

1 B. Trent Webb (*pro hac vice*)
Lauren E. Douville (*pro hac vice*)
2 Ryan J. Schletzbaum (*pro hac vice*)
SHOOK, HARDY & BACON LLP
3 2555 Grand Boulevard
Kansas City, Missouri 64108-2613
4 T: 816.474.6550
F: 816.421.5547
5 bwebb@shb.com
ldouville@shb.com
6 rschletzbaum@shb.com

7 Amelia Murray (*pro hac vice*)
SHOOK, HARDY & BACON LLP
8 111 S. Wacker Dr., Suite 4700
Chicago, IL 60606
9 T: 312.704.7700
F: 312.558.1195
10 aemurray@shb.com

11 Megan Whyman Olesek (SBN 191218)
Keeley I. Vega (SBN 259928)
12 TURNER BOYD SERAPHINE LLP
155 Bovet Road, Suite 750
13 San Mateo, CA 94402
T: 650.521.5930
14 olesek@turnerboyd.com
vega@turnerboyd.com

15 *Attorneys for Defendant HP Inc.*

16
17 **UNITED STATES DISTRICT COURT**
NORTHERN DISTRICT OF CALIFORNIA
18 **OAKLAND DIVISION**

19 UNIVERSAL CONNECTIVITY
TECHNOLOGIES INC.,

20 Plaintiff,

21 v.

22 HP INC.,

23 Defendant.

Case No. 4:24-cv-04097-YGR

**HP'S INVALIDITY CONTENTIONS
PERSUANT TO PATENT LOCAL RULE
3-3**

24
25 **HP'S INVALIDITY CONTENTIONS**

26 Pursuant to Patent Local Rule 3-3 and Court's Scheduling Order (Dkt. 93), defendant HP
27 Inc. ("HP" or "Defendant") respectfully submits these invalidity contentions with respect to the
28 asserted claims of U.S. Patent Nos. 7,154,905 ("the '905 Patent"), 7,187,307 ("the '307 Patent"),

HP'S INVALIDITY CONTENTIONS PERSUANT TO PATENT LOCAL RULE 3-3
Case No. 4:24-cv-04097-YGR

- 1 -

1 7,746,798 (“the ’798 Patent”), 9,232,265 (“the ’265 Patent”), 8,680,712 (“the ’712 Patent”),
2 7,856,520 (“the ’520 Patent”), 7,921,231 (“the ’231 Patent”), and 9,852,103 (“the ’103 Patent”)
3 (collectively, the “Asserted Patents”) identified by Plaintiff Universal Connectivity Technologies
4 Inc., (“UCT” or “Plaintiff”) in its Infringement Contentions.

5 The currently Asserted Claims, as reflected in Plaintiff’s Infringement Contentions, are:

- 6 • ’905 Patent: Claims 1, 8, 15, 21, 22, 23
- 7 • ’307 Patent: Claims 1, 7, 23, 25, 53, 68, 72
- 8 • ’798 Patent: Claims 11, 19, 22, 25, 28
- 9 • ’265 Patent: Claims 1, 8, 11, 14
- 10 • ’712 Patent: Claims 1, 5, 7
- 11 • ’520 Patent: Claims 1, 8, 10, 12, 13, 15, 19
- 12 • ’231 Patent: Claims 1, 2, 10, 11, 12, 13, 16
- 13 • ’103 Patent: Claims 1, 3, 11, 17, 21, 22

14 As detailed further below, the asserted claims of each Asserted Patent are anticipated by
15 and/or obvious in view of one more of the prior art references being produced concurrently pursuant
16 to 35 U.S.C. § 102 and/or § 103, are directed to ineligible subject matter under 35 U.S.C. § 101, and
17 are otherwise invalid under 35 U.S.C. § 112.

18 **I. PRELIMINARY STATEMENT**

19 These invalidity contentions and eligibility contentions are based on Defendant’s current
20 knowledge, understanding, and belief of the Asserted Patents and prior art, of Plaintiff’s
21 infringement theories (inasmuch as they can be inferred from its Infringement Contentions), and of
22 the facts and other information available as of the date of these invalidity contentions. Defendant’s
23 investigation, discovery, and analysis of information related to this action is ongoing. Additional
24 discovery, elucidation of Plaintiff’s impermissibly vague infringement contentions, Plaintiff’s
25 expert report(s) regarding infringement, and/or orders of the Court may require Defendant to amend
26 or supplement these invalidity contentions and eligibility contentions, and Defendant expressly
27 reserves the right to do so, subject to the Court’s procedures, as the case proceeds. These contentions
28 represent Defendant’s good-faith effort to provide a comprehensive identification of prior art

1 relevant to this case, but Defendant reserves the right to modify or supplement its prior art list and
2 invalidity contentions and eligibility contentions at a later time with, or based upon, pertinent
3 information that may be subsequently discovered.

4 In addition to the prior art that is charted in the attached exhibits, additional background
5 references are discussed herein and/or included in the document production accompanying these
6 Invalidity Contentions. Defendant reserves the right to rely on these background references for any
7 purpose including, for example, to demonstrate the knowledge of a POSITA at the relevant time.

8 To the extent UCT believes anything in these contentions (including the attached claim chart
9 exhibits and accompanying document production) are insufficient, Defendant invites UCT to bring
10 such issues to its attention as soon as reasonably possible.

11 **A. No Waiver**

12 Nothing in these contentions is intended, nor should be construed, as a waiver of any non-
13 infringement position or argument under 35 U.S.C. §§ 101 or 112. To the contrary, as set forth
14 herein, Defendant contends that all Asserted Claims are invalid pursuant to 35 U.S.C. §§ 101 and/or
15 112. Defendant's statements herein (including the accompanying claim charts) reflect Defendant's
16 present understanding of the purported scope of the claims as alleged by Plaintiff in its Infringement
17 Contentions (as best those contentions can be understood in light of their present deficiencies).

18 The patent claims have yet to be construed. As a result, Defendant has based these
19 contentions upon its knowledge and understanding of the potential scope of the Asserted Claims at
20 this time, and, in part, upon the apparent interpretations of the Asserted Claims advanced by Plaintiff
21 in its Infringement Contentions. Defendant may disagree with Plaintiff's interpretation of the
22 meaning of many terms and phrases in the Asserted Claims. Defendant has provided these
23 contentions based in part on its present understanding of Plaintiff's apparent constructions and
24 interpretations of the Asserted Claims. These contentions do not represent Defendant's agreement
25 or view as to the proper interpretation of any claim term contained therein. Any similarity between
26 any apparent claim interpretation in any of Defendant's charts of prior art references and Plaintiff's
27 contentions is not an admission or agreement with Plaintiff about the meaning of any claim term,
28 but rather a reflection of the fact that the subject matter Plaintiff believes is claimed is present in the

1 prior art, or that the claims are otherwise invalid. These contentions are made in the alternative, and
2 should not be interpreted to rely upon, or in any way affect, the non-infringement arguments
3 Defendant may assert in this case. Defendant reserves the right to amend, supplement, or materially
4 modify these contentions as the case proceeds, including to identify additional prior art and/or
5 additional bases for invalidity under 35 U.S.C. §§ 101 and/or 112. Defendant also reserves the right
6 to amend, supplement, or materially modify its contentions based on any infringement and/or
7 additional claim construction positions that Plaintiff may take in this case.

