


**Exhibit C-1 to Defendants' Invalidity Contentions:
Comparison of U.S. Patent No. 7,594,052 ("’052 patent") and U.S. Patent Publication No. 2002/0129173 ("Weber")**

As described in the following claim chart, claim 6 of U.S. Patent No. 7,594,052 ("’052 patent") is invalid under 35 U.S.C. §§ 102 and/or 103 in view of U.S. Patent Publication No. 2002/0129173 ("Weber"). The citations below are to Weber and compares Weber to the ’052 patent claims.¹


To the extent Plaintiff alleges that Weber does not disclose any particular limitation of claim 6 (the "Asserted Claim") of the ’052 patent, either expressly or inherently, it would have been obvious to a person of ordinary skill in the art as of the priority date of the ’052 patent to modify Weber and/or to combine the teachings of Weber with other prior art references, including but not limited to the prior art references presently found in Defendants' Invalidity Contentions and any supplements thereto and the relevant section of charts for other prior art for the ’052 patent in a manner that would have rendered the Asserted Claim invalid as obvious. The inclusion of the particular limitation into the methods and systems described in Weber would have been obvious for at least the following reasons: it would have been obvious to modify the methods and/or systems disclosed in Weber, alone or in combination with the references identified in Defendants' Invalidity Contentions, to meet this limitation based on: (a) the teachings in Weber itself; (b) a person of ordinary skill in the art implementing a predictable variation of the methods and/or systems disclosed in Weber or in the other identified references; (c) the logic, judgment, intelligence, creativity, innovation and common sense a person of ordinary skill in the art would employ; (d) the inferences and creative and routine steps that a person of ordinary skill in the art would employ; (e) market demand for methods and/or systems to include this limitation; and/or (f) design needs or market pressures to solve a problem with a finite number of solutions that could have produced the desired outcome.

¹ The citations are to the issued patent version of Weber, and they also disclose the corresponding citations to the patent application publication version of Weber, which are omitted for brevity. Nothing in these claim charts should be construed as an admission regarding infringement, either literally or under the doctrine of equivalents, or as an admission regarding Defendants' understanding of the proper scope of the Asserted Claim. Given Plaintiff's vague and ambiguous infringement contentions, the exemplary citations herein necessarily account for a variety of possible infringement arguments and claim constructions. Defendants reserve the right to rely on additional citations or sources of evidence that also may be applicable, or that may become applicable in light of claim construction, changes in Plaintiff's infringement contentions, and/or information obtained during discovery as the case progresses. Various positions put forth in this document are predicated on Plaintiff's overly broad interpretation of its claims as evidenced by the infringement allegations received to date. Those positions are not intended to and do not necessarily reflect Defendants' interpretation of the true and proper scope of Plaintiff's claims, and Defendants reserve the right to adopt claim construction positions that differ from or even conflict with various positions put forth in this document.


A. Claim 6

'052 Patent Claim Limitation	Weber
[6pre-a] Method of communication service mapping in an integrated circuit, having a plurality of processing modules (M, S),	To the extent the preamble is limiting, Weber discloses and/or renders obvious this part of the preamble based on at least the following citations. See, e.g.: • Weber at FIG. 1  FIG. 1 • Weber at para. [0037] FIG. 1 is a block diagram of a complex electronics system 100. Shared communications bus 112 connects sub-systems 102, 104, 106, 108, and 110. Subsystems are typically functional blocks including a interface module for interfacing to a shared bus.

'052 Patent Claim Limitation	Weber
	• Weber at para. [0038] System 100 is an example of a system that may consist of one or more integrated circuits or chips. • Weber at para. [0127] In the case of an initiator functional block, a one-to-one static correspondence may exist between Thread ID and CONNID. For example if the Thread ID is “1”, a single CONNID is mapped for a particular interface module, solving the problem of multiple, identical functional blocks. Furthermore, to the extent the preamble is not disclosed in Weber, the preamble (and the claim as a whole) would have been obvious based on the combinations identified in Defendants’ Invalidity Contentions.
[6pre-b] wherein at least one first of said processing modules (M) requests at least one communication service to at least one second processing module (S)	To the extent the preamble is limiting, Weber discloses and/or renders obvious this part of the preamble based on at least the following citations. See, e.g.: • Weber at para. [0048] An initiator interface module issues and receives read and write requests to and from functional blocks other than the one with which it is associated. In one embodiment, an initiator interface module is typically connected to a CPU, a digital signal processing (DSP) core, or a direct memory access (DMA) engine. • Weber at FIG. 4


