of modules groups includes at least one write-group, wherein each write-group of the at least one write-group includes a respective read port and a common write port.
8. The semiconductor package according to claim 7, wherein the at least one write-group includes a first write group and a second write group, wherein the first write group

comprises about a first half of the plurality of modules groups and the second write group comprises about a second half of the plurality of modules groups.

9. The semiconductor package according to claim 8, wherein the first write group is the modules group.

10. The semiconductor package according to claim 1, wherein the modules group is a write group and the write port is a common write port for the write group.

11. The semiconductor package according to claim 1, wherein each module of the plurality of modules includes an independent read port for concurrent reading via a respective independent read port of any of the plurality of modules whereby the first read port is configured to read from the first module concurrently with another read of the second module using the second read port.

12. The semiconductor package according to claim 1, wherein the modules group is configured to have a single write address space configured to write to the plurality of modules.

13. The semiconductor package according to claim 12, wherein the first read port is configured to have a first read address space and the second read port is configured to have a second read address space.

14. The semiconductor package according to claim 13, wherein the first read address space numerically overlaps with the second read address space.

15. The semiconductor package according to claim 13, wherein the first read address space is coextensive with the second read address space.

16. The semiconductor package according to claim 13, wherein the first read address space is contiguous with the second read address space.

17. The semiconductor package according to claim 1, wherein the application wafer includes a plurality of processing modules.

18. The semiconductor package according to claim 1, further comprising a redistribution layer on one of the memory wafer and the application wafer to provide power.

19. The semiconductor package according to claim 1, further comprising another wafer bonded to the memory wafer to provide power supply routing to all of the plurality of modules.

20. The semiconductor package according to claim 1, wherein the application wafer includes an application programming interface.

21. The semiconductor package according to claim 1, further comprising a plurality of vias coupled to the memory wafer and the application wafer.

22. The semiconductor package according to claim 21, wherein the plurality of vias are on a peripheral of the memory wafer.
23. The semiconductor package according to claim 21, wherein the plurality of vias is on a center of the memory wafer.
24. A method of forming a semiconductor package, comprising:
forming a memory wafer having a plurality of modules groups including a modules group, the modules group having:
a plurality of modules including a first module and a second module;
a write port configured to write to the modules group;
a first read port configured to read from the first module; and
a second read port configured to read from the second module; and
bonding an application wafer to the memory wafer, the application wafer being in operative communication with the plurality of modules groups.
25. The method of claim 24, further comprising forming the memory wafer and the application wafer such that an outer perimeter of the memory wafer and an outer perimeter of the application wafer are coextensive relative to each other.
26. The method of claim 24, wherein the application wafer includes at least one of an FPGA core, a microcontroller, a GPU processing element, an ASIC, a CPU core, a neural network accelerator, a digital signal processor (DSP), a reconfigurable processing unit, a machine learning inference engine, a tensor processing unit (TPU), a system-on-chip (SoC), a network processing unit (NPU), an embedded processor, a security module, and a programmable logic device.
27. The method of claim 24, wherein the write port is a shared write port configured to write to the first and second modules.
28. The method of claim 27, further comprising forming the shared write port to write to an address space, wherein the shared write port is configured to write to the first module via a first portion of the address space and write to the second module via a second portion of the address space.
29. The method of claim 24, wherein the write port is a shared write port configured to write to only the modules group.
30. The method of claim 24, further comprising forming at least one write-group in the plurality of modules groups, wherein each write-group includes a respective read port and a common write port.

31. The method of claim 30, further comprising forming a first write group and a second write group in the at least one write-group, wherein the first write group comprises about a first half of the plurality of modules groups and the second write group comprises about a second half of the plurality of modules groups.
32. The method of claim 31, wherein the first write group is the modules group.
33. The method of claim 24, further comprising forming the modules group as a write group and the write port as a common write port for the write group.