8 Defendant also reserves the right to amend, supplement, or materially modify its contentions
9 in response to any claim construction or interpretation positions that Plaintiff may take or based on
10 any claim construction the Court may further adopt in this case. Defendant also reserves the right
11 to assert that a claim is indefinite, not enabled, or fails to meet the written description requirement
12 of 35 U.S.C. § 112 based on any claim construction or interpretation position Plaintiff may take in
13 this case or based on any claim construction the Court may further adopt in this case.

14 **B. No Admission**

15 Nothing disclosed herein is an admission or acknowledgement that any product accused of
16 infringement by Plaintiff in its Infringement Contentions (the “Accused Products”), or any of
17 Defendant’s other products or services, infringes any of the Asserted Claims.

18 Defendant further notes that Plaintiff appears to rely upon overly broad interpretations of the
19 Asserted Claims. At the same time, Plaintiff’s Infringement Contentions are in most places too
20 general and vague to discern Plaintiff’s infringement theories and how exactly Plaintiff contends
21 each Accused Product meets or practices each element of the Asserted Claims. As a result,
22 Defendant has been prejudiced in its ability to prepare these invalidity contentions. In addition,
23 Plaintiff’s Infringement Contentions, in many cases, continue to fail to put Defendant on notice of
24 Plaintiff’s interpretation of the Asserted Claims, further prejudicing Defendant’s ability to identify
25 relevant prior art.

26 In its Contentions, Plaintiff identified no specific claim elements as allegedly being present
27 under the doctrine of equivalents. In preparing these invalidity and eligibility contentions, Defendant
28 has relied on Plaintiff’s representation that it has no doctrine of equivalents theories for any specific

1 claim element, and any attempt by Plaintiff to present an untimely doctrine of equivalents argument
2 or allegation would be severely prejudicial to Defendant. To the extent Plaintiff attempts to or is
3 later permitted by the Court to amend its contentions to pursue any currently undisclosed doctrine
4 of equivalents theories for additional claim elements, Defendant expressly reserves the right to
5 supplement or amend these invalidity contentions to account for such amendments.

6 To the extent that any of the prior art references discloses the same functionality or feature
7 of any of the Accused Products, Defendant reserves the right to argue that said feature or
8 functionality does not practice any element of any of the Asserted Claims, and to argue, in the
9 alternative, that if said feature or functionality is found to practice any element of any of the Asserted
10 Claims, then the prior art reference demonstrates that the element is not novel, is obvious, and/or is
11 otherwise not patentable.

12 Attached hereto are representative claim charts that identify examples of where the elements
13 of the Asserted Claims may be found in the prior art and further identifying why those claims are
14 not novel or non-obvious. Defendant's claim charts cite particular teachings and disclosures of the
15 prior art as applied to features of the Asserted Claims. Providing citations to portions of the prior
16 art references should not be construed to mean that other portions of the prior art references are not
17 also relevant to the validity of the claims. Other passages of the identified prior art may additionally
18 disclose, either expressly or inherently, information relevant to the validity of one or more elements
19 of the Asserted Claims. And, persons having ordinary skill in the art generally may view an item
20 of prior art in the context of other publications, literature, products, and understanding. As such,
21 Defendant reserves the right to rely on uncited portions of the prior art references and on other
22 publications and expert testimony to establish the state of the art and to aid in understanding and
23 interpreting the cited portions, as providing context thereto, showing reasons to combine references
24 and as additional evidence that the prior art discloses a claim limitation. Defendant further reserves
25 the right to rely on uncited portions of the prior art references, other publications, and testimony to
26 establish bases for combinations of certain cited references that render the asserted claims obvious.
27 Defendant also reserves the right to rely on uncited portions of the prior art references that describe
28 substantially the same functionality or features that are described in cited portions. Further, for any

1 combination, Defendant reserves the right to rely on information generally known to those skilled
2 in the art and/or common sense.

3 Where Defendant cites a particular drawing or figure in the accompanying charts, the citation
4 encompasses the description of the drawing or figure, as well as any text associated with the drawing
5 or figure (even if the associated text is not itself expressly cited). Similarly, where Defendant cites
6 particular text concerning a drawing or figure in the accompanying charts, the citation encompasses
7 that drawing or figure (even if the associated drawing or figure is not expressly cited).

8 The references cited in the attached claim charts may disclose the limitations of the Asserted
9 Claims expressly and/or inherently. Defendant may rely on inherency to demonstrate the invalidity
10 of the Asserted Claims. Moreover, certain prior art references and solutions may inherently
11 anticipate certain features of the asserted claims as construed by Plaintiff. Defendant may rely on
12 cited or uncited portions of the prior art, other documents, and expert testimony to establish the
13 inherency of certain features of the prior art to invalidate the Asserted Claims. Defendant also may
14 rely on any reference identified in these invalidity contentions or any other reference to prove that
15 the references are enabling, to provide a reason to combine the references, as an indication of the
16 state of the art, or to explain the meaning of a term used in any reference.

17 The suggested obviousness combinations may be presented in conjunction with or in the
18 alternative to Defendant's contentions regarding anticipation. These obviousness combinations
19 should not be construed to suggest that any reference included in any combination is not anticipatory
20 in its own right. Further, to the extent that Plaintiff contends that any of the references identified
21 does not constitute prior art under 35 U.S.C. § 102, Defendant reserves the right to rely upon other
22 prior art references in the same patent family with substantially identical disclosures as evidence of
23 invalidity based on the same theories as those disclosed below.

24 Throughout the charts, Defendant provides exemplary citations indicating where the cited
25 references disclose subject matter recited in the preambles of the claims, without regard to whether
26 the preambles are limitations of the claims. Defendant does not take any position at this time
27 regarding whether or not the preambles of the Asserted Claims are limitations, and will address this
28 topic, if necessary, during the scheduled claim construction process.

1 With regard to the prior art that Defendant identifies as rendering the Asserted Claims
2 obvious, non-limiting examples of the reasons why the prior art renders the Asserted Claims
3 obvious, including exemplary identification of combinations of prior art showing obviousness, are
4 provided below. Generally, for all prior art combinations addressed herein, a reason or motivation
5 to combine the prior art includes that each prior art reference is relevant to data communications
6 and/or power delivery between devices, and one of ordinary skill in the art tasked with implementing
7 such a system would have been motivated to investigate the various existing systems, patents, and
8 other publications, including those identified herein, to address his or her particular needs. The
9 combinations and modifications of the prior art to invalidate the Asserted Claims would have arisen
10 from ordinary innovation, ordinary skill, or common sense and/or would have been obvious to try
11 or otherwise predictable in the field of data communications and/or power distribution between
12 devices.

13 A person having ordinary skill in the art would have had a reason to combine the prior art
14 based on the nature of the problem to be solved, the teachings of the prior art, and the knowledge of
15 persons having ordinary skill in the art. For example, each of the prior art references cited for a
16 given Asserted Patent is relevant to the subject of the Asserted Patent, including, for example, data
17 communications methods and systems and/or power delivery systems and methods, including
18 addressing problems such as communications with multilevel functionality, synchronizing and
19 ordering data communications, negotiating power delivery, supplying power, and arbitrating control
20 of a control bus. In addressing these problems, one of ordinary skill in the art would have been
21 motivated to combine the teachings of research in the field.