'052 Patent Claim Limitation	Weber
	 FIG. 4 • Weber at para. [0048] FIG. 4 is a block diagram of a part of a computer system 1000 according to one embodiment. FIG. 4 is useful in illustrating multilevel connection identification. System 1000 includes initiator functional block 1002, which is connected to initiator interface module 1004 by interconnect 1010. Initiator interface module 1004 is connected to target interface module 1006 by shared communications bus 1012. Target interface module 1006 is connected to target functional block 1008 by an interconnect 1010. Typically, shared communications bus 1012 is analogous to shared communications bus 112 of FIG. 1 or to shared communications bus 114 of FIG. 1. Interconnects 1010 are typically analogous to interconnect 115 of FIG. 1 in that they connect functional blocks to interface modules and are point-to-point, rather than shared, interconnects. Interconnects 1010 are typically physically shorter than shared communications


'052 Patent Claim Limitation	Weber
	bus 1012 because of their local nature. As will be explained more fully below, system 1000 uses two different levels of connection identification depending upon the requirements of a particular functional block. “Global” connection identification information is sent on shared communications bus 1012, while “local” connection information, or thread identification information, is sent in interconnects 1010. • Weber at FIG. 11

'052 Patent Claim Limitation	Weber
	 FIG. 11 Weber at para. [0074] FIG. 11 shows an interconnect 1010, which is a point-to-point interconnect as shown in FIG. 4. Interconnect 1010 includes additional signals as compared to the protocol described with


'052 Patent Claim Limitation	Weber
	reference to FIG. 5. As will be explained below, some of the additional signals are particularly useful as signals sent over point-to-point interconnects such as interconnects 1010. The protocol of interconnect 1010 controls point-to-point transfers between a master entity 1102 and a slave entity 1104 over a dedicated (non-shared) interconnect. Referring to FIG. 5, a master entity may be, for example, initiator functional block 1002 or target interface module 1006. A slave entity may be, for example, initiator interface module 1004 or target functional block 1008. Furthermore, to the extent the preamble is not disclosed in Weber, the preamble (and the claim as a whole) would have been obvious based on the combinations identified in Defendants' Invalidity Contentions.
[6pre-c] based on specific communication properties and at least one communication service identification, wherein said at least one communication service identification comprises at least one communication thread or at least one address range, said address range for identifying one or more second processing modules (S) or a memory region within said one or more second processing modules (S), comprising the steps of:	To the extent the preamble is limiting, Weber discloses and/or renders obvious this part of the preamble based on at least the following citations. See, e.g.: • Weber at para. [0012] One embodiment also includes a bus configured to carry multiple signals, wherein the signals include a connection identifier signal that indicates a particular connection that a data transfer between an initiator functional block and a target functional block is part of. The connection identifier includes information about the connection, such as which functional block is the source of a transmission, a priority of a transfer request, and transfer ordering information. One embodiment also includes a thread identifier, which provides a subset of the information provided by the connection identifier. In one embodiment, the thread identifier is an identifier of local scope that identifies transfers between an interface module and a connected functional block, where in some embodiments, an interface module connects a functional block to a shared communications bus. • Weber at para. [0085] Request Thread Identifier (ReqThreadID), in one embodiment, is a four-bit signal that provides the thread number associated with a current transaction intended for slave 1104. All commands executed with a particular thread ID must execute in order with respect to one another, but they may execute out of order with respect to commands from other threads. Response Thread Identifier (RespThreadID) provides a thread number associated with a current response. Because responses in a thread may return out of order with respect to other threads, RespThreadID is necessary to identify

'052 Patent Claim Limitation	Weber
	which thread's command is being responded to. In one embodiment, ReqThreadID and RespThreadID are optional signals, but if one is used, both should be used. • Weber at para. [0089] The ReqThreadBusy and RespThreadBusy signals described above are an example of a per-thread flow control mechanism. Other per-thread flow control mechanisms exist. Alternately a credit-based mechanism may be used. In a credit-based mechanism, the slave informs the master how many or approximately how many transfers it can accept. In one embodiment, this information is communicated by the slave raising a credit signal for a single cycle for each transfer it can accept. Thus, the number of “credits” or transfers the slave can accept is communicated by the number of cycles that the signal is active. • Weber at FIG. 1

