

ABOVE MOTHERBOARD INTERPOSER WITH PERIPHERAL CIRCUITS

BACKGROUND

Central processing units in all different types of computing devices continue to
5 increase in operating speed, with similar gains occurring in peripheral circuits such as
memories, peripheral computing interface (PCI) circuits, graphics processors, and
many others. However, the gains in speed have begun to run into external physical
limitations that hamper performance by capping the data transfer speed between the
processor and its associated peripheral devices, including the memory. These
10 limitations include signal integrity problems due to high capacitance/high resistance
signal lines, tight line spacing on printed circuit boards which leads to crosstalk,
connector and circuit trace discontinuities in the signal path, and increased power
consumption leading to higher device temperatures and heat management problems.
Dealing with these factors often results in lengthened, more expensive design cycles,
15 more costly motherboard layout and production, and more complicated and costly
cooling techniques.

One can view signal integrity as the measure of confidence that a data value
received at one end of a signal line matches the data value transmitted from the other
end. As the transmission speed increases, previously inconsequential circuit non-
20 idealities such as cross-talk, line capacitance, and impedance discontinuities, have a
greater impact on the signal fidelity, which limits the achievable speed to a level
below that what the native semiconductor devices could otherwise support.

The number of elements within CPU and memory devices, and their native
switching speed, has increased along the way to increasing processing speed, and the
25 number of Input-Output (I/O) pins has similarly increased to provide wider data paths

and, hence, transfer rate, to the outside world; however, the general package size remains the same or smaller. This results in more and more dense CPU connector pin arrays, which in turn forces the signal lines between external devices and the CPU to be placed closer and closer together. This leads to increased line-to-line coupling,
5 resulting in greater cross-talk and reduced signal integrity.

Other characteristics associated with current architectures also contribute to reduced signal integrity. The push to increase memory capacity results in more and more memory devices being connected to the memory bus, often resulting in longer bus lengths to accommodate more memory modules. Longer lines, however, come at a
10 price, which includes relatively high capacitance in the signal paths, relatively high series resistance of the signal paths, an increasing number of discontinuities in the signal path, and greater disparity in time delays between the driving end of the bus and the nearest and farthest load or destination devices.

In a typical memory array, precise synchronization of send and receive
15 activities is vital to the proper overall operation. As more and more memory or other peripheral circuits or devices are connected along a bus, the disparity in the time delay between the CPU and the nearest destination or load device, and the CPU and the farthest destination or load device, becomes problematic. For proper synchronization, all devices must wait until certain signals have been received, thus limiting the speed
20 of the entire network to that of the slowest (farthest away) device. Various techniques are used to attempt to mitigate this so-called differential delay problem, but in general, the longer the line the slower the overall performance.

The signal lines between the CPU and memory or other peripherals may be
crafted as transmission lines having a specific line impedance, with driving and load
25 end terminations, or as simple wires. Transmission lines generally provide the highest

speed signal transmission environment, but the requirement that they be resistively terminated for proper signal integrity mandates the use of extra power supplies and termination resistors, and an attendant increase in power. If properly terminated an ideal transmission line can provide good signal fidelity independent of length;

5 however, in the case of a memory system, each connection, or “Tap”, along the line, and the series resistance inherent in the wire itself, disrupt or attenuate the signal along the line and place a limit on the useful length at any given frequency of operation.

Signal lines may also be designed as simple wires, in which case they do not
10 require termination; however, their performance is now limited by the series resistance of the metal, and the parasitic capacitance to ground and/or other nearby wires.

The series resistance in each signal line consumes a portion of the signal power present at the input end of the line, converting the electrical energy into heat
15 and in the process attenuating the signal. The longer the line the greater the signal attenuation, until at some point the receiving end can no longer recognize the information encoded on the signal. This imposes limits on the length of the signal path, which will probably be different than in the case of a transmission line environment described above and, therefore, the overall size of the memory.
20 Amplifiers, buffers, and repeaters may alleviate this issue, but these all contribute to increased overall component count and board area, increased power and thermal load, and increased cost and complexity of the system.