22 Additionally, some prior art references cite or discuss other prior art references, illustrating
23 the close technical relationship among the prior art. To the extent that any prior art reference refers
24 to or discusses another prior art reference, either expressly or inherently, it would have been obvious
25 to combine those references for at least that reason.

26 As shown in the accompanying claim charts, numerous books and research papers show that
27 methods were being used to streamline packetized data transmission and power distribution across
28 device-to-device connections. Design incentives and other market forces would have prompted

1 combinations and modifications of the prior art. Design needs and market pressures provide ample
2 reason to combine prior art elements in the manner combined. *See KSR Int'l Co. v. Teleflex Inc.*,
3 550 U.S. 398, 401 (2007). Defendant reserves the right to supplement these disclosures if such
4 evidence of design needs and/or market pressures become apparent through discovery in this case.

5 Moreover, the Asserted Claims are also obvious because they merely arrange old elements,
6 with each performing the same function that had been known, to perform and yield no more than
7 what one having ordinary skill in the art would expect from such an arrangement.

8 Because there were a finite number of predictable solutions in the art of packetized data
9 transmission and power distribution across device-to-device connections with regard to any claimed
10 feature, it would have been obvious to a person of ordinary skill in the art to pursue the known
11 options. A person skilled in this art would have been familiar with all the claim elements that the
12 patentees used to distinguish the prior art during prosecution. The identified prior art references
13 merely use those familiar elements for their primary or well-known purposes and in a manner well
14 within the ordinary level of skill in the art. Accordingly, common sense and the knowledge of one
15 having ordinary skill in the art render the Asserted Claims invalid.

16 **C. Reservation of Rights**

17 Prior art not currently included in this disclosure may become relevant. Defendant is
18 currently unaware of the extent, if any, to which Plaintiff will contend that limitations of the
19 Asserted Claims are not disclosed in the prior art identified by Defendant or contend that prior art
20 identified by Defendant does not qualify as prior art. Defendant reserves the right to identify other
21 references that would have made the addition of the allegedly missing limitation to the disclosed
22 device or method obvious or show that the allegedly missing limitation would have been known or
23 readily apparent to one of ordinary skill in the art at the time of the invention in light of the disclosure
24 of the prior art at issue or other prior art. Defendant further reserves the right to rely on any of the
25 references produced concurrently with these contentions in order to demonstrate the state of art at
26
27
28

1 the alleged times of invention and as evidence of the knowledge of one of ordinary skill in the art¹
2 in support of any motivations to modify or combine the charted prior art references with other
3 references or knowledge.

4 To the extent that Defendant charted a published patent application that later issued as a
5 patent that includes substantially the same disclosure, Defendant contends that the issued patent
6 discloses the claims for at least the same reasons explained in connection with the published
7 application and Defendant reserves its rights to rely on the issued patent to establish invalidity.
8 Similarly, to the extent that Defendant charted an issued patent the application for which was
9 previously published (and which contains substantially the same disclosure), Defendant contends
10 that the published application discloses the claims for at least the same reasons explained in
11 connection with the issued patent and Defendant reserves the right to rely on the published
12 application to establish invalidity (including to establish an earlier prior art date).

13 Consistent with Scheduling Order (Dkt. 93), and in view of the following, Defendant
14 reserves the right to supplement or otherwise amend these Invalidity Contentions. This is
15 particularly true for all claims for which Plaintiff has not provided adequate infringement
16 contentions.

17 Plaintiff may also be aware of additional prior art that is not known to Defendant. To the
18 extent Plaintiff produces additional prior art responsive to Defendant's discovery requests after these
19 contentions are served, Defendant may supplement its Invalidity Contentions with prior art
20 contained in such production once it has had a fair opportunity to review, analyze, and chart such
21 prior art. Defendant reserves the right to amend its Invalidity Contentions with any additional
22 potential prior art known by Plaintiff but not yet disclosed to Defendant.

23 Defendant provides these Invalidity Contentions only for the claims that are asserted by
24 Plaintiff, but reserves the right to seek invalidation of all claims in each of the Asserted Patents.

27 ¹ For purposes of these Invalidity Contentions, the term "general knowledge" includes the common
28 knowledge and common sense of a POSITA as well as the general knowledge of POSITAs in the
field.

1 In addition to the prior art that is set forth in the charts, additional prior art references are
2 being produced concurrently with these contentions. This additional art may be relevant to the
3 validity of the Asserted Claims depending on the positions taken by Plaintiff, and in any event may
4 be referenced to show the state and evolution of technology in these fields. Accordingly, for the
5 foregoing reasons, Defendant reserves the right to supplement these Invalidity Contentions with
6 specific claim charts and obviousness combinations with the additional references based on the
7 allegations and contentions asserted by Plaintiff as these allegations and corresponding
8 constructions of the claims become clear.

9 **D. Ongoing Investigation**

10 Defendant's investigation is ongoing, and Defendant expressly reserves the right to amend
11 its disclosures and document production to account for evidence uncovered as its investigation
12 continues. Such amendments include identifying and relying on additional references that may
13 result from Defendant's further search and analysis. Defendant reserves the right to supplement
14 these contentions in light of any additional prior art of which Plaintiff is aware and has not disclosed
15 to Defendant in discovery, or that might be subsequently disclosed by Plaintiff in response to
16 Defendant's discovery requests. An example of such prior art may include products sold and offered
17 for sale by Silicon Image and/or Lattice Semiconductor. Such products may have been offered for
18 sale or publicly shown in a manner that renders them prior art under 35 U.S.C. §§ 102 and 103.
19 Defendant reserves the right to supplement these contentions to rely on those and other products
20 from Plaintiff. Defendant's investigation into prior art systems and products is ongoing, and
21 Defendant reserves the right to amend these contentions to reflect the identification of and
22 information about such systems and products learned through discovery.

23 Defendant anticipates issuing subpoenas to third parties believed to have knowledge,
24 documentation and/or corroborating evidence concerning some of the prior art listed herein and/or
25 additional prior art. These third parties include, but are not limited to, the authors, employers of
26 inventors, assignees, or former or current employees of assignees or previous assignees, of the
27 references identified in these Invalidity Contentions. Defendant reserves the right to supplement
28

these contentions in light of any newly discovered information produced by these or other companies from which Defendant may seek discovery.

E. Incorporation by Reference

Defendant hereby incorporates by reference all arguments made in all IPR petitions challenging any Asserted Claim, including any decision instituting IPR that is filed after the date of these contentions, including the following IPR proceedings:

- IPR2024-01482;
- IPR2024-01480;
- IPR2024-01479;
- IPR2024-01478;
- IPR2024-01481;
- IPR2024-01429;
- IPR2024-01428.

Defendant further incorporates by reference the file history of each Asserted Patent and the prior art cited therein.

