

Petitioner Qualcomm's Trial Hearing Demonstratives

Qualcomm Incorporated,
Petitioner

v.

Network System Technologies, LLC
Patent Owner

IPR2024-00343 ('052 patent)

March 20, 2025



The Instituted Grounds

	Claims	Basis for Unpatentability
Ground 1	1–2, 5–7	Anticipated by Rădulescu2004
Ground 2	1–7	Obvious based on Rădulescu2004 in view of OCP
Ground 3	1–7	Anticipated by Weber
Ground 4	1–7	Obvious based on Weber in view of Rădulescu2004

Radulescu2004 is Prior Art

- Evidence of Publication Prior to March 17, 2004 is overwhelming
 1. Dutt Declaration (Ex-1012)
 - Personally attended DATE 04 Conference in February 2004 and presented papers there.
 - Testified that “[a] copy of the proceedings for DATE 04 was provided to conference participants upon registrations at the start of the conference”
 2. MacPherson Declaration (Ex-1018)
 - IEEE custodian confirmed DATE 04 proceedings were distributed at DATE 04 Conference in February 2004 and posted to IEEE Xplore on March 8, 2004
 3. Conventional indicia of publication
 4. Goossens Website (Ex-1062)
 - Website of co-author identifying Radulescu2004 as published at DATE in February, 2004
 - Goossens is retained by PO in the related district court litigation (Ex-1034)
 5. Second Hall-Ellis Declaration (Ex-1068)
 - Professor of Library Science attaching various Library of Congress documents, including copy of DATE 04 proceedings with Radulescu2004 and a **received stamp of February 26, 2004**

Radulescu2004 is Prior Art



Ex-1066

Slide 4

DEMONSTRATIVE EXHIBIT – NOT EVIDENCE

Radulescu2004 is Prior Art

- Radulescu2004 was available to POSAs via both DATE 04 Conference and IEEE Xplore
 1. Dutt Reply Declaration (Ex-1038)
 - DATE Conferences advertised to POSAs (¶¶ 6-7)
 - IEEE Xplore well-known and used by POSAs (¶¶ 12-16)
 2. Second Hall-Ellis Declaration (Ex-1068)
 - IEEE Xplore is a well-known public source for technical publications and proceedings (¶¶ 11-18)
 3. Goossens Website (Ex-1062)
 - Inventors submitted multiple NoC publications to DATE
 4. DATE Website (Exs-1031, -1033, -1036)
 - Discussing focus on SoC
- No contrary evidence from PO

OCP 1.0 is Prior Art

- Both the '052 Patent and Radulescu2004 identify OCP as an existing protocol

The network interfaces NI are responsible for packetization/depacketization, for implementing the connections and services, and for offering a **standard interface** (e.g., AXI or **OCP**) to the IP modules connected to the NoC. In particular, the network interfaces NI may implement adapter units A to **existing on-chip protocols, such as AXI, OCP and DTL**, to seamlessly connect existing IP modules to the NoC. These adapter units A may also be arranged between the IP blocks and their associated network interfaces.

The communication of the IP block described above may be implemented by **existing protocols** like, DTL Device Transaction Level (DTL) Protocol Specification, MTL Protocol Specification, AMBA AXI Protocol Specification, **OCP International Partnership, Open Core Protocol Specification**, or the like.

'052 Patent, col. 6:30-39, 8:41-46

*In this paper we present a network interface for an on-chip network. Our network interface decouples computation from communication by offering a shared-memory abstraction, which is independent of the network implementation. We use a transaction-based protocol to achieve backward compatibility with **existing bus protocols such as AXI, OCP and DTL**. Our network interface*


[18] OCP International Partnership. **Open Core Protocol Specification. 2.0 Release Candidate, 2003.**

Radulescu2004 (Ex-1005), pp. 1, 6

PO's Expert: "It says Release 2.0. I think **Release 2.0 came after Release 1.0.**"

Ex-1050 (Bagherzadeh Dep.), p. 161:1-2

OCP 1.0 is Prior Art



Ex-1006
Slide 7

OCP 1.0 is Prior Art

- Other Evidence of Publication Prior to March 17, 2004
 1. Conventional indicia of publication
 2. First Frank-White Declaration (Ex-1013)
 - Internet Archive custodian showing copy of OCP Specification 1.0 available on professor's website in 2003
 3. Second Frank-White Declaration (Ex-1019)
 - Showing OCP Specification 1.0 was available from OCP website in 2003 via click-through license
 4. Goossens Textbook Chapter (Ex-1009)
 - Another paper authored by inventor citing "Open core protocol specification, 2001"
 5. Dutt Reply Declaration (Ex-1038)
 - Confirming public availability of OCP Specification 1.0 prior to 2004
- No contrary evidence from PO

IPR2024-00343
('052 Patent IPR)

Ground 1: Anticipation by
Radulescu2004
(Claims 1-2 and 5-7)

Radulescu2004 (Ex-1005)

- System-on-Chip (SoC) that includes:
 - Processing Modules
 - Network
 - Network Interfaces (NIs)
 - Adapter Units in NIs
 - Shells in NIs





Figure 8. NI configuration

Radulescu2004


The Æthereal NoC offers its services on *connections*, which can be point to point (one master, one slave), *multicast* (one master, multiple slaves, all slaves executing each transaction), and *narrowcast* (one master, multiple slaves, a transaction is executed by only one slave) [23]. Connections are composed of unidirectional point-to-point *channels* (between a single master and a single slave). To each channel, *properties* are attached, such as guaranteed message delivery or not, in order or un-ordered message delivery, and with or without timing guarantees. As a result, different properties can be attached to the request and response parts of a connection, or for different slaves within the same connection.