Signal lines crafted as simple wires are also sensitive to, and limited by, their capacitance to ground and surrounding conductors. As the length of the line increases,
25 or the physical distance to ground or other conductors is reduced, or the number of

taps increases, the line capacitance increases. Higher capacitance, coupled with higher resistance, results in lower signal path bandwidth, which in the case of a digital pulse train slows the rising and falling edges. As the edges slow the pulse stretches out in time, so fewer bits of data can be transmitted in a given interval and data transfer rate drops or become limited. As the pulse transmission rate is increased the data bits can run together resulting in errors due to so-called inter-symbol interference.

Discontinuities typically consist of transitions in the size and thickness of the metal or other conductive material forming the signal path. For example, a signal line transition to the contact pad of a through-via on a printed circuit board, a trace turning a corner, a signal line intersecting the contact pad for a memory connector, or for a pin on the ball grid array or land grid array on an integrated circuit, all cause discontinuities. These change the local impedance of a transmission line on the one hand, or add capacitance to a simple wire on the other hand, both of which distort the shape of the signal due to reflections or simply reduce the bandwidth and, hence, the amplitude.. All of these effects impact the signal integrity, in turn leading to data errors that require compensation such as periodic retransmission of the signals, confirmation signaling, or other typical error correction approached.

One approach to handling signal integrity problems involves boosting the power of the signal to compensate for attenuation or distortion, resulting in higher power consumption for the CPU and the peripheral circuits. This also produces a higher level of heat to manage. Heat in general causes a degradation in the semiconductor devices due to increased current leakage within the transistors, as well as a general reduction in transistor operating speed due to decreased conductor mobility. This mitigation approach can result in a cycle of boosting the power to overcome the loss of speed, generating more heat because of the increased power,

which then requires a further increase in power and heat, until one reaches a practical or economic limit.

Other error mitigation approaches involve reducing the speed of the transmissions, thereby reducing the impact of capacitive loading on the lines, as well
5 as reducing the power required to drive the lines at higher speeds. However, this approach runs counter to the desire for faster speeds and more performance.

Still other mitigation approaches stack the memory chips and/or the processor chips to shorten the communication paths. However, these require custom designed motherboards and non-traditional system architectures, making them impractical for
10 typically uses.

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 shows an embodiment of an above motherboard interposer.

Figure 2 shows a plan view of an embodiment of an above motherboard interposer.

15 Figure 3 shows an embodiment of a motherboard having a connector.

Figure 4 shows an embodiment of a CPU mounted into the connector with an interposer.

Figure 5 shows an embodiment of a flexible interposer.

Figures 6A and 6B shows an embodiment of an alternative packaging
20 arrangement.

Figures 7A and 7B show an alternative embodiment of an alternative packaging arrangement.

Figure 8 shows a side view of an embodiment of an interposer with spring contacts.

Figures 9 and 10 show side views of an embodiment of an interposer having elastomer contacts.

Figure 11 shows a prior art fanout length of DIMM connectors on a motherboard.

5 Figure 12 shows a fanout length of an above board interposer.

Figure 13 shows a prior art embodiment of a signal path having multiple different discontinuities.

Figure 14 shows the TDR response of an embodiment of an above board interposer signal line.

10 Figure 15 displays a TDR signature of the signal path of a conventional motherboard to memory module path.

Figures 16-19 show embodiments of the above board interposer configured to ensure equal path lengths.

Figure 20 shows an embodiment of a motherboard with typical DRAM
15 memory wiring.

Figure 21 shows an embodiment of a motherboard with DRAM memory wiring eliminated.

Figure 22 shows an embodiment of an above board interposer having peripheral circuits to function as off chip DRAM cache.

20 Figure 23 shows an embodiment of an interposer having short length memory thermal paths to a CPU heat sink.

Figure 24 shows an embodiment of the interposer having peripheral circuits in direct contact with a heat sink.

Figure 25 shows an embodiment of an interposer having peripheral circuits connected to the heat sink with a graphite thermal sheet material or other similar heat transfer media.

Figure 26 shows an embodiment of an interposer having an air-cooled heat
5 sink.

Figure 27 shows an embodiment having an interposer coupled to another circuit.

Figure 28 shows an embodiment of an interposer mated to a CPU pin grid array.