II. PRIORITY DATE OF THE ASSERTED PATENT CLAIMS

Plaintiff asserted the following priority dates for the Asserted Claims in its Infringement Contentions:

Asserted Patent	“Priority Date”
’905 Patent	“at least a priority date of November 22, 2000”
’307 Patent	“at least a priority date of June 12, 2003”
’798 Patent	“at least a priority date of November 22, 2000”
’265 Patent	“at least a priority date of March 31, 2011”
’712 Patent	“at least a priority date of December 11, 2008”
’520 Patent	“at least a priority date of January 4, 2008”
’231 Patent	“at least a priority date of January 4, 2008”
’103 Patent	“at least a priority date of April 14, 2014”

1 It is Plaintiff's burden to show entitlement to its asserted priority dates, and Defendant asserts
2 that Plaintiff has failed to meet that burden. Plaintiff has further failed in its obligations under PLR
3 3-1(f) to disclose "the priority date to which each asserted claim allegedly is entitled," at least
4 because it discloses ambiguous "priority dates" qualified with the phrase "at least as early as." See
5 UCT's Infringement Contentions at 7-8, served October 21, 2024. To the extent Plaintiff attempts
6 to rely on a different date from the specifically enumerated "Priority Dates" listed above (i.e., under
7 the "at least" language), Defendant reserves the right to supplement these contentions. Moreover,
8 because discovery has not yet been completed, Defendant has relied on Plaintiff's assertions for the
9 purposes of identifying relevant prior art in these Invalidity Contentions, but Defendant does not
10 concede that those priority dates are accurate. Defendant reserves the right to supplement these
11 Invalidity Contentions should discovery reveal that Plaintiff is not entitled to the asserted priority
12 date. In addition, Defendant also reserves the right to supplement these contentions should Plaintiff
13 subsequently attempt to amend its infringement contentions in any way, or to otherwise modify its
14 infringement allegations or seek to establish an earlier date of invention.

15 **III. THE ASSERTED CLAIMS OF THE '905 PATENT ARE INVALID**

16 **A. Identification of Prior Art**

17 Defendant contends that the prior art references charted in Exhibits A-1 through A-8,
18 discussed in any incorporated IPR petition, and/or described below anticipate and/or render obvious,
19 alone or in combination, one or more of the Asserted Claims of the '905 Patent.

20 **1. The '905 Patent Claims Are Anticipated by the Prior Art**

21 Some or all of the Asserted Claims of the '905 Patent are invalid as anticipated under 35
22 U.S.C. § 102 in view of each of the prior art references identified in the claim charts included in
23 Exhibit A-1 through A-8 and identified in each incorporated IPR petition, which identify specific
24 examples of where each limitation of the Asserted Claims is found in the prior art references. As
25 explained above, the cited portions of prior art references identified in the attached claim charts are
26 exemplary in nature and representative of the content and teaching of the prior art references, and
27 should be understood in the context of the reference as a whole and as they would be understood by
28 a person of ordinary skill in the art ("POSITA").

1 not clear whether this term should be understand to mean “included,” “represented by,”
2 “related to,” “translated to,” “occurs as a result of,” or some other meaning.

3 These claim terms, as apparently construed by Plaintiff, violate the definiteness requirements
4 of 35 U.S.C. § 112 at least because they fail to inform, with reasonable certainty, those skilled in
5 the art about the scope of the invention. Based on Defendant’s present understanding of Plaintiff’s
6 Infringement Contentions, at least these claim terms are indefinite because they are inconsistent
7 with and/or broader than the alleged invention disclosed in the specification and during prosecution,
8 and given Plaintiff’s apparent constructions of the claims, a POSITA at the time of the invention
9 would not understand what is claimed with reasonable certainty, even when the claims are read in
10 light of the specification and prosecution history.

11 **V. THE ASSERTED CLAIMS OF THE ’798 PATENT ARE INVALID**

12 **A. Identification of Prior Art**

13 Defendant contends that the prior art references charted in Exhibits C-1 through C-11,
14 discussed in any incorporated IPR petition, and/or described below anticipate and/or render obvious,
15 alone or in combination, one or more of the Asserted Claims of the ’798 Patent.

16 **1. The ’798 Patent Claims Are Anticipated by the Prior Art**

17 Some or all of the Asserted Claims of the ’798 Patent are invalid as anticipated under 35
18 U.S.C. § 102 in view of each of the prior art references identified in the claim charts included in
19 Exhibit C-1 through C-11 and/or identified in any incorporated IPR petition, which identify specific
20 examples of where each limitation of the Asserted Claims is found in the prior art references. As
21 explained above, the cited portions of prior art references identified in the attached claim charts are
22 exemplary in nature and representative of the content and teaching of the prior art references, and
23 should be understood in the context of the reference as a whole and as they would be understood by
24 a person of ordinary skill in the art.

25 In addition to references identified in any incorporated IPR petition, Defendant identifies the
26 following references as anticipating one or more of the Asserted Claims of the ’798 Patent under 35
27 U.S.C. § 102. The table of anticipating references below is exemplary, and it does not constitute an
28 admission that any reference not listed below does not also anticipate the claims of the ’798 Patent.

Further, Defendant contends that any prior art reference in the attached charts that is charted for each limitation of any given claim, anticipates that claim, regardless of whether that prior art reference is listed in the following tables.

a) Prior Art Patents and Applications

Number	Published/Issued Date	Short Name	Exhibit Number
U.S. Patent No. 5,592,487	Filed: Mar. 20, 1995 Issued: Jan. 7, 1997	Knecht	C-1
U.S. Patent No. 5,686,965	Filed: Sep. 18, 1995 Issued: Nov. 11, 1997	Auld	C-2
European Patent Application No. 1,137,278	Filed: Mar. 20, 2001 Pub'd: Sep. 26, 2001	Kotani	C-3
U.S. Patent No. 5,790,607	Filed: Nov. 28, 1995 Issued: Aug. 4, 1998	Burke	C-4
U.S. Patent No. 6,609,167	Filed: Mar. 15, 2000 Issued: Aug. 19, 2003	Bastiani	C-5
U.S. Patent No. 7,251,256	Filed: May 18, 2000 Issued: July 31, 2007	Barry	C-6
WIPO Patent Pub. No. 1999/023776	Filed: Oct. 31, 1997 Pub'd: May 14, 1999	Hui	C-7
U.S. Patent No. 5,917,830	Filed: Oct. 18, 1996 Issued: June 29, 1999	Chen	C-8
U.S. Patent No. 6,430,230	Filed: June 30, 1999 Issued: Aug. 6, 2002	Cunningham	
U.S. Patent No. 5,801,781	Filed: June 19, 1996 Issued: Sep. 1, 1998	Hiroshima	
WIPO Patent Pub. No. 2000/011831	Filed: Aug. 17, 1999 Pub'd: Mar. 2, 2000	Haartsen	
U.S. Patent No. 5,974,464	Filed: Sep. 30, 1996 Issued: Oct. 26, 1999	Shin	
U.S. Patent No. 6,498,568	Filed: Sep. 24, 1997 Issued: Dec. 24, 2002	Austin	
WIPO Patent Pub. No. 1992/09162	Pub'd: May 29, 1992	Hornak	C-10

b) Prior Art Publications

Author or Publisher	Title	Publication/Use Date	Short Name	Exhibit Number
ISO (International Organization for Standardization)	Information Technology – Generic coding of moving pictures and associated audio information: Systems	April 1, 1996	MPEG-2, Part 1	See IPR 2024-01481

Defendant additionally identifies and relies on any system, product, or public knowledge or use that embodies or otherwise incorporates any of the prior art patents and publications listed above. Defendant reserves the right to identify and rely on systems that represent different versions or are otherwise related variations of the identified systems. Defendant also incorporates by reference any and all other prior art systems identified in any other case brought by Plaintiff and/or involving the Asserted Patents.