Ex-1005, p. 1


Radulescu2004



Radulescu2004




'052 Patent



Ex-1005, FIGS. 1, 5 (annotated by PO)

In Figures 5 and 6, we show a master and slave shells that implement a simplified version of a protocol such as AXI. The basic functionality of such a shell is to sequentialize commands and their flags, addresses, and write data in request messages, and to desequentialize messages into read data, and write responses. Examples of the message structures (i.e., after sequentialization) passing from NI shells and NI kernel are shown in Figure 7. In full-fledged master and slave shells, more blocks would be added to implement e.g., the unbuffered writes at the master side, and read linked, write conditional at the slave side.

Ex-1005, p. 5




Ex-1005, FIGS. 1, 3 (annotated by PO)

In Figure 3, we show an example of a narrowcast shell. Narrowcast connections are connections between one master and several slaves, where each transaction is executed by a single slave selected based on the address provided in the transaction [23]. Narrowcast connections provide a simple, low-cost solution for a single shared address space mapped on multiple memories. It implements the splitting/merging of data going to/coming from these memories.

We implement the narrowcast connection as a collection of point-to-point connections, one for each master-slave pair. Within a narrowcast connection, the slave for which the transaction is destined is selected based on the address (Conn block). The address range assigned to a slave is configurable in the narrowcast module. To provide in-order response delivery, the narrowcast must also keep a history of connection identifiers of the transactions including responses (e.g., reads, and acknowledged writes), and the length of these responses. In-order delivery per slave of request messages is already provided by the point-to-point connections.

Ex-1005, p. 4

Radulescu2004



The NI kernel (see Figure 2) receives and provides messages, which contain the data provided by the IP modules via their protocol after sequentialization. The message structure may vary depending on the protocol used by the IP module. However, the message structure is irrelevant for the NI kernel, as it just sees messages as pieces of data to be transported over the NoC.

The NI kernel communicates with the NI shells via *ports*. At each port, point-to-point connections can be configured, their maximum number being selected at NI instantiation time. A port can have multiple connections to allow differentiated traffic classes, in which case there are also `connid` signals to select on which connection a message is supplied or consumed.

Ex-1005, p. 3

Ex-1005, FIGS. 3, 2 (annotations added)

Claim 6

Independent Claim 6	
[6.pre.1]	Method of communication service mapping in an integrated circuit, having a plurality of processing modules (M,S),
[6.pre.2]	wherein at least one first of said processing modules (M) requests at least one communication service to at least one second processing module (S) 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:
[6.a]	coupling said plurality of processing modules (M,S) by an interconnect means (N) and
[6.b]	enabling a connection based communication having a set of connection properties,
[6.c]	controlling the communication between said at least one first of said plurality of processing modules (M) and said interconnect means (N) by at least one network interface (NI) associated to said at least one first of said processing modules,
[6.d]	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.

Disputed claim language (limitations 6.pre.2 and 6.d)

“6. A method of communications service mapping...

wherein at least one first of said processing modules (M) **requests at least one communication service** to at least one second processing module (S) **based on**

[1] **specific communication properties** and [2] **at least one communication service identification,**

wherein **said at least one communication service identification comprises** [a] at least one communication thread or [b] at least one address range,

* * * *

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.**

Service request “based on” on address range in the ‘052 patent

- Dr. Hassoun’s conclusion


For example, a POSA would understand that a request with address 0004 (which is within the range 0000-7FFF assigned to slave S1) would be mapped by the network interface to connection id cid=0, and a request with address 8FFF (which is within the range 8000-BFFF²) would be mapped by the network interface to connection id cid=1. A POSA would understand that the service request sent by master M is “based on” an address range in the sense that the request contains an address that corresponds to one of the address ranges stored in the network interface that identifies the slave S for which the message is intended. Ex. 1001, 8:7-20; *see also* 8:21-40 (reproduced above).

Ex-1030 (Hassoun Reply Dec.), ¶26

Service request “based on” on address range in the ‘052 patent

156. I note that Rădulescu2004 uses the term “addr” when referencing the address. Ex. 1005 at FIGs. 5, 6. A POSITA would immediately recognize that “addr” represents an address, and not an address range. The same authors of

Ex-2070 (Bagherzadeh Dec.), ¶ 156



‘052 Patent

Service request “based on” on address range in the ‘052 patent

The invention is based on the idea to offer differentiated services for protocols such as DTL, MTL, AXI, and OCP, by mapping the identification means of these protocols to connections. The identification means in the existing protocols are: communication threads and addresses, i.e. the threads or addresses are mapped to connections through the interconnect based on specific connection properties.

‘052 Patent, col. 4:15-21

A master module M requests a communication service, i.e. a request is issued to the slave module S. This request may comprise a command as well as some data. The slave S receives the request and performs the respective processings based on the received command and returns a response possibly also with some data. The adapter units A serve to connect the IP modules and their communication protocol to the network interface which in turn serve to connect to the network N. In other words, the adapter units A translate the requested communication service into messages msg, which are processed by the network interface NI and are translated into flits fit to be transmitted over the network N, in particular the adaptor unit A maps the identification means of these protocols to connections. The identification means in the existing protocols are: communication threads and addresses, i.e. the threads or addresses are mapped to connections through the interconnect based on specific connection properties.

‘052 Patent, col. 5:53-6:3

Service request “based on” on address range in the ‘052 patent

FIG. 4 shows a basic block diagram of a network on chip for illustrating the operation of the mapping of IP block communication via address ranges to connections. The network on chip comprises a master M and two slaves S1, S2 each being connected to a network via a network interface NI, respectively. The two slaves S1, S2 are implemented as memories, wherein the first slave S1 comprises the address range 0000-7FFF and the second slave S2 comprises the address range 8000-FFFF. Accordingly, the network interface NI associated to the master M may map the communication to the first slave S1, i.e. to the address range 0000-7FFF, to a connection with the connection id cid=0 and the communication to the second slave S2, i.e. to the address range 8000-FFFF, to a connection with the connection id cid=1.