10 DETAILED DESCRIPTION OF THE EMBODIMENTS

Figure 1 shows an embodiment of an “above motherboard” or “above board” interposer. The term “above motherboard” refers to the location of the interposer as being up above the socket or mounting plane for the central processing unit. The term “central processing unit” may be the main processor for a computing device or may be
15 the processor on any kind of accessory board, such as a video or audio board, although the term may be referred to as a CPU. The term “above motherboard” also refers to the configurations of the connections between the processing unit and peripheral circuits. The peripheral circuits here are those that reside upon the above board interposer, to which the connections are made through the interposer, rather than
20 through the motherboard. While several examples are given below, the peripheral circuits may be any type of integrated circuit and no limitation to any particular type is implied.

Figure 1 shows a CPU package 10 mounted to a motherboard 12 through a connector 16. In this instance the connector consists of a land grid array (LGA) using
25 springs such as 20 to make connections between the array of connections on the CPU

package and the array of connectors such as 18 and 22 on the motherboard to which the processor will connect. While in this embodiment the connectors are LGA connections, but may also consist of a ball grid array (BGA), a pin grid array (PGA), elastomer contacts, etc., for either the connector between the CPU and the above
5 board interposer or between the above board interposer and the motherboard. For purposes of this discussion, any mechanisms that connect between the motherboard, the interposer, and the CPU package will be referred to as interconnects, whether they be pads, BGAs, PGAs, LGAs, etc.

In prior art, all of the connections between the processor and the motherboard
10 would pass through the socket to the motherboard, such as 26. With the use of the above board interposer 14, selected ones of the connections between the CPU 10 and traces on the motherboard 12 become rerouted through the interposer to the peripheral circuits on the interposer, not shown here. The route 24 shows an example of one of the rerouted paths. As will become apparent in further discussion, this reduces the
15 connection length, the number of discontinuities, and eliminates many issues with routing the connections through the motherboard.

In some instances, some connections between the CPU and the motherboard remain untouched in their routing. For example, connection path 28 between the connector 18 and the motherboard 12 will pass through. On the other hand, the
20 connection between connector 22 and the motherboard has been rerouted to pass through the interposer 14, and will not make connection to the motherboard.

Figure 2 shows a plan view of the same configuration. Motherboard 12 has a connector 16 that may consist of a LGA, BGA, etc., forming an array of connections to which the CPU 10 will connect to the motherboard 12. The above board interposer
25 14 lies between the CPU 10 and the connector 16. The above board interposer 14

includes peripheral circuits 30. Peripheral circuits may comprise memory, PCIe circuits, solid state disk drives, etc. Figures 3-5 show a side view of an embodiment of the above board interposer.

Figure 3 shows the motherboard 12 upon which resides the connector 16. In this embodiment, the peripheral circuits 30 comprise at least one memory circuit. In the prior art, the memory utilized by the CPU 10 would typically reside somewhere on the motherboard 12, such as in the dual inline memory module (DIMM) connectors 40. By moving the memory onto the above board interposer 14, one can significantly reduce access times, signal path length, discontinuities, capacitive and resistive effects, as well as increasing the signal integrity.

In Figure 4, the CPU has been mounted into the connector 16 with the interposer 14 in place. In Figure 5, one can see that the interposer is flexible, allowing the portion or portions of the interposer to fold up perpendicular to the connector. This allows the interposer to be used with any socketed or otherwise insertable processor without impacting the current architecture on the motherboard.

Figures 6A and 6B show an alternative packaging arrangement in which there is a LGA/spring connection 16 between the motherboard 12 and the interposer 14, as well as the connector 46 between the interposer and the CPU package 10. Figure 6A shows the CPU, the interposer, the sockets 16 and 46 as separate components. Figure 6B shows the components assembled together. The resulting structure has pass through connections such as 28, "T" connections such as 27, and rerouted connections such as 24. In this embodiment, the peripheral circuits 30 reside on one side of the interposer.

Figures 7A and 7B show a similar arrangement as Figures 6A and 6B, except that the interposer has connectors 50. The peripheral circuits 30 in this embodiment

are attached to auxiliary interposer substrates such as 44. The auxiliary interposer substrates connect to the above board interposer 14 connects through connectors 50. These connectors such as 50 may be populated or not, as desired. This configuration allows highly flexible arrangements of peripheral circuits in close proximity with the CPU. Further, the auxiliary interposers may also be flexible to allow them to be bent to accommodate a particular architecture, or may be pre-bent or flexed to allow their arrangement.