To the extent any item of prior art cited above is deemed not to disclose, explicitly or inherently, any limitation of an asserted claim of the '798 Patent, Defendant contends that any difference between that prior art and the corresponding patent claim would have been either inherent in the art or obvious to a person of ordinary skill in the art, even if Defendant has not specifically denoted that the art is to be combined with the knowledge of a person of ordinary skill in the art.

c) Prior Art Systems

On information and belief, systems implementing the protocols identified below were in public use, on sale, or otherwise available to the public before the priority date of the '798 Patent. Defendant's further investigation and/or subsequent discovery from Plaintiff or third parties with knowledge regarding prior art systems may reveal additional relevant prior art systems and/or further information regarding the systems identified below, or other relevant prior art systems. Defendant reserves the right to supplement these Invalidity Contentions, based on subsequent investigation and discovery, including from third parties.

System	Available At Least As Of	Short Name	Exhibit Number
Gigabit Ethernet employing IEEE 802.3z (1998)	1998	802.3z-1998	C-9

1	HP G-Link Chipset implementing Serial-HIPPI	1992	G-Link	C-11
---	---	------	--------	------

2 Defendant additionally identifies and relies on any system, product, or public knowledge or
3 use that embodies or otherwise incorporates any of the prior art patents and publications listed
4 above. Defendant reserves the right to identify and rely on systems that represent different versions
5 or are otherwise related variations of the identified systems. Defendant also incorporates by
6 reference any and all other prior art systems identified in any other case brought by Plaintiff and/or
7 involving the Asserted Patents.

8 To the extent any item of prior art cited above is deemed not to disclose, explicitly or
9 inherently, any limitation of an asserted claim of the '798 Patent, Defendant contends that any
10 difference between that prior art and the corresponding patent claim would have been either inherent
11 in the art or obvious to a person of ordinary skill in the art, even if Defendant has not specifically
12 denoted that the art is to be combined with the knowledge of a person of ordinary skill in the art.

13 2. The '798 Patent Claims Are Obvious Over the Prior Art

14 To the extent Plaintiff argues that any limitation of the asserted claims of the '798 Patent are
15 not explicitly disclosed by an item of prior art identified above, in Exhibits C-1 through C-11, or
16 discussed in any incorporated IPR petition, any purported differences are such that the claimed
17 subject matter as a whole would have been obvious to one skilled in the art at the time of the alleged
18 invention, in view of the state of the art and knowledge of those skilled in the art. The item of prior
19 art would, therefore, render the relevant claims invalid as obvious under 35 U.S.C. § 103.

20 It would have been obvious to combine any of the prior art in Exhibits C-1 through C-11
21 with each other, at least because all of these references relate to synchronizing serial
22 communication(s) between two devices.

23 For example, U.S. Patent No. 5,592,487 ("Knecht") teaches that "[a] communications
24 protocol is established with data packets, control packets, checksum packets, and synchronization
25 symbols. One of the control packets is transmitted after a synchronization symbol is transmitted.
26 One of the synchronization symbols, data packets, or checksum packets is transmitted after one of
27 the control packets is transmitted. One of the synchronization symbols is transmitted after one of
28 the checksum packets is transmitted, and one of the synchronization symbols or another of the data

1 packets is transmitted after one of the data packets is transmitted.” Knecht, 2:3–13. Knecht further
2 teaches that the “[c]ommunication protocol is established and maintained with four signals, namely:
3 sync symbol, data packet, control packet, and checksum packet. In the preferred embodiment,
4 control, data and checksum packets comprise 16 bits of information before encoding and 20 bits
5 after encoding.” *Id.* at 3:28–32.

6 U.S. Patent No. 5,686,965 (“Auld”) teaches that “[a]n MPEG ISO system bitstream typically
7 includes a sequence of packs 90. A typical pack 90, in turn, includes a sequence pack layer headers
8 92, a system header packet 94, a sequence of packets 96, 98, and 100, and an end code 102. The
9 pack layer header 92 contains a pack start code, or sync code, used for synchronization purposes,
10 and a system clock value. The system header packet 94 contains a variety of information commonly
11 referred to as system housekeeping data, and a system stream ID which is used to differentiate
12 among multiple system streams. A sequence of one or more packets typically contains either
13 encoded audio or encoded video stream data. The final element of a pack 90 is an end code 102. It
14 should be noted that any one packet 96, 98, 100 may carry either audio or video data, but not both
15 simultaneously. With respect to a packet containing video data, typically a video packet header and
16 one or more video frames is contained within the packet.” Auld, 10:20–36.

17 European Patent Application No. 1,137,278 (“Kotani”) teaches “the structure of a packet
18 generated by a VCR. A sync pattern of about 2 bytes, which indicates the head of the packet, and
19 an ID of about 3 bytes, which includes address information on data, are added to the packet.” Kotani,
20 ¶ 4. Kotani further teaches that “packets for commercialized household digital VCRs can be
21 classified into two types, i.e., a 12-byte sub-code packet including a 2-byte sync pattern, and a 90-
22 byte AV data packet including a 2-byte sync pattern.” *Id.* ¶ 5. Kotani further teaches that “[w]hen
23 there are two types of time-division multiplexed HD data (for example, when there are a sub-code
24 and AV data), it is necessary to add different sync patterns to the two types of data in order to
25 discriminate between each type of data.” *Id.* ¶ 35.

26 U.S. Patent No. 5,790,607 (“Burke”) teaches that “for asynchronous data transmission and
27 reception, it is highly desirable to organize the data into predetermined formats or packets for
28 accurate and reliable data reception. In addition to providing the actual data, which is highly variable

1 and essentially random, the data packet should also provide appropriate timing and synchronizing
2 information needed for accurate data reception, without creating excessive overhead and without
3 decreasing data throughput. In addition, for asynchronous data transmission and reception, an
4 apparatus and method is also necessary to quickly, reliably and accurately detect such timing and
5 synchronizing information prior to the reception of any actual data, without excessive or
6 unnecessary processor complexity, speed, and cost.” Burke, 2:38–51.

7 U.S. Patent No. 6,609,167 (“Bastiani”) teaches that “a method for communicating serial data
8 between a host computer and a peripheral device is disclosed. The packet for communication
9 between the host computer and the peripheral device includes: (a) a synchronization field; (b) a
10 packet type (PT) field following the synchronization field, the packet type field defining an error
11 correction type for the packet; (c) a byte count (BC) field for defining a length of data for the packet;
12 (d) a data type (DT) field for defining whether the data is one of link control, device control, and
13 device data; (e) a data field; and (f) a CRC field for error checking the data in the data field.”
14 Bastiani, 2:29–39.