‘052 Patent, col. 8:7-20




FIG.4

Prosecution History

6. (Currently Amended) Method of communication service mapping in an integrated circuit, having a plurality of processing modules (M, S), wherein at least one first of said processing modules (M) requests at least one communication service to at least one second processing module (S) 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:

coupling said plurality of processing modules (M, S) by an interconnect means (N) and
enabling a connection based communication having a set of connection properties, controlling the communication between said at least one first of said plurality of processing modules (M) and said interconnect means (N) by at least one network interface (NI) associated to said at least one first of said processing modules,
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.

**“thread or address range” added
to gain allowance**

The cited portions of Radilescu fail to disclose or suggest “a mapping means (A) for 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, 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)**”, as recited in Claim 1 (Emphasis Added).

★ ★ ★ ★

Applicants respectfully submit that, **in contrast to Radilescu**, the invention is based on the idea of offering differentiated services for protocols such as DTL, MTL, AXI, and OCP, by mapping the **identification means** of these protocols to connections. **The identification means in the existing protocols are: communication threads and addresses**, i.e. the **threads or addresses** are mapped to connections through the interconnect based on specific connection properties. Applicants submit that this is different from Radilescu which describes end-to-end flow control and guaranteed throughputs, but is silent with respect to offering differentiated services for protocols such as DTL, MTL, AXI, and OCP, by mapping the identification means of these protocols to connections.

Argument made with amendment

Ex-1002, pp. 138-141

Slide 22

Limitation 6.d (mapping step)

- PO's description:

The claimed “mapping” requires: “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.” EX1001, 10:28-32; EX2070, ¶164. As

Paper 38, p. 43

Limitation 6.d (mapping step)

The claimed “mapping” requires: “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.” EX1001, 10:28-32; EX2070, ¶164. As

refer to a specific example in Applicant’s specification. Fig. 4 is a block diagram for specifically illustrating the operation of mapping (i.e., mapping means) of IP block communication via address ranges to connections. As described below, the network interface, NI, associated to the master M maps the communication to the first slave to an address range 000-7FFF to a connection id cid=0 and performs a similar operation for mapping the communication to the second slave.

Ex-1002 (Prosecution History), p.142





FIG.4

Limitation 6.d (mapping step)

- Radulescu2004 teaches mapping “to a connection based on a set of connection properties according to said at least one communication service identification”




In Figure 3, we show an example of a narrowcast shell. Narrowcast connections are connections between one master and several slaves, where each transaction is executed by a single slave selected based on the address provided in the transaction [23]. Narrowcast connections provide a simple, low-cost solution for a single shared address space mapped on multiple memories. It implements the splitting/merging of data going to/coming from these memories.

We implement the narrowcast connection as a collection of point-to-point connections, one for each master-slave pair. Within a narrowcast connection, the slave for which the transaction is destined is selected based on the address (Conn block). The address range assigned to a slave is configurable in the narrowcast module. To provide in-order response delivery, the narrowcast must also keep a history of connection identifiers of the transactions including responses (e.g., reads, and acknowledged writes), and the length of these responses. In-order delivery per slave of request messages is already provided by the point-to-point connections.

Ex-1005, p. 3

Limitation 6.d (mapping step)

- Radulescu2004 teaches mapping “to a connection based on a set of connection properties according to said at least one communication service identification”



Ex-1005, FIGS. 3, 2 (annotations added)

The NI kernel (see Figure 2) receives and provides messages, which contain the data provided by the IP modules via their protocol after sequentialization. The message structure may vary depending on the protocol used by the IP module. However, the message structure is irrelevant for the NI kernel, as it just sees messages as pieces of data to be transported over the NoC.

The NI kernel communicates with the NI shells via *ports*. At each port, point-to-point connections can be configured, their maximum number being selected at NI instantiation time. A port can have multiple connections to allow differentiated traffic classes, in which case there are also *connid* signals to select on which connection a message is supplied or consumed.

Ex-1005, p. 3

IPR2024-00343
('052 Patent IPR)


Ground 2: Obviousness
Radulescu2004 in view of OCP
(Claims 1-7)

Open Core Protocol v. 1 (Ex-1006)

The OCP defines a point-to-point interface between two communicating entities such as IP cores and bus interface modules (bus wrappers). One entity acts as the master of the OCP instance, and the other as the slave. Only the master can present commands and is the controlling entity. The slave responds to commands presented to it, either by accepting data from the master, or presenting data to the master. For two entities to communicate in a peer-to-peer fashion, there need to be two instances of the OCP connecting them - one where the first entity is a master, and one where the first entity is a slave.

Figure 1 shows a simple system containing a wrapped bus and three IP core entities: one that is a system target, one that is a system initiator, and an entity that is both.

Figure 1 System Showing Wrapped Bus and OCP Instances



Ex-1006, p. 14

Open Core Protocol (OCP)

- OCP signals include:
 - MAddrSpace – Address Range
 - MThreadID – Thread

Simple Extensions

Table 4 lists the simple OCP extensions. The extensions add address spaces, byte enables, burst support, datahandshake, and response flow control.

Table 4 Simple OCP Extensions

Name	Width	Driver	Function
MAddrSpace	1-8	master	Address space
MBurst	3	master	Burst code
MByteEn	1/2/4/8/16	master	Byte enables
MDataValid	1	master	Write data valid
MRespAccept	1	master	Master accepts response
SDataAccept	1	slave	Slave accepts write data

MAddrSpace

Address Space. This field is an extension of the MAddr field and is used to indicate the address region of a transfer. Examples of address regions are the register space versus the regular memory space of a target core or the user versus supervisor space for an initiator core.