'052 Patent Claim Limitation	Weber
	 The diagram illustrates a system 100. A central horizontal bus 112 connects several components. On the left, a block 102 contains sub-blocks A, B, and C, each connected to a control unit (120, 121, 122) via lines 115. These control units are further connected to a block 114, which then connects to sub-block D. Below block 102 is an ASIC block containing sub-blocks A, B, and C, connected to a control unit 123 via line 115, which in turn connects to sub-block E. To the right of the ASIC is a block 104 containing sub-blocks A, B, and C, connected to a control unit 125 via line 115. On the far right, a block 110 (FPGA) contains sub-blocks A and B, connected to a control unit 128 via line 114. Above the FPGA are two memory blocks: DRAM 106 (connected via line 126) and EEPROM 108 (connected via line 127). A block 124 is also connected to the bus 112 via line 115. The entire system is labeled 100 at the bottom right. FIG. 1 • Weber at para. [0037] Sub-system 106 is a dynamic random access memory (DRAM) • Weber at FIG. 11


'052 Patent Claim Limitation	Weber
	 The diagram illustrates a point-to-point interconnect 1010 between a MASTER 1102 and a SLAVE 1104. The signals are as follows: - Clock (bidirectional) - Cmd [3] (MASTER to SLAVE) - Addr [N] (MASTER to SLAVE) - DataOut [N] (MASTER to SLAVE) - ReqAccept (SLAVE to MASTER) - Resp [3] (SLAVE to MASTER) - DataIn [N] (SLAVE to MASTER) - RespAccept (MASTER to SLAVE) - Width [3] (I) (MASTER to SLAVE) - Burst [4] (I) (MASTER to SLAVE) - InterruptOut (I) (MASTER to SLAVE) - ErrorOut (I) (MASTER to SLAVE) - InterruptIn (I) (SLAVE to MASTER) - ErrorIn (I) (SLAVE to MASTER) - Reset (I) (SLAVE to MASTER) - ReqThreadID [4] (A) (MASTER to SLAVE) - RespThreadID [4] (A) (SLAVE to MASTER) - ReqThreadBusy [16] (AI) (MASTER to SLAVE) - RespThreadBusy [16] (AI) (SLAVE to MASTER) - ReqConnID [8] (AI) (MASTER to SLAVE) FIG. 11 - Weber at para. [0074] FIG. 11 shows an interconnect 1010, which is a point-to-point interconnect as shown in FIG. 4. Interconnect 1010 includes additional signals as compared to the protocol described with

'052 Patent Claim Limitation	Weber
	reference to FIG. 5. As will be explained below, some of the additional signals are particularly useful as signals sent over point-to-point interconnects such as interconnects 1010. The protocol of interconnect 1010 controls point-to-point transfers between a master entity 1102 and a slave entity 1104 over a dedicated (non-shared) interconnect. Referring to FIG. 5, a master entity may be, for example, initiator functional block 1002 or target interface module 1006. A slave entity may be, for example, initiator interface module 1004 or target functional block 1008. Furthermore, to the extent the preamble is not disclosed in Weber, the preamble (and the claim as a whole) would have been obvious based on the combinations identified in Defendants' Invalidity Contentions.
[6a] coupling said plurality of processing modules (M, S) by an interconnect means (N) and	Weber discloses and/or renders obvious this limitation based on at least the following citations. See, e.g.: • Weber at Fig. 1

'052 Patent Claim Limitation	Weber
	 FIG. 1 • Weber at para. [0037] FIG. 1 is a block diagram of a complex electronics system 100. Shared communications bus 112 connects sub-systems 102, 104, 106, 108, and 110. Subsystems are typically functional blocks including a interface module for interfacing to a shared bus. Furthermore, to the extent this limitation is not disclosed in Weber, this limitation (and the claim as a whole) would have been obvious based on the combinations identified in Defendants' Invalidity Contentions.
[6b] enabling a connection based communication having	Weber discloses and/or renders obvious this limitation based on at least the following citations. See, e.g.:

'052 Patent Claim Limitation	Weber
a set of connection properties,	• Weber at paras. [0053-0054] Connection identifier (CONNID) lines 226 carry a multi-bit signal driven by an initiator bus owner to indicate which connection the current transfer is part of. A connection is a logical state, established by an initiator, in which data may pass between the initiator and an associated target. The CONNID typically transmits information including the identity of the functional block initiating the transfer and ordering information regarding an order in which the transfer must be processed. In one embodiment, the information conveyed by the CONNID includes information regarding the priority of the transfer with respect to other transfers. In one embodiment the CONNID is a eight-bit code. An initiator interface module sends a unique CONNID along with an initial address transfer of a connection. Later transfers associated with this connection (for example, data transfers) also provide the CONNID value so both sender and receiver (as well as any device monitoring transfers on shared communications bus 1012) can unambiguously identify transfers associated with the connection. One advantage of using a CONNID is that transfers belonging to different transactions can be interleaved arbitrarily between multiple devices on a per cycle basis. In one embodiment, shared communications bus 1012 implements a fully pipelined protocol that requires strict control over transaction ordering in order to guarantee proper system operation. Without the use of a CONNID, ordering constraints within a particular transaction may be violated because transfers associated with a particular connection are not identified. Because a first command may be rejected by a BUSY response while a later command is already in flight, it is essential to provide mechanisms that allow full control over which commands complete. If such control is not present, ambiguous system behavior can result. For instance, if a single initiator interface module issues a sequence of dependent read and write commands, a busy response to one of the commands could result in later commands returning the wrong data. One solution to such problems is to avoid overlapping dependent commands. This solution, however, increases the latency of every dependent command in order to ensure proper results. The present invention uses a CONNID signal, in part, to allow overlapping of dependent commands. Therefore, use of a CONNID improves system performance and efficiency. Another benefit of the CONNID of the present invention is that communication system predictability is enhanced because it allows a shared functional block to respond to requests based upon quality of service guarantees that may vary

'052 Patent Claim Limitation	Weber
	between connections. For example, data requested to operate a computer display cannot tolerate unpredictable delay because delay causes the display to flicker. Therefore, the CONNID may be used to prioritize data requests from a display controller so that requests from the display controller to a common resource are serviced before other requests. The present invention also allows for flexible reconfiguration of the CONNID to retune system performance. • Weber at paras. [0125] Using a networking analogy, the Thread ID is a level-2 (data link layer) concept, whereas the CONNID is more like a level-3 (transport/session layer) concept. Some functional blocks only operate at level-2, so it is undesirable to burden the functional block or its interface module with the expense of dealing with level-3 resources. Alternatively, some functional blocks need the features of level-3 connections, so in this case it is practical to pass the CONNID through to the functional block. • Weber at paras. [0128] In the case of a target functional block, there is a one-to-one dynamic correspondence between a Thread ID and a CONNID. If a target functional block supports two simultaneous threads, the target interface module acquires the CONNID of an open connection and associates it with a thread as needed. For example, a target interface module receives a CONNID of “7”, and then maps CONNID 7 to thread “0”. Thereafter, all transfers with CONNID 7 are associated with thread 0 until connection 7 is closed. Furthermore, to the extent this limitation is not disclosed in Weber, this limitation (and the claim as a whole) would have been obvious based on the combinations identified in Defendants’ Invalidity Contentions.
[6c] controlling the communication between said at least one first of said plurality of processing modules (M) and said interconnect means (N) by at	Weber discloses and/or renders obvious this limitation based on at least the following citations. See, e.g.: • Weber at FIG. 1

'052 Patent Claim Limitation	Weber
least one network interface (NI) associated to said at least one first of said processing modules,	 The diagram illustrates a system architecture with several interconnected components. A central block 102 contains functional blocks A, B, C, D, and E, along with interface modules 120, 121, 122, 123, and 124. Block 102 is connected to a bus 112. A peripheral block 104 contains functional blocks A, B, and C, with interface module 125. A block 106 contains DRAM 106 and EEPROM 108, connected to the bus 112. A block 110 contains functional blocks A and B, with interface module 128. The bus 112 connects all these components. Reference numerals 114, 115, 126, 127, and 128 point to specific connection points or components within the system. FIG. 1 • Weber at para. [0040] Interface modules 120-127 are connected to various functional blocks as shown. In this embodiment, interface modules 120, 122, 123 and 124 are physically separated from their connected functional block (A, B, C, E and 102, respectively). Interface modules 121, and 125-128 are essentially part of their respective functional blocks or sub-systems. Some functional blocks, such as 102D, do not require a dedicated interface module. The arrangement of sub-systems, functional blocks and interface modules is flexible and is determined by the system designer.