15 U.S. Patent No. 7,251,256 (“Barry”) teaches that “[a]n LMAC-compatible link has the
16 following characteristics: (a) the two devices at the endpoints of the link are not synchronized, e.g.
17 the link is asynchronous; (b) the data transported on the link is in the form of packets, cells, or
18 frames; and (c) the MAC layer encoding of the data is such that devices called LMACs can insert
19 synchronization symbols anywhere into the data stream at the ingress end of the link and remove
20 them at the egress end so that the synchronization symbols are not visible to higher-layer
21 components that process data packets.” Barry, 3:61–4:3. Barry further teaches that “[i]n the LMAC
22 protocol, the MAC layer encoding of the data is such that devices called LMACs can insert
23 synchronization symbols anywhere into the data stream at the ingress end of an LMAC-compatible
24 link and remove them at the egress end so that the synchronization symbols are not visible to higher-
25 layer components that process data packets.” Barry, 4:65–5:4.

26 WIPO Patent Publication No. 1999/023776 (“Hui”) teaches that “[i]n data retrieval or
27 communication systems, digital data or digitized multimedia information which can be, potentially,
28 large or variable in size is normally packetized before storage or transmission. This process breaks

1 the data into smaller packets where each packet contains one or more detectable synchronization
2 patterns called sync-word, necessary information about the packet, and a payload of a smaller
3 partition of the original data. The information about the packet may consist of a packet identification
4 (ID), size of the payload contained in the current packet, timing information, and any necessary side
5 information.” Hui, 1:13–19.

6 U.S. Patent No. 5,917,830 (“Chen”) teaches “the insertion of digital video messages such as
7 commercials into a pre-existing compressed packetized data stream.” Chen, 1:7–9. Chen further
8 teaches that “a transport stream, or transport multiplex is made up of a contiguous set of fixed length
9 packets. Each packet is 188 total bytes in length, with the first four of those bytes being defined as
10 the packet header.” *Id.* at 14:47–50. Chen further discloses that “[v]arious timing and identification
11 information is provided in different portions of the transport stream. These include a packet
12 identifier (PID) found in the transport header of each transport packet” and “a program clock
13 reference (PCR) which effectively represents a sample of the system time clock (STC) time base
14 that underlies the service composed of the PIDs referenced in the service map.” *Id.* at 14:57–59,
15 15:5–8.

16 U.S. Patent No. 6,430,230 (“Cunningham”) teaches “a method of encoding payload bits for
17 transmission over a communications link such that the running disparity of the transmitted symbol
18 stream is constrained within a predetermined limit, and said payload bits comprising either data or
19 control information.” Cunningham, 1:58–63. Cunningham further teaches that “a data frame
20 intended to be communicated over an optical fibre channel 20 typically comprises a binary digital
21 message (payload) to be transferred between stations on the network; this message may comprise
22 user data, or control information for managing operation of the network and its components.” *Id.* at
23 2:65–3:3. Cunningham further teaches that “[c]ontrol information is conveyed by using selected n-
24 bit data block values, together with a distinctive value of 10 for the pair of label bits a and b to
25 indicate that the associated data block represents control information rather than actual data.” *Id.* at
26 4:23–27.

27 U.S. Patent No. 5,801,781 (“Hiroshima”) teaches “the details of the MPEG2-TS. Each
28 packet of the MPEG2-TS has a prefix length of 188 bytes. The header 194 consists of four bytes

1 and the payload 196 consists of 184 bytes. The header 194 is constructed by sync bytes 198, an error
2 indicator 200, a unit start indicator 202, a transport priority 204, a packet ID 206, a scramble control
3 208, an adaptation field control 210, a continuity counter 212, and an adaptation field 214.”
4 Hiroshima, 10:51–58. Hiroshima further teaches “a process for forming the MPEG2-TS from the
5 PES packet 82 by a multiplexing process.” *Id.* at 11:13–14.

6 WIPO Patent Publication No. 2000/011831 (“Haartsen”) teaches “a multi-part digital
7 preamble for use in transmitting digital data packets. Advantageously, a preamble according to the
8 invention substantially reduces the overall length of DC-free sequences that are transmitted with
9 each data packet in a digital FM system.” Haartsen, 5:3–6. More specifically, Haartsen teaches that
10 “a digital preamble can be strategically divided into three or more parts so that synchronization and
11 other information can be included within the preamble while still enabling accurate DC estimation.”
12 *Id.* at 12:13–16.

13 U.S. Patent No. 5,974,464 (“Shin”) teaches “a high-speed digital video signal transmission
14 system in which high-speed video transmission is effected between a host computer system and a
15 system display.” Shin, 4:56–59. Shin further teaches that “various synchronization and other
16 special characters are defined using the high-transition and low-transition sets of out-of-band
17 characters. . . . The difference in the number of transitions between in-band and out-of-band
18 characters allows selected out-of-band characters to serve as control characters, and to be readily
19 distinguished from in-band characters within the transmitted data stream.” *Id.* at 9:64–10:8.

20 U.S. Patent No. 6,498,568 (“Austin”) teaches a “system [that] uses up to 5 pre-defined data
21 packet sizes and up to 16 user defined data message types within each packet. Each packet has
22 Forward Error Connection (FEC) and bit interleaving applied to it on transmission. Error checking
23 is performed using a standard CRC. The system uses standard run in and byte synchronisation
24 techniques. The byte synchronisation word used is 32 bits long and there are 5 different words used,
25 one to identify each of the different packet types.” Austin, 4:37–45.

26 ISO/IEC 13818-1, first edition 1996-04-15 (“MPEG-2, Part 1”) is a standard specification
27 that was published in April 1996 by the International Organization for Standardization (ISO).
28 MPEG-2, Part 1 “addresses the combining of one or more elementary streams of video and audio,

1 as well as other data, into single or multiple streams which are suitable for storage or transmission.”
2 MPEG-2, Part 1 at viii. MPEG-2, Part 1 describes a “Program Stream” which “can be decoded with
3 synchronization.” *Id.* at ix. The standard further includes a specification of the “Program Stream,”
4 including its syntax and semantics. *Id.* at Section 2.5.3. Video and audio data compressed according
5 to the specification of the “Program Stream” “are packetized to produce PES [packetized elementary
6 stream] packets.” *Id.* at viii.

7 IEEE Std 802.3 (“802.3z-1998”) is a working group of standard specifications for Ethernet
8 that describes a method of packet-based physical communication in a local area network maintained
9 by the Institute of Electrical and Electronics Engineers (IEEE). Its 1998 Edition (“802.3z-1998”)
10 was first implemented in systems in the United States by at least 1998. 802.3-1998 accordingly is
11 prior art to the ’798 patent under pre-AIA 35 U.S.C. 102(a)(1), and for at least the reasons described
12 below, invalidates the asserted claim 11, 19, 22, 25, 28 of the ’798 patent. Additionally, devices
13 employing the IEEE Std 802.3 were offered for sale before the ’798 patent. For example, the
14 1000BASE-T Gigabit Ethernet Connection employing the IEEE Std 802.3 was included in Apple’s
15 Power Mac G4 no later than October 2000.