Ex. 1006, 26

Open Core Protocol (OCP)

- OCP signals include:
 - MAddrSpace – Address Range
 - MThreadID – Thread

Complex Extensions

Table 6 shows a list of complex OCP extensions. The extensions add support for threads and connections.

Table 6 Complex OCP Signal Extensions

Name	Width	Driver	Function
MConnID	1-8	master	Connection identifier
MDataThreadID	1-4	master	Write data thread identifier
MThreadBusy	1-16	master	Master thread busy
MThreadID	1-4	master	Request thread identifier
SThreadBusy	1-16	slave	Slave thread busy
SThreadID	1-4	slave	Response thread identifier

* * * *

MThreadID

Request thread identifier. This variable-width field provides the thread identifier associated with the current transfer request. If threads is greater than 1, this field is enabled. The field width is the next whole integer of $\log_2(\text{threads})$.

Ex. 1006, 27-28

Motivation to Combine – MAddrSpace

In this paper we present a network interface for an on-chip network. Our network interface decouples computation from communication by offering a shared-memory abstraction, which is independent of the network implementation. We use a transaction-based protocol to achieve backward compatibility with existing bus protocols such as AXI, OCP and DTL. Our network interface

Ex-1005 (Radulescu2004), Abstract

MAddrSpace

Address Space. This field is an extension of the MAddr field and is used to indicate the address region of a transfer. Examples of address regions are the register space versus the regular memory space of a target core or the user versus supervisor space for an initiator core.

Ex. 1006 (OCP Specification), p. 26

provide such analysis in my First Declaration. My First Declaration provides that a POSA would have been motivated to achieve backwards compatibility with OCP, as expressly taught by Radulescu2004. Ex. 1003, ¶¶154-155. A POSA would know that in order to provide backward compatibility with IP modules that use the combination of MAddr and MAddrSpace, the adapter units in the network interfaces of Radulescu2004's system would need to work with both MAddr and MAddrSpace. Ex. 1006, 38 (annotated figure below indicating that MAddr and MAddrSpace are "always" part of a request if these signals are required):

* * * *

This is because some architectures take advantage of the MAddrSpace signals to identify the "address region of a transfer" such as "the register space versus the regular memory space." Ex. 1006, 26. The MAddr signal would then provide the location within the "address region" in order to identify the proper data. This is

Ex. 1030 (Hassoun Reply Dec., ¶ 39

Motivation to Combine – MAddrSpace

We implement the narrowcast connection as a collection of point-to-point connections, one for each master-slave pair. Within a narrowcast connection, the slave for which the transaction is destined is selected based on the address (Conn block). The address range assigned to a slave is configurable in the narrowcast module. To provide in-order response delivery, the narrowcast must

Ex-1005 (Radulescu2004), p. 4

MAddrSpace

Address Space. This field is an extension of the MAddr field and is used to indicate the address region of a transfer. Examples of address regions are the register space versus the regular memory space of a target core or the user versus supervisor space for an initiator core.

Ex. 1006 (OCP Specification), p. 26

40. For the system of Radulescu2004 to be compatible with an OCP IP module that uses MAddr and MAddrSpace, the mapping would need to be based on both addressing signals. Dr. Bagherzadeh argues that “OCP’s standard MAddr

★ ★ ★ ★

the two sentences to be discussing the same process. A POSA would understand Radulescu2004 discloses a connection to a slave is chosen via the address in the request and such an identification of which slave such an address belongs to is through the address range of that slave that has been configured in the narrowcast module. For an OCP IP module that uses MAddr and MAddrSpace, a POSA would be motivated to use the MAddrSpace, in addition to MAddr, for mapping in order to identify the address region of the transfer. Ex. 1006, 26. A POSA would have been motivated to use both the MAddr and MAddrSpace signals for such an architecture in the mapping described by Radulescu2004.

Ex. 1030 (Hassoun Reply Dec., ¶ 40

Motivation to Combine – MThreadID

157. One of the Request signals is “MThreadID,” which “provides the thread identifier associated with the current transfer request.” Ex-1006 at 28. OCP teaches many advantages of using threads, including “support[ing] many different burst styles” and “support[ing] concurrent activity, and out-of-order completion of transfers”:

* * * *

Id. 181. A POSA would have been motivated to use the OCP threads in Rădulescu2004 to get these known benefits of threading. See *supra*, paragraph 55.

Ex-1003 (Hassoun Dec.) ¶ 157

The burst type is thread and connection specific. By using multiple threads or connections, a slave can support many different burst styles, such as different cache line sizes for different processors.

Threads and Connections

All transfers within a single thread must remain ordered. (See “Phase Ordering Between Transfers” on page 29.)

Using multiple threads, it is possible to support concurrent activity, and out-of-order completion of transfers. All transfers within a given thread must remain strictly ordered, but there are no ordering rules for transfers that are in different threads. Mapping of individual requests and responses to threads is handled through the MThreadID and SThreadID fields respectively. If datahandshake has been enabled when thread ID has been enabled, there must also be an MDataThreadID field to annotate the datahandshake phase.

Ex-1006 (OCP Specification), p. 46

Motivation to Combine – MThreadID

157. One of the Request signals is “MThreadID,” which “provides the thread identifier associated with the current transfer request.” Ex-1006 at 28. OCP teaches many advantages of using threads, including “support[ing] many different burst styles” and “support[ing] concurrent activity, and out-of-order completion of transfers”:

* * * *

Id. 181. A POSA would have been motivated to use the OCP threads in Rădulescu2004 to get these known benefits of threading. *See supra*, paragraph 55.