'052 Patent Claim Limitation	Weber
	Furthermore, to the extent this limitation is not disclosed in Weber, this limitation (and the claim as a whole) would have been obvious based on the combinations identified in Defendants' Invalidity Contentions.
[6d] mapping the requested at least one communication service based on said specific communication properties to a connection based on a set of connection properties according to said at least one communication service identification.	Weber discloses and/or renders obvious this limitation based on at least the following citations. See, e.g.: • Weber at paras. [0053-0054] Connection identifier (CONNID) lines 226 carry a multi-bit signal driven by an initiator bus owner to indicate which connection the current transfer is part of. A connection is a logical state, established by an initiator, in which data may pass between the initiator and an associated target. The CONNID typically transmits information including the identity of the functional block initiating the transfer and ordering information regarding an order in which the transfer must be processed. In one embodiment, the information conveyed by the CONNID includes information regarding the priority of the transfer with respect to other transfers. In one embodiment the CONNID is a eight-bit code. An initiator interface module sends a unique CONNID along with an initial address transfer of a connection. Later transfers associated with this connection (for example, data transfers) also provide the CONNID value so both sender and receiver (as well as any device monitoring transfers on shared communications bus 1012) can unambiguously identify transfers associated with the connection. One advantage of using a CONNID is that transfers belonging to different transactions can be interleaved arbitrarily between multiple devices on a per cycle basis. In one embodiment, shared communications bus 1012 implements a fully pipelined protocol that requires strict control over transaction ordering in order to guarantee proper system operation. Without the use of a CONNID, ordering constraints within a particular transaction may be violated because transfers associated with a particular connection are not identified. Because a first command may be rejected by a BUSY response while a later command is already in flight, it is essential to provide mechanisms that allow full control over which commands complete. If such control is not present, ambiguous system behavior can result. For instance, if a single initiator interface module issues a sequence of dependent read and write commands, a busy response to one of the commands could result in later commands returning the wrong data. One solution to such problems is to avoid overlapping dependent commands. This solution, however, increases the

'052 Patent Claim Limitation	Weber
	latency of every dependent command in order to ensure proper results. The present invention uses a CONNID signal, in part, to allow overlapping of dependent commands. Therefore, use of a CONNID improves system performance and efficiency. Another benefit of the CONNID of the present invention is that communication system predictability is enhanced because it allows a shared functional block to respond to requests based upon quality of service guarantees that may vary between connections. For example, data requested to operate a computer display cannot tolerate unpredictable delay because delay causes the display to flicker. Therefore, the CONNID may be used to prioritize data requests from a display controller so that requests from the display controller to a common resource are serviced before other requests. The present invention also allows for flexible reconfiguration of the CONNID to retune system performance. • Weber at paras. [0127-0128] In the case of an initiator functional block, a one-to-one static correspondence may exist between Thread ID and CONNID. For example if the Thread ID is “1”, a single CONNID is mapped for a particular interface module, solving the problem of multiple, identical functional blocks. In the case of a target functional block, there is a one-to-one dynamic correspondence between a Thread ID and a CONNID. If a target functional block supports two simultaneous threads, the target interface module acquires the CONNID of an open connection and associates it with a thread as needed. For example, a target interface module receives a CONNID of “7”, and then maps CONNID 7 to thread “0”. Thereafter, all transfers with CONNID 7 are associated with thread 0 until connection 7 is closed. • Weber at para. [0130-0131] A DMA engine is an example of an initiator functional block that also functions as a target functional block. When the DMA engine is programmed by software, it acts as a target. Thereafter, the DMA engine is an initiator. Because a DMA engine performs both read and write operations, two connections can be associated with a single DMA engine. If some buffering is available in the DMA engine, read and write operations may be decoupled so that both types of operations can be performed concurrently. A read may occur from a long latency storage device which requires the read data to be buffered on the DMA engine before a write operation writes the data. In one

'052 Patent Claim Limitation	Weber
	embodiment, each of DMA engines 1202 uses a Thread ID to identify the read stream and a different Thread ID to identify the write stream. The DMA engine does not require more information, such as what other functional block participates in a transaction. Therefore, a CONNID is not required to be sent from the DMA engine 1202 to a connected interface module 1204. Mapping of a Thread ID to a CONNID occurs in the interface module 1204. In one embodiment, each initiator interface module 1204 maps a unique CONNID to each of two Thread Ids from a connected DMA engine 1202. Each of DMA engines 1202 use a single bit, for example, Thread ID of FIG. 11, to distinguish between its two threads. For each transfer over shared communications bus a unique CONNID is sent to target interface module 1206. The CONNID may include priority information, for example, assigning high priority to requests for graphics data. The high priority graphics data request is immediately serviced by DRAM controller 1208 while lower priority request may be required to wait. Furthermore, to the extent this limitation is not disclosed in Weber, this limitation (and the claim as a whole) would have been obvious based on the combinations identified in Defendants' Invalidity Contentions.