Figure 8 shows a side view of the interposer in contact with spring contacts. Again the connection can consist of spring contacts such as in an LGA connector, or other types of connectors, as mentioned previously. A possible point of confusion may result from the above board interposer being confused with what is typically referred to as a test interposer. Test interposers are typically employed to allow a device to be tested without damaging the device, and typically do not remain in place during normal operation; most importantly, however, test interposers are not required for proper operation of the device, in contrast to the above board interposer which is a key, required, functional element of the system.

Test interposers typically have one of two types of connections, a T-type connection or an 'out-and-back' connection. In a T-connection, a pass-through connection may have a further branch off of the side of it, similar to the letter T. An out-and-back connection has a connection that goes out from the connector on the CPU and then back to the same connector. It should be noted, that in contrast to the above board interposer, the connection either passes through or gets rerouted through the interposer, typically not both, although the T connections could be used as well. Figure 8 shows this in comparison.

At least three layers exist in the interposer, a layer of nonconductive material 64, such as polyimide or FR-4, a layer of metal used to make connections in the above board interposer, and a solder mask 66 used to insulate the metal layer 68 from unwanted contact with other devices or features, and define access points to the metal layer 68 at locations at which connections to other devices or features will be formed. For pass through contacts, such as 28, none of these three layers are involved. For a connection that mimics a test interposer T connection, such as 68, the metal layer 68 is involved. For a rerouted connection such as that from contact pad 62, the polymer or other nonconductive material 64 ensures that the connection from pad 62 does not reach the motherboard. Instead, the metal layer 68 will route that connection through the interposer to the peripheral circuits, not shown.

Figures 9 and 10 show a side view of another embodiment of an interposer substrate, using elastomer connections such as TYCO MPI (metalized particle interconnect) conductive elastomer bumps. The use of flexible connectors or contact media in the arrays of connectors may assist with mechanical conformity and connection integrity.

Using the interposer, with whichever kinds of connectors or contacts are used, reduces many of the issues with memory connectors on motherboards, such as DIMMs. It must be noted that while the discussions may use DIMMs as an example of memory connectors, the same reasoning applies to any on-motherboard memory or peripheral circuit connectors. One such issue with on-motherboard connectors such as DIMMs is the fanout length of the signal paths.

The length of the signal path has a huge effect on the signal integrity of signals traveling the path. Resistive and capacitive effects, as well as discontinuities, reduce the power of the signal and can cause errors when received if the signal does not

‘make it’ to the other end with the correct voltage. Figure 11 shows a prior art fanout length of DIMM connectors on a motherboard, a single-line, multi-drop, fanout with four memory modules.

The CPU 10 connects to a line 70 which in turn connects to a line 72. Line 72
5 traverses the four DIMM connectors, each of which represent one drop off of the line. The fanout length for four DIMMs is approximately 52 inches, for six DIMMs it is approximately 80 inches. In contrast, as shown in Figure 12, the fanout length for the above board interposer to a memory module example is between 1.5 to 2.5 inches from the CPU 10 through the on-board connector 70 through the paths 82a-c or 84a-c.
10 These segments will be discussed further with regard to Figures 14 and 15.

Further, the above board interposer has a path length that is within a quarter wavelength of the signals that will traverse the path. The advantages of being within a quarter wavelength include no need for transmission lines, which saves power; a simpler signal path; lower capacitive load requiring smaller drive signals, saving
15 power; lower read/write latency; simpler wire layout; lower part counts without the need for terminating resistors, capacitors, etc.; smaller printed circuit board (PCB) footprint with fewer layers, making it less expensive; and enhanced reliability with fewer vias and solder joints.

Typically, to determine the $\frac{1}{4}$ wavelength, one divides 300 by the clock
20 frequency of the signal, producing the wavelength of the signal in meters. For a clock frequency of 1.066 GHz, one would divide 300 by 1066 MHz, resulting in .281 meters, or 281 millimeters, or about 11 inches. One quarter of 11 inches is approximately 2.73 inches or 70 millimeters. Therefore, using the above board interposer provides the above advantages.