Ex-1003 (Hassoun Dec.) ¶ 157

Threads

The thread capability relies on a thread ID to identify and separate independent transfer streams (threads). The master labels each request with the thread ID that it has assigned to the thread. The thread ID is passed to the slave on MThreadID together with the request (MCmd). When the slave returns a response, it also provides the thread ID (on SThreadID) so the master knows which request is now complete.

The transfers in each thread must remain in-order with respect to each other (as in the basic OCP), but the order between threads can change between request and response.

The thread capability allows a slave device to optimize its operations. For instance, a multi-bank SDRAM controller could respond to a second read request referencing an open page before opening a new page in a different bank to service a first read on a different thread.

In multi-threaded, multi-initiator systems, it is frequently useful to associate a transfer request with a thread operating on a particular initiator. Initiator identification enables busy targets (such as multi-bank SDRAM controllers) to prioritize requests. For instance, a memory controller that is completing a first burst transaction needs to determine which of several pending requests to service next. In many systems, the decision should be based upon priority; for example, real-time frame buffer transfers have higher priority than polygon data transfers. The priority is typically associated with the requesting initiator. In some cases, the initiator has several assigned priorities based upon the type of request. If different request types are assigned to different system initiator threads, the priority is dependent upon a thread within a particular initiator.

Ex-1006 (OCP Specification), p. 181

Motivation to Combine – MThreadID

above with respect to Ground 1. *Supra*, Ground 1, element [6.d]. Second, by mapping OCP threads to Rădulescu2004's connections, the NoC can perform thread-based ordering and priority can be communicated through the NoC "between a system initiator and a system target," and implement Rădulescu2004's multiple connections to a port in order to allow differentiated traffic classes through the connid signals. See Ex-1006 at 180–181; see also, for example, Ex-1005 at §4.1.

Ex-1003 (Hassoun Dec.) ¶ 177

for the reasons I discussed in my First Declaration. For example, in my First Declaration, I explain OCP teaches MThreadID is mapped with a connection identifier (MConnID), and that a POSA would have been motivated to have the OCP adapter in the network interface of Radulescu2004 map the MThreadID in requests to the analogous connections in Radulescu2004. Ex. 1003, ¶158. Mapping OCP threads to Radulescu2004's connections enables thread-based ordering and priority to be communicated through the NoC between system initiators and a system target and enables implementation of Radulescu2004's differentiated traffic classes. Ex. 1003, ¶177. A POSA would have been motivated to make such a mapping. The whole point of having connections with properties is to send requests for which particular properties are desired (e.g., ordering or high priority) over a connection with the desired property (e.g., ordering or high priority).

Ex-1030 Hassoun Reply Dec., ¶ 50

Motivation to Combine

The Board rejected Intel's known-technique rationale. The Board stated that "[i]f . . . Kabemoto already addresses [the] problem [of cache coherency] through the use of a known technique similar to that of Bauman's, [it] fail[ed] to see why one of ordinary skill in the art would regard Bauman's technique as an obvious improvement to Ka-bemoto." *Final Written Decision*, 2021 WL 3503434, at *11. ***But the Board's reasoning belies its conclusion. That Ka-bemoto and Bauman address the same problem and that Bauman's cache was a known way to address that problem is precisely the reason that there's a motivation to combine under KSR and our precedent.***

* * * *

And contrary to the Board's suggestion, ***Intel never had to show that replacing Kabemoto's secondary cache with Bauman's secondary cache was an "improvement" in a categorical sense.*** See *Final Written Decision*, 2021 WL 3503434, at *11. Intel just had to show that Bauman's secondary cache ***was a "suitable option" to replace Kabemoto's secondary cache.***

Intel Corp. v. PACT XPP Schweiz AG, 61 F.4th 1373, 1380 (Fed. Cir. 2023).

IPR2024-00343
('052 Patent IPR)

Ground 3: Anticipation by Weber
(Claims 1-7)

Weber (Ex-1007)

- System-on-Chip (SoC) that includes:
 - Functional blocks (A-E)
 - Interfaces (120-128)
 - Point-to-Point buses (115)
 - Shared buses (112, 114)
 - Bridge between shared buses (F)




FIG. 1

Sole Issue for Claim 6

**Does Weber disclose an “integrated circuit”
(single chip) with an “interconnect means”?**

“Interconnect Means”

Using **networks** for on-chip communication when designing systems on chip (SoC), however, raises a number of new issues that must be taken into account. This is because, **in contrast to existing on-chip interconnects (e.g., buses, switches, or point-to-point wires)**, where the communicating modules are **directly connected**, in a NoC the modules communicate remotely via network nodes. As a result, intercon-

‘052 Patent, col. 1:47-54

Introducing **networks as on-chip interconnects** radically changes the communication when compared to **direct interconnects, such as buses or switches**. This is because of the multi-hop nature of a network, where communication modules are not directly connected, but separated by one or more network nodes. This is in contrast with the prevalent **existing interconnects (i.e., buses)** where modules are directly connected. The implications of this change reside in the arbitration (which must change from centralized to distributed), and in the communication properties (e.g., ordering, or flow control).

‘052 Patent, col. 2:32-42

“Network”

- PO’s construction: “Multi-hop interconnection between modules separated by one or more network nodes”


Introducing networks as on-chip interconnects radically changes the communication when compared to direct interconnects, such as buses or switches. This is because of the multi-hop nature of a network, where communication modules are not directly connected, but separated by one or more network nodes. This is in contrast with the prevalent existing interconnects (i.e., buses) where modules are directly connected. The implications of this change reside in the arbitration (which must change from centralized to distributed), and in the communication properties (e.g., ordering, or flow control).

‘052 Patent, col. 2:32-42


The '052 patent uses “hop” to describe the **link** between components

“This is because of the **multi-hop** nature of a network, where communication modules are **not directly connected**, but separated by **one** or more network nodes.