16 G-Link HDMP-1000 chip set (G-Link) is designed for network data transmission developed
17 by HP, first sold or offered for sale in the United States by at least October 1992, implementing
18 High-Performance Parallel Interface (HIPPI) and IEEE 802.2-6 protocols. Aspects of G-Link are
19 described in various publications, including, but not limited to:

- 20 • James R. Hoffman, “HIPPI-6400 Technology Dissemination,” SPIE Photonics East (Nov.
21 18, 1992) (“Hoffman”);
- 22 • “The G-Link Chipset for Gigabit-Rate Data Communication,” pp. 103-116 Hewlett-Packard
23 Journal (Oct. 1992) (“HPJ”);
- 24 • Walker et al., “A Two-Chip 1.5 G-Bd Serial Link Interface,” IEEE Journal of Solid State
25 Circuits (Dec. 1992) (“Walker”);
- 26 • Peter Steenkiste, “Gigabit Nectar: Architecture and performance” (Sept. 1995) (“Gigabit”);
- 27 • Tolmie et al., “HIPPI/Serial-HIPPI,” (1992);
- 28 • Katz, “High Performance Network and Channel-Based Storage” (1991);

- “IEEE Standard for Scalable Coherent Interface (SCI),” IEEE Std 1596-1992 (Mar. 19, 1992) (“SCI”)

Accordingly, because each of the identified references are directed to similar devices in the similar field of endeavor, a person of ordinary skill in the art would be motivated to combine the teachings of two or more references together because the references disclose similar systems that solve the same problems in similar ways. Thus, it would have been obvious to one of ordinary skill in the art to combine any of the foregoing references and one of ordinary skill would have expected such a combination to be successful.

For example, to the extent that it is determined that any of the above references do not disclose a “different synchronization symbol” for each “packet type,” it would further have been obvious to modify the teachings of that reference to include this feature in view of the following teachings: *See, e.g.*, Auld, 10:53–61 (“The packet data field of the packet layer 96 may contain emulations of the sync codes of other bitstream data, such as those comprising audio packets. The system synchronizer 62, however, analyzes the bitstream to determine whether a legitimate video packet sync code has been located, as at 126. If either the video sync code has not been located or a sync code emulation is located, the system synchronizer 62 will continue to predict the occurrence of a legitimate video sync code in the bitstream.”); Kotani, ¶ 35 (“When there are two types of time-division multiplexed HD data (for example, when there are a sub-code and AV data), it is necessary to add different sync patterns to the two types of data in order to discriminate between each type of data.”); Bastiani, 24:58–63 (“A packet type (PT) immediately follows the SYNC field of every ASP packet. A PT consists of a four bit packet type field followed by a four-bit check field as shown in FIG. 18. The PT indicates the type of packet and, by inference, the format of the packet and the type of error detection applied to the packet.”); Hui, 13:6–7 (“The sync-word is defined based on applications or packet types, and is normally a digital pattern not easily found within the payload.”); Chen, 14:56–15:3 (“Various timing and identification information is provided in different portions of the transport stream. These include a packet identifier (PID) found in the transport header of each transport packet to provide a reference number for identifying the transport packets carrying a specific service component. . . . Each type of packet will have a different PID identifying the packet

1 type.”); Cunningham, 4:23–27 (“Control information is conveyed by using selected n-bit data block
2 values, together with a distinctive value of 10 for the pair of label bits a and b to indicate that the
3 associated data block represents control information rather than actual data.”); Hiroshima, 11:16–
4 22 (“The TS header 194 is formed in a manner such that the parameters such as error indicator 200
5 indicative of the presence or absence of the bit error in the packet, unit start indicator 202 indicating
6 whether a new PES packet appears in the TS payload or not, packet ID 206 indicating whether the
7 packet in the payload is video or audio”); Haartsen, 12:19–13:2 (“Thereafter, a synchronization
8 part of the preamble, which is not necessarily DC free, but which is protected by code to prevent
9 detection errors, can be used not only to synchronize detection of a data packet, but also to carry
10 timing and/or other useful information (e.g., channel identification, destination address, time of day,
11 user identification, etc.).”); Austin, 4:42–45 (“The byte synchronisation word used is 32 _bits long
12 and there are 5_different words used, one to identify each of the different packet types.”). Therefore,
13 a person of ordinary skill in the art would be motivated to make such a combination.

14 To the extent that it is determined that any of the above references do not disclose
15 “synchronization information,” it would further have been obvious to modify the teachings of that
16 reference to include this feature in view of the following teachings: *See, e.g.*, Burke, 3:41–48 (“In
17 the preferred embodiment of the invention, data for transmission is organized into packets beginning
18 with a ‘preamble’ 101 comprising a known, predetermined bit sequence, followed by a transition
19 102 comprising an indicator that data will follow, such as synchronizing information (a SYNC word
20 preferably [sic] having a high autocorrelation), followed by a data stream 103 having a variable
21 length”); Hui, 13:6–9 (“The sync-word is defined based on applications or packet types, and is
22 normally a digital pattern not easily found within the payload. After detection of the sync-word, the
23 process checks at step 503 for the correct packet ID which may be used to differentiate packets from
24 different data sources, types, or versions.”); Haartsen, 13:16–19 (“The synchronization word 520,
25 which can include timing and/or other information and which is coded to protect against detection
26 errors, is in turn followed by a trailing DC-free sequence 530.”). Therefore, a person of ordinary
27 skill in the art would be motivated to make such a combination.

28

1 To the extent that it is determined that any of the above references do not disclose packets
2 containing “in-band symbols” that are “transition optimized” and synchronization symbols that are
3 “out-of-band symbols” and “not transition optimized,” it would further have been obvious to modify
4 the teachings of that reference to include this feature in view of the following teachings: *See, e.g.,*
5 Auld, 11:14–19 (“A typical sync word may be represented by a 24-bit or 32-bit word made up of
6 consecutive zeros and ending with a ‘1’ bit. Alignment involves the recognition of word or byte
7 boundaries within a system bitstream in accordance with the syntax of the selected compression
8 algorithm.”); Knecht, 4:46–49 (“The 5-bit codes, with the exception of the sync symbol, are
9 formatted in a 40/60 percent ratio of 0 and 1 bits, with either end of the 5-bit code being made up of
10 no more than two logical ‘0’ or ‘1’ bits.”), 5:37–38 (“Two representations of sync symbols namely
11 00000 and 11111 are necessary.”); Kotani, ¶ 6 (“Ten 1’s or ten 0’s are repeated in the respective
12 [sync] patterns. In this way, erroneous detection of data bits and sync patterns is prevented since the
13 number of consecutive 0’s is restricted in the data bits.”); Shin, 9:66–10:8 (“Each character
14 corresponding to one of these ‘out-of-band’ characters associated with the high-transition set
15 includes less than four logical transitions, and each of the out-of-band characters associated with the
16 low-transition set of code characters includes more than four logical transitions. The difference in
17 the number of transitions between in-band and out-of-band characters allows selected out-of-band
18 characters to serve as control characters, and to be readily distinguished from in-band characters
19 within the transmitted data stream.”). Therefore, a person of ordinary skill in the art would be
20 motivated to make such a combination.

21 To the extent that it is determined that any of the above references do not disclose packet
22 types that include a “control packet,” it would further have been obvious to modify the teachings of
23 that reference to include this feature in view of the following teachings: *See, e.g.,* Knecht, 3:28–32
24 (“Communication protocol is established and maintained with four signals, namely: sync symbol,
25 data packet, control packet, and checksum packet. In the preferred embodiment, control, data and
26 checksum packets comprise 16 bits of information before encoding and 20 bits after encoding.”);
27 Kotani, ¶ 5 (“For example, packets for commercialized household digital VCRs can be classified
28 into two types, i.e., a 12-byte sub-code packet including a 2-byte sync pattern, and a 90-byte AV

1 data packet including a 2-byte sync pattern.”); Bastiani, 26:48–50 (“Data type 1 will be used for
2 control and status of the device, and Data type 2 will be used to transfer data to the device.”); Chen,
3 14:65–15:3 (“In the case of video, audio or isochronous control data, the stream of packets labeled
4 with a single PID represents a single video, audio or isochronous data service elementary stream,
5 respectively. Each type of packet will have a different PID identifying the packet type.”);
6 Cunningham, 4:23–27 (“Control information is conveyed by using selected n-bit data block values,
7 together with a distinctive value of 10 for the pair of label bits a and b to indicate that the associated
8 data block represents control information rather than actual data.”). Therefore, a person of ordinary
9 skill in the art would be motivated to make such a combination.

10 Additionally, to the extent that it is determined that any of the above references do not
11 disclose transmitting symbols to a “switch network,” it would further have been obvious to modify
12 the teachings of that reference to include this feature in view of the following teachings: *See, e.g.,*
13 Burke, 1:31–36 (“[D]igital data may be transmitted both in the downstream direction, from the head
14 end (connected to a network) to the individual user (subscriber access unit), and in the upstream
15 direction, from the individual user to the head end (and to the network).”); Bastiani, 6:3–6 (“ASP is
16 a star network that supports data exchange between a host computer and storage devices such as
17 disks, tapes and CD’s. The attached peripherals share ASP bandwidth through a host scheduled
18 polling based protocol.”), 6:20–23 (“In this example, the host adapter 104 is configured to include
19 ASP circuitry 104 a for processing ASP packet data and communicating between the host computer
20 102 and the ASP device 106.”); Barry, 12:17–23 (“The output of packet processor 54 enters the
21 switch fabric 50. One suitable packet switch is the MMC Networks model nP5400 Packet Switch
22 Module. In one embodiment, four such switches are connected in each switching card for faster
23 throughput. This embodiment is used on the switching card. The capabilities of this switch are not
24 critical to this specification.”); Chen, 4:63–67 (“A compressed digital packetized data stream which
25 includes the inserted message combined with the main program is output from the DAIM 130 to an
26 ATM transmitter (Tx) interface 126, and then provided to an access network (e.g., cable system
27 headend) 140.”); Cunningham, 2:65–3:1 (“Referring to FIG. 1, a data frame intended to be
28 communicated over an optical fibre channel 20 typically comprises a binary digital message

1 (payload) to be transferred between stations on the network. . . .”); Hiroshima, 5:40–43 (“The video-
2 on-demand system corresponds to MPEG2 having a transmission rate of 6.144 Mbps and transmits
3 data by a prefix bit rate of the MPEG2-TS through the network 22.”). Therefore, a person of
4 ordinary skill in the art would be motivated to make such a combination.

5 To the extent that it is determined that any of the above references do not disclose a
6 “synchronization primitive” made up of “a sequence of bit-string synchronization symbols,” it
7 would further have been obvious to modify the teachings of that reference to include this feature in
8 view of the following teachings: *See, e.g.*, Bastiani, 24:43–46 (“A minimum of two and maximum
9 of four SYNC characters must be sent to assure byte synchronization. It is used by the receive
10 circuitry to align incoming data bytes and is defined to be eight bits in length.”); Cunningham, 7:5–
11 8 (“There are many potential ways to achieve code block synchronization (i.e. correct alignment in
12 the receiver of the label and payload fields). For example, one technique would be to transmit a
13 ‘comma sequence’ during start-up.”); Shin, 15:20–26 (“During each preamble period, three
14 repetitions of the same synchronization character are produced by the encoder 18. As is described
15 herein, by processing the 21 bits most recently produced by the encoder 14, synchronization module
16 114 is able to detect at least one of the three repetitions of the synchronization character transmitted
17 during a given preamble period.”). Therefore, a person of ordinary skill in the art would be
18 motivated to make such a combination.

19 A person skilled in the art would have been motivated to combine any of the above
20 references. In particular, the combinations identified herein would merely have been (a) a
21 combination of prior art elements according to known methods to yield predictable results; (b) a
22 simple substitution of one known element for another to obtain predictable results; (c) a use of
23 known technique to improve similar devices or methods in the same way; (d) application of a known
24 technique to a known device or method ready for improvement to yield predictable results; (e)
25 obvious to try; and/or (f) known work in one field of endeavor prompting variations of it for use in
26 either the same field or a different one based on design incentives or other market forces since the
27 variations are predictable to one of ordinary skill in the art. Moreover, a person of ordinary skill in
28

1 the art had a reason to combine these references because—at a minimum—they all pertain to serial
2 communication between a transmitting device and a receiving device using packets.

3 As discussed above, the motivation to combine the teachings of the patents and other prior
4 art referenced above is found in the art as specified in the charts and in: (1) the nature of the problem
5 being solved, (2) the express, implied and inherent teachings of the prior art, (3) the knowledge of
6 persons of ordinary skill in the art; and (4) the fact that the prior art is generally directed towards
7 communication between a transmitting device and a receiving device using packets, and (5) the
8 predictable results obtained in combining the different elements of the prior art. The Supreme Court
9 has held that a showing of a motivation to combine within the references themselves is not required
10 to prove obviousness. See *KSR*, 550 U.S. at 401 (“Any need or problem known in the field of
11 endeavor at the time of invention may provide the motivation to combine.”).

12 For example, a person skilled in the art would have been motivated to combine Auld with
13 Haartsen, Kotani, Shin, and/or additional references discussed above and disclosed in Exhibits C-1
14 through C-11 for the reasons provided above. Moreover, the combination of references is no more
15 than a combination of systems all directed to serial transmission of data using packets. Additionally,
16 the combination reflects no more than a simple substitution of known elements for one another to
17 obtain predictable results. For example, a simple substitution of Haartsen’s digital preamble for
18 Auld’s sync code would achieve the predictable result of incorporating synchronization information
19 into a synchronization symbol. As another example, a person of ordinary skill in the art would have
20 been motivated to implement Kotani’s disclosure of a unique synchronization code for subcode (i.e.,
21 control) packets into Auld, which discloses a need to distinguish between packet types.
22 Additionally, the straightforward addition of Shin’s preamble sequence to Auld would enhance
23 the capability of Auld’s system to synchronize and align the transmitted data packets.

24 A person skilled in the art would have been motivated to combine Bastiani with Haartsen,
25 Kotani, Knecht, Shin, and/or additional references discussed above and disclosed in Exhibits C-1
26 through C-11 for the reasons provided above. Moreover, the combination of references is no more
27 than a combination of systems all directed to serial transmission of data using packets. Additionally,
28 the combination reflects no more than a simple substitution of known elements for one another to

1 obtain predictable results. For example, a simple substitution of Haartsen's digital preamble for the
2 non-data fields of the packets in Bastiani would achieve the predictable result of incorporating
3 synchronization information into a synchronization symbol. As another example, a person of
4 ordinary skill in the art would have been motivated to implement the out-of-band synchronization
5 symbols in Kotani and Knecht in order to avoid the erroneous detection of packets and
6 synchronization symbols. Additionally, the straightforward addition of Shin's preamble sequence
7 to Bastiani would enhance the capability of Bastiani's system to synchronize and align the
8 transmitted data packets.