Further, the use of the interposer eliminates many of the signal discontinuities that exist in current architectures where the processor interacts with peripheral circuits through the motherboard. Figure 13 shows a prior art example of a signal path having multiple different discontinuities. The path provides the connection between the CPU
5 10 and one of the DIMM connectors such as 99.

The path of Figure 13 includes multiple discontinuities of several different types. The CPU connects to the motherboard through the spring 86 introducing a first, inductive, discontinuity, followed by inductive discontinuities 93a-d at the memory module sockets, among others. Capacitive discontinuities include through-
10 hole via stubs such as 88, 92a-e and 95a-d. Potential wire impedance mismatches occur throughout the signal path, at locations such as 90a-c and 91a-d, within the circuitry of the memory modules, and at resistor 94 on the motherboard.

Once the path reaches the connectors for the modules, further through-hole via stubs 95a-d exist in the motherboard, on the way to additional inductive
15 discontinuities due to the contact metal in the plug-in connectors 93a-d. In the modules themselves, capacitive discontinuities occur at the edge connector contact pads 96a-d, and further through-hole via stubs 97a-d and 98a-d. All of these discontinuities, whether capacitive, inductive, or resistive, result in impedance mismatches and have negative effects on the signal integrity. In contrast, referring
20 back to Figure 12, there are very few signal discontinuities in the signal paths of the above board interposer substrate.

Time Domain Reflectometry ("TDR") is a standard measurement method used to characterize the quality of a high speed signal path, which method provides insight into both the magnitude and physical location of discontinuities in a signal path. The
25 method operates by sending a fast-rising step waveform down a signal path, and

displaying the resulting reflections on an oscilloscope screen. The time (X-) axis relates to the physical location along the path, and the amplitude (Y-) axis displays the impedance at each point along the line, measured in ohms.

The image of Figure 14 shows the TDR response of the above board interposer
5 signal line of Figure 12, viewed looking into the line from the end that would contact the CPU pin. In a typical measurement setup, the line under test would be connected to the TDR instrument using a length of 50 ohm transmission line 101 which measures as 50 ohms. Ideally, the entire signal path from CPU to memory IC would be 50 ohms, under which conditions there would be no signal degradation due to
10 discontinuities.

When the signal encounters the line 83 between the CPU pin and the point 70 at which it begins to fan out to the memory elements, one sees that the typically very narrow line 83 needed to escape from the dense array of CPU pins exhibits an impedance of 80 ohms. When the signal reaches junction point 70, it encounters the
15 beginning of 84b, an 80 ohm line to the left, and 82b, an 80 ohm line to the right; together the two present an impedance of $80/2=40$ ohms, as shown in Fig. 14. Very shortly after passing point 70, lines 82b and 84b branch again into lines 82a and 84a, which cause the impedance to drop to $80/3=27$ ohms.

The 40 ohm impedance level lasts only as long as it takes the signal to travel to
20 the ends of stubs 82a and 84a, where it encounters a further drop to 20 ohms as a result of the signal encountering the pin capacitance of the memory ICs located at the ends of stubs 82a and 84a. After the pulse edge passes the junction points of 82a-82b, and 84a-84b, the impedance returns to the 82b, 84b level of 40 ohms until it reaches the junctions with paths 82c and 84c, where a process holds sway that is similar to
25 that at the junctions with 82a and 84a. The time delay between the arrival of the

signal at the first memory IC, and the ends of stubs 82a and 84a, and its arrival at the memory ICs at the ends of stubs 82c and 84c, ultimately plays a role in determining the maximum operating speed of the system: Since all devices must see stable “1” or “0” signal levels, the system has to wait at least as long as the approximately 100ps
5 required for the signal to traverse the distance between the closest and farthest memory ICs.

Figure 15 displays the TDR signature of the signal path of Figure 13, namely that of a conventional motherboard-memory module. In like fashion to Figure 14, the TDR signal is launched from a 50 ohm coax, 103. Unlike the case with the above
10 board interposer, the signal first encounters the spring of the CPU socket, 88, and a via in the motherboard, 86, which introduce inductive and capacitive discontinuities, respectively. The signal then traverses a length of narrow line, 85, similar to 83 in Fig. 12 and used for the same purpose, namely to escape from the dense array of pins beneath the CPU socket. Because line 85 is narrow, its impedance is well above 50
15 ohms, shown as >90 ohms in Fig. 15.