'052 Patent, col. 2:34-37



Direct connection = one hop




“**One** or more” nodes = “multi-hop”

Weber discloses an “integrated circuit” with an “interconnect means”

Owner asserts that “interconnect means” is not a Section 112(6) term and should be construed as “multi-hop interconnection between modules separated by one or more network nodes.” Ex-1017 at 25. While I disagree with that construction, Weber discloses such a “multi-hop” interconnection. For example, Weber also indicates that system 100 can include a bridge, which a POSA would understand refers to a network node that serves as an interface between buses in a multi-hop interconnection. Ex-1007 at [0041]. Figure 1 of Weber illustrates such a multi-hop interconnection (annotated in red) between sub-system A (highlighted in yellow) of block 102 and sub-system A (highlighted in yellow) of block 110 through interfaces 120, 124, 128, and block F (highlighted in yellow).

Ex-1003 (Hassoun Dec.) ¶199



Ex-1007 (Weber)

FIG. 1

Weber discloses an “integrated circuit” with an “interconnect means”

- PO’s assertion that integrating all subsystems in Fig. 1 was “impractical,” and/or “unusual” is both irrelevant and wrong.

does not disclose them. Moreover, it was known, even at the time of Weber, that integrating DRAM with a processor was both possible and beneficial. See Ex. 1053 (e.g. at abstract “A *microprocessor integrated with DRAM on the same die* has the potential to *improve system performance* by reducing the memory latency and improving the memory bandwidth.”)

Ex-1030 (Hassoun Dec.) ¶55

57. This was well-known in the industry. For example, an integrated circuit product from an industry leader in 2002 shows a single integrated circuit with DRAM, SRAM, ROM, FLASH, Microprocessor, DSP, Peripherals and ASIC.




Figure 1 System-on-a-Chip for Embedded Applications


Ex. 1054, 8.

Ex-1030 (Hassoun Dec.) ¶56

Weber discloses an “integrated circuit” with an “interconnect means”

72. For example, sub-system 102 includes a “multi-hop interconnection between modules separated by one or more network nodes,” because the functional blocks A, C, E, and F are “separated by one or more network nodes” due to the respective interface 120, 122, 123, 124 in front of each of those blocks. The interconnect coupling those blocks together is also “multi-hop” because the communications must hop to and from these interfaces 120, 122, 123, 124 to get to and from those functional blocks. *See* Ex. 1003, ¶194-197, 199-201; Ex. 2068, 22:8-12, 26:6-12, 111:20-114:13. For example, I showed an example of this at ¶194 of my First Declaration, for the path from A to C:

Ex-1030 (Hassoun Reply Dec.) ¶72




Ex-1007 (Weber)

FIG. 1

Interfaces are “network nodes”

modules to the NoC.” *Id.* at 6:33-36. An example of the network interface with adapter units is shown in FIG. 3, reproduced below.



Ex. 1001 at FIG. 3 (annotated)

Ex-2070 (Bagherzadeh Dec.), ¶86
(annotations in original)

5. Weber’s Bus-with-Interfaces is a “Network”

85. As mentioned above, Weber’s system 100 discloses an interconnect including multiple links/buses (112, 114, 115), multiple interfaces (120, 122, 123, 124, 125, 126, 127, 128), and a bridge (F) interconnecting a plurality of functional blocks. This bus-with-interface topology is a “network” as that term is used in the ‘052 patent. As described above, each of the interfaces (and the bridge F) create a hop making Weber a “multi-hop interconnect.” *Supra*, §V.B, §VI.C.4. Each of the interfaces (and the bridge F) are also part of the network and therefore “network nodes,” thus the functional blocks are “separated by one or more network nodes.” *Supra*, §V.B, §VI.C.4; Ex. 1003, ¶¶193-196, 209.

Ex-1030 (Hassoun Reply Dec.), ¶85

Interfaces are “network nodes”

- The only “network nodes” in Figure 4 are the network interfaces




FIG.4

Interfaces are “network nodes”


16	Q. So in that paragraph 86, are you saying
17	that in the alternative where the interconnect means
18	is a network, that the network interface can't be
19	part of the network?
20	MR. CHRISTOPHER: Objection. Form.
21	Leading.
22	THE WITNESS: No, that's not what I'm
23	saying here.
24	BY MR. VANDENBURGH:
25	Q. What are you saying in that paragraph
Page 299	
1	relative to the network interface?
2	A. What I'm saying here is a network
3	interface on its own is not a network in the context
4	of the '052 Patent. And I'm going to read what it
5	says. "But excluding the network interface NI
6	because it is separately claimed in claim 6."
7	Q. Is there any question in your mind that
8	the '052 Patent discloses the network interfaces as
9	being part of the disclosed networks?
10	A. No. I have no doubt.

Ex-2084 (Hassoun Dep.), pp. 298-99

Weber discloses an “integrated circuit” with an “interconnect means”

203. Regarding the discussion of Figure 14 in the above chart for element [6.a], the travel from a first communication subsystem to a second communication subsystem via an interface is a “multi-hop interconnection between modules separated by one or more network nodes.” Referring to Figure 14, the data flow highlighted in red includes at least two “hops.” A first hop is from the initiator functional block 1410 to the interface 1460 and a second hop is from the interface 1460 to the target functional block 1430. Thus, the interconnection is “multi-hop”. The interface 1460 is also a network node and separates the initiator functional block 1410 and the target function block 1430. Thus, the interconnection between functional blocks 1410, 1430 is “separated by one or more network nodes”. Thus, Weber’s Figure 14 discloses a “multi-hop interconnection between modules separated by one or more network nodes”.

Ex-1003 (Hassoun Dec.) ¶203




Ex-1007 (Weber)

Weber discloses an “integrated circuit” with an “interconnect means”

203. Regarding the discussion of Figure 14 in the above chart for element [6.a], the travel from a first communication subsystem to a second communication subsystem via an interface is a “multi-hop interconnection between modules separated by one or more network nodes.” Referring to Figure 14, the data flow highlighted in red includes at least two “hops.” A first hop is from the initiator functional block 1410 to the interface 1460 and a second hop is from the interface 1460 to the target functional block 1430. Thus, the interconnection is “multi-hop”. The interface 1460 is also a network node and separates the initiator functional block 1410 and the target function block 1430. Thus, the interconnection between functional blocks 1410, 1430 is “separated by one or more network nodes”. Thus, Weber’s Figure 14 discloses a “multi-hop interconnection between modules separated by one or more network nodes”.

Ex-1003 (Hassoun Dec.) ¶203



Ex-1007 (Weber)


PO's Improper Attempt to Read Unclaimed Limitations Into “Network”

PO's construction: “Multi-hop interconnection between modules separated by one or more network nodes”

- PO argues that to be a “network”:
 1. Nodes must be routers or switches
 2. Must provide protocol conversion
 3. Must provide “route variability”
 4. Must transmit packets

PO's Sur-Reply

Likewise, in the example below, one hop occurs when a packet traverses bridge 101 and another occurs when traversing bridge 104. EX1030, ¶83; EX2010, 461 (describing multiple hops between stations 1 and 5); EX2084, 259:23-260:12.



Paper 55, pp. 3-4

Another problem with Dr. Hassoun's new opinion is that he defines a "hop" using a generic "node." EX1030, ¶18. For instance, Dr. Hassoun asserts there is one "hop" between two directly linked modules. *Id.* However, a "network node" is a component (e.g., router or switch) "in a network," i.e., the interconnect means. EX1030, ¶21; EX1050, 58:21-59:3; EX2084, 202:6-10 (explaining that "a node, if it's not part of a network, then it's not a network node"). The hop must occur *within* the interconnect means. EX2070, ¶¶ 113-22; EX2068, 26:22-27:21.

Paper 55, pp. 5-6

PO's Improper Attempt to Read Unclaimed Limitations Into "Network"

PO's construction: "Multi-hop interconnection between modules separated by one or more network nodes"

- Ordinary meaning of "network"

network (1) (measuring longitudinal balance of telephone equipment operating in the voice band) A combination of elements or devices. (COM/TA) 455-1985w


(2) (A) (data transmission) A series of points interconnected by communication channels. **(B) (data transmission)** The switched telephone network is the network of telephone lines normally used for dialed telephone calls. **(C) (data transmission)** A private network is a network of communications channels confined to the use of one customer.

(COM/SUB/PE) 999-1992, 599-1985

(3) (A) (software) An interconnected or interrelated group of nodes. *See also:* documentation; node. **(B) (software)** In connection with a disciplinary or problem oriented qualifier, the combination of material, documentation, and human resources that are united by design to achieve certain objectives, for example, a social science network, a science information network. *See also:* node; documentation. (C/SE) 729-1983

(4) (distribution of electric energy) An aggregation of interconnected conductors consisting of feeders, mains, and services. *See also:* alternating-current distribution; network analysis. (T&D/PE) [10]

(5) (A) (data management) A data structure in which components are allowed to have more than one superordinate component. *See also:* hierarchy; graph. **(B) (data management)** A graph in which the edges connecting the nodes are assigned weights representing some characteristic, such as cost or quantity, related to the edge.



(C) 610.5-1990

(6) A collection of interconnected hosts. The term network in this document is used to refer to the network of hosts. The batch system is used to refer to the network of batch servers. (C/PA) 1003.2d-1994

(7) A network is any set of devices or subsystems connected by links joining (directly or indirectly) a set of terminal nodes. (C/BA) 1355-1995

(8) An arrangement of components, or nodes, and interconnecting branches. *See also:* computer network; circuit; store-and-forward switched network; public data network; public circuit-switched network; net; decentralized computer network; centralized computer network; packet switching network; distributed computer network; communications network; public switched network; switched network; value-added network; circuit-switched network; private network; public packet switching network; software defined network. (C) 610.7-1995, 610.10-1994w

(9) A communication mechanism for interconnecting multiple Smart Transducers. (IM/ST) 1451.1-1999

Ex-1052 (IEEE Dictionary), pp. 5-6

Slide 54

PO's Improper Attempt to Read Unclaimed Limitations Into "Network"

PO's construction: "Multi-hop interconnection between modules separated by one or more network nodes"

Ordinary meaning of "hop"

hop 1. A small, quick jump. 2. In networks, the journey a signal makes across a transmission link between two devices such as bridges, hubs, switches, or routers. There often are multiple links in an end-to-end circuit between two devices. Therefore, a frame may make multiple hops as it transverses a physical path from one workstation to another. An IP packet typically makes a significant number of hops as it transverses the Internet. See also *bridge, circuit, frame, hub, Internet, LAN, link, packet, router, and switch*. 3. In radio communications, one skip of a radio wave from an earth station to the ionosphere and back. See also *ionosphere, radio, and waveform*. 4. In satellite communications, one roundtrip of a signal from an Earth station to a space station (i.e., satellite) and back to an Earth station. See also *satellite and signal*. 5. In frequency hopping spread spectrum (FHSS) radio communications, a small, quick jump from one frequency channel to another in a carefully choreographed hop sequence. See also *channel, FHSS, and frequency*.

Ex-1051 (Webster's Telecom Dictionary), pp. 6-7

Dr. Hassoun's Testimony on "Packets"

what exactly he contends is required. In any case, I do not believe the claims require any of these aspects.

75. With respect to packets, none of the claims use the word "packet" and Dr. Bagherzadeh could not identify any language in the specification that indicated that the claims required packets. Ex. 1050, 65:24-74:12.