At the end of segment 85 the signal runs into a series of short lengths of line 90a-c of varying width and dielectric environments due to their locations within the circuit substrate stack up, which lines change between layers of the PCB through vias 92a-c. The varying widths of the lines, their dielectric environments, and the
20 capacitive vias all lead to impedance discontinuities that are less than 50 ohms.

Along the way to the memory modules, the signal passes through terminating resistor 94, which is located on the top or bottom side of the circuit substrate and is connected to the buried signal paths through vias 92d-e. The resistor 94 in this case presents a resistance of 34 ohms, while the vias exhibit their typical capacitive
25 discontinuities 92d-e in Fig. 15.

The signal traverses a short length of line 91a, at which point it reaches the location of the first memory module 87a and its associated socket 93a, where one sees a discontinuity due to via 95a an inductive discontinuity due to connector spring contact 96a, and a capacitive discontinuity at the metal pad on the memory module
5 corresponding to the contact spring in the connector. Once on the memory module, the signal path traverses additional vias 97a and 98a, and various short lengths of line on the way to the memory ICs mounted on the module.

It is important to note that the above description covers only the signal path to and on the first of several memory modules; each succeeding module in the chain
10 adds its own set of similar reflections, as seen in subsequent groups of similar reflections in Figure 15. Ultimately, the signal arrives at terminating resistor 89, which may present very low impedance.

As in the case of Figure 14, the time delay between closest and farthest memory ICs plays a role in setting the maximum operating speed of the system. In the
15 case of the conventional CPU-to-motherboard configuration, the delay amounts to an estimated 800ps – eight times that for the above board interposer-based system of Figures 12 and 14.

It is important to note that the *only* reason we are able to see the discontinuities in the signal path is that each discontinuity reflects energy backward to the source,
20 away from the load. All energy reflected backward represents signal information that does not reach the load and, hence, cannot contribute to information transfer. Making up for losses in the signal path, whether from dissipative resistive effects, series wiring parasitics, or impedance discontinuities, requires the expenditure of extra power to increase the signal amplitude to make up for losses.

In addition to the advantages of path lengths under $\frac{1}{4}$ wavelength and fewer signal discontinuities, it is desirable that the signal paths to the components be equal for signaling and clocking coordination. Figures 16-19 show examples of how the above board interposer can be configured to ensure equal path lengths. Figure 16 shows the interposer 14 connected to a peripheral circuit 100 that has two sets of connections 102 and 104. By configuring the interposer as shown, the components have equal path lengths to the CPU.

Several variations on this configuration exist. Figure 17 shows a configuration in which the interposer is folded to provide equal path lengths to the connections 102/104 and 106/108. Figure 18 shows a multi-layer interposer resulting from the folded configuration of Figure 17, which is then joined. Figure 19 shows a similar folded and joined configuration for the peripheral circuits, which are then stacked. The flexibility and capacity adaptation of these configurations is a further advantage of the flexible, above board interposer.

This discussion this far has focused on an above board interposer used on current motherboards for computing devices. However, if one were to assume the presence of the interposer at the design point of the motherboard, one could eliminate the fanout area on the motherboard such as that area shown as 110 in Figure 20. The board 112 in this example is a single socket server. Figure 21 shows a resulting motherboard 114 with the DRAM memory wiring and connectors eliminated. The motherboards designed with the assumption that the interposer will be used will be referred to here as 'streamlined' motherboards. Simulations have shown that this reduces the board size by 28%. The below table shows the projected board reduction sizes for different motherboard architectures based upon the number of processor

sockets. Note that some architectures have ranges depending upon the layout of the board.

Architecture	Board area reduction (%)
Single socket server	28
Double socket server	33-40
Quadruple socket server	66

In addition to assuming a retrofit application, the discussion up to this point has not addressed the possible configuration of the peripheral circuits on the interposer, except in a general sense. As mentioned above, the peripheral circuits on the interposer may consist of memory circuits such as system memory, cache, etc. Figure 22 shows an embodiment where the peripheral circuits can function as ultra-fast, off chip, meaning off the CPU, DRAM cache.

A typical memory architecture has some predetermined number of 'lanes' for memory access. For example, an Intel® X58 motherboard has 3 memory lanes and 6 memory DIMMs. In Figure 22, the peripheral circuits 116 on the interposer consist of 64 to 128 GB cache memory, with further access via the pass-through connections mentioned before to a 64 to 128 GB main memory 118. Obviously, many other memory capacities could be envisioned, this is merely one example.

Another advantage gained by elimination of the DIMMs in the streamlined motherboards comes from the reduction of the number of components that can fail under thermal cycling. Heat management has become a critical issue in motherboard design and operation as processing power has increased. The number of components reduces the number of possible points of failure. By eliminating the DIMM connectors, or other types of peripheral circuit connectors, the number of through-hole vias, solder joints, wiring runs, edge contacts, connectors, DIMM PCBs, resistors and

capacitors removed from the board number in the tens of thousands. By employing the above board interposer, these possible points of failure vanish.

In addition to reducing the number of components that may fail during thermal cycling, the above board interposer architecture may employ short length memory
5 thermal paths to the CPU heat sink. Figure 23 shows an example of this. In Figure 23, the heat from the peripheral circuit 120 on the interposer 14 is delivered via a PGS graphite sheet 121 to an existing heat sink 122. In this embodiment, the heat sink 122 is coupled to a cooler 124 via a heat pipe 123

Figure 24 shows a direct contact embodiment in which the flexibility of the
10 interposer 14 is used to put the peripheral circuit chip 120 into direct contact with a heat sink 128. In this embodiment, the heat sink consists of a water or other liquid cooled block.

Figure 25 shows another embodiment in which the interposer 14 holds the peripheral circuits such as 120. The peripheral circuits in turn are in contact with the
15 PGS graphite sheet 121 by the surface of the circuits away from the interposer. The PGS graphite sheet 121 then forms a thermal/mechanical interface with the heat sink 128, possibly including a copper mounting block 130. Materials other than graphite may be used in the sheet 121, including various metals, polymers, etc.

Many other forms of heat sinks exist than the water cooled block or the heat
20 pipe to a remote cooler block discussed this far. Figure 26 shows a similar embodiment to Figure 25, except that this embodiment employs an air-cooled heat sink block 132 with heat pipes such as 134 thermally coupled to cooling fins 134. The flexibility of the interposer architecture allows for several different possibilities for thermal management.

Another variation mentioned but not discussed in any detail addresses the number of locations of the peripheral circuits. Returning to Figure 2, the interposer 14 shown has an array of peripheral circuits on one side of the array of connectors for the CPU. No limitation of this type of configuration should be implied and none is intended. For example, the peripheral circuits may be on one side, two sides, three sides or all four sides of the array of connector for the CPU. Further, the array of peripheral circuits may have more than one bank of circuits. The interposer may make multiple folds when inserted into the motherboard socket, allowing for multiple banks of peripheral circuits. Similarly, multiple banks may be achieved using the connectors on the interposer shown in Figures 7A and 7B.

In addition, the interposer may connect in a coplanar fashion with another circuit. Figure 27 shows an embodiment in which the above board interposer 14 couples to another circuit 140. The connection may be made in several ways. This particular embodiment uses PCB edge connector pads 146 on the interposer 14 and the target circuit 140. These contacts are covered with a 'Zebra strip' 144. A Zebra strip normally consists of a rubber or plastic strip with carbon, or other conductive material, bands that connect between the edge connector pads 146 on each circuit. Zebra strips typically mount to the boards with an adhesive and pressure plates such as 142. This type of mounting can be very tolerant of circuit substrate thickness mismatches, mechanical misalignments, gaps, etc.

As mentioned above, other variations include the different type of connectors used in connecting the CPU to the interposer and between the interposer and the motherboard. One such variation includes pin grid arrays. Figure 28 shows an example of connector pins such as 150 inserting into a center opening of the copper capture pad in the socket on the motherboard. Openings in the interposer 152 allow

the pins to penetrate the interposer, while having traces in the interposer to connect to the pins.

The advantages of the above board interposer are varied and numerous, as discussed above. Although there has been described to this point a particular
5 embodiment for an above board interposer, it is not intended that such specific references be considered as limitations upon the scope of this invention except in-so-far as set forth in the following claims.

WHAT IS CLAIMED IS:

1. An interposer substrate, comprising:

an array of interconnects in the interposer substrate, the array of connectors arranged in accordance with an array of interconnects for a processor on a circuit

5 substrate;

at least one conductive trace in the interposer substrate in connection with at least one connector in the array of interconnects, the conductive trace arranged parallel to the interposer substrate such that no electrical connection exists between the connector in the interposer substrate and a corresponding one of the interconnects for

10 the processor on the circuit substrate; and

at least one peripheral circuit residing on the interposer substrate in electrical connection with the conductive trace.

2. The interposer substrate of claim 1, wherein the interposer substrate comprises a flexible substrate.

15 3. The interposer substrate of claim 1, wherein the peripheral circuit comprises at least one memory chip.

4. The interposer substrate of claim 1, wherein the peripheral circuit further comprises at least one Peripheral Component Interconnect Express (PCIe) unit.

5. The interposer substrate of claim 1, wherein the peripheral circuit comprises a
20 solid state disk drive.

6. The interposer substrate of claim 1, wherein the peripheral circuit is arranged on an edge of the interposer such that the interposer can fold when inserted between the processor and the circuit substrate.

7. The interposer substrate of claim 1, wherein the peripheral circuit attaches to
25 the interposer by a connector.

8. The interposer substrate of claim 7, wherein the connector further comprises a memory micromodule connector.

9. The interposer substrate of claim 7, wherein the connector further comprises a coplanar connector.

5 10. The interposer substrate of claim 1, wherein the interposer substrate has multiple folds, each fold forming a plane of the interposer, each plane of the interposer holding at least one peripheral circuit.

11. The interposer substrate of claim 10, wherein the path length to each peripheral circuit is the same.

10 12. The interposer substrate of claim 1, wherein the array of interconnects on the interposer substrate comprise one of a land grid array, ball grid array, pin grid array or elastomer contacts.

13. The interposer substrate of claim 1, wherein the array of interconnects for the processor comprise one of a land grid array, ball grid array, pin grid array or a socket.

15 14. The interposer substrate of claim 1, wherein the conductive trace has a length of less than or equal to one quarter a wavelength of a highest frequency signal transmitted over the trace.

15. The interposer substrate of claim 1, wherein the interposer further comprises at least one connector of the array of interconnects in the interposer substrate that
20 connects to one of the interconnects in the array of interconnects for the processor on the circuit substrate.

16. The interposer substrate of claim 15, wherein the one connector connects to memory circuits on the circuit substrate and the peripheral circuit on the interposer is configured as cache memory.

25 17. An apparatus comprising:

a processor having an array of interconnects arranged to connect the processor to conductive paths;

a circuit substrate having an array of interconnects arranged to provide connections between a processor and the circuit substrate, the circuit substrate having
5 conductive paths connected to the array of interconnects;

an interposer substrate arranged between the processor and the circuit substrate, at least one conductive trace in the interposer substrate in connection with at least one connector in the array of interconnects on the processor, the conductive trace arranged parallel to the interposer substrate such that no electrical connection exists
10 between the connector in the interposer substrate and a corresponding one of the interconnects on the processor on the circuit substrate; and

at least one peripheral circuit connected to the at least one conductive trace.

18. The apparatus of claim 17, the apparatus further comprising a heat sink thermally coupled to the processor and the peripheral circuit.

15 19. The apparatus of claim 18, wherein the heat sink thermally coupled to the processor comprises one of either the heat sink directly contacting the processor, or a thermal transfer sheet between the heat sink and the processor

20. The apparatus of claim 17, wherein the interposer further comprises at least one connector of the array of interconnects in the interposer substrate that connects to
20 one of the interconnects in the array of interconnects for the processor on the circuit substrate.

21. The apparatus of claim 17, wherein the circuit substrate comprises one of a motherboard having an array of memory interconnects, or a motherboard without any memory interconnects.