* * * *

The IEEE dictionary also suggests that packets are not required for all networks. It's definitions of "network" do not require packets and it indicates that the more specific term "packet switching network" is used for networks that use packets.

Supra, ¶80; Ex. 1052, 7 ("packet switching network A network that uses packet switching techniques for transmission of data. *Synonyms*: packet switched network, packed data network.")

Ex-1030 (Hassoun Dec.), ¶¶ 74-75

16 Q. I'm just trying to understand what the hop is.
17 Would the communication pathway between two network
18 nodes and a network be considered a hop?
19 A. I don't know what you mean by a "communication
20 pathway." Can you give me an example?
21 Q. Sure.
22 A network on a chip, packets are transmitted
23 over the network, right?
24 A. Yes.
25 Q. And the network can include multiple network

Page 27


1 nodes, right?
2 A. Correct.
3 Q. And so the packets can travel from one network
4 node to a second network node, right?
5 A. Correct.
6 Q. So is the hop -- let me say it this way.
7 Strike that.
8 When a packet travels from a first network
9 node to a second network node is that a hop?
10 A. If I understand your example correctly, what
11 you described is a direct connection between one network
12 node to another network node.

Ex-2068 (Hassoun Dep.), pp. 26-27

Additional Issue for Claims 1-5, 7

Does Weber disclose the same or equivalent structure for the claimed “mapping means”?

Weber Shows the “Mapping Means” Structure (Claims 1-5, 7)



Structure in '052 Patent Linked to Mapping Function (Hassoun Dec. ¶¶ 85)

224. Weber teaches such an adapter unit. The interface modules of Weber perform the function of the claimed mapping as discussed with respect to element 6.d. As explained below, Weber's interface modules also translate command,

* * * *


225. Weber discloses that the address, data, and control signals can be multiplexed over common wires on the shared communication buses 112, 114. Ex-

* * * *

higher on a single integrated circuit chip.”). Multiplexing the address, data and control signals over common wires is sequentialization. Based on this description in Weber, a POSA would at once envisage that the interface modules of system 100 were intended to perform such multiplexing. Thus, the interface modules of system 100 are the “mapping means” recited in element [1.e].

Ex-1003 (Hassoun Dec.), ¶¶ 224-225

Weber Shows the “Mapping Means” Structure (Claims 1-5, 7)



Structure in '052 Patent Linked to Mapping Function (Hassoun Dec. ¶¶ 85)

up a basic set of interface module signals.”). For example, as shown in FIG. 4 and FIG. 11, a POSA would understand that this would require the initiator interface module 1004 to multiplex the parallel signals on interconnect 1010 from initiator functional block 1002 for placement onto the bus 1012 and would also require target interface module 1006 to de-multiplex the signals from bus 1012 for placement as parallel signals onto the interconnect 1010 to the target functional block 1008. Ex. 1001, FIGs. 4, 11 (annotated below).




FIG. 4

Hassoun Reply Dec., ¶¶ 62

IPR2024-00343
('052 Patent IPR)

Ground 4: Obviousness
Weber in view of Radulescu2004
(Claims 1-7)

Combination was Explained

How to Combine

Rădulescu2004 is directed to the same goals addressed by Weber, and a POSA would have been motivated to combine Weber's disclosures with Rădulescu2004's disclosures regarding its network on chip (NoC) to obtain the benefits of Rădulescu2004's known solution to "the interconnect problem for highly complex chips ..." See Ex-1005, 1; Ex-1003, ¶237. In particular, a POSA would have looked to Rădulescu2004's disclosure of a NoC to perform the interconnect role of the shared communications bus disclosed by Weber. Ex-1003, ¶241. Like Weber's

Paper 3 (Petition), p. 69

Motivation to Combine

241. Rădulescu2004 confirms these benefits. As Rădulescu2004 states, one of the known advantages of Rădulescu's NoC was improved scalability. Ex-1005 at 1 §1 ("NoCs ... scale better than buses ..."). A POSA would have understood that using Rădulescu2004's NoC would improve the scalability and performance of Weber's communication system. A POSA also would have understood that a router network like Rădulescu2004's NoC can support more concurrent data transmissions than a bus-based system such as that of Weber, and thus enable a larger amount of data to be communicated between the various modules and blocks in an SoC. A POSA would have been motivated to use Rădulescu2004's NoC (including its network interfaces made up of shells and a kernel) in place of the shared buses of Weber in order to achieve these known benefits, thereby enabling the same interconnect to be used even as more and more functional blocks are integrated into an SoC.

Ex-1003 (Hassoun Dec.) ¶241

Combination was Explained





Figure 8. NI configuration



“Integrated Circuit”

Combination

A POSA would have been motivated to combine the teachings of Weber with those of Rădulescu2004 in order to create SoCs with greater scalability and performance. Ex-1003, ¶¶236-241.

Radulescu2004’s Teaching

Rădulescu2004 is directed to the same goals addressed by Weber, and a POSA would have been motivated to combine Weber’s disclosures with Rădulescu2004’s disclosures regarding its network on chip (NoC) to obtain the benefits of Rădulescu2004’s known solution to “the interconnect problem for highly complex chips ...” See Ex-1005, 1; Ex-1003, ¶237. In particular, a POSA would have looked

Weber’s Teaching

Weber describes that advances in semiconductor technology have enabled complex SoC designs that incorporate large numbers of functional blocks, such as processor cores, digital signal processing cores, communication interfaces, and application-specific logic. Ex-1007, [0002]; Ex-1003, ¶236. Weber explains that

Integrated Circuit Element

1. [6.pre.1]

Both Weber and Rădulescu2004 disclose this part of the preamble for the reason in Grounds 1 and 3. See *supra* X.A.1, VIII.A.1.

Paper 3 (Petition), p. 68-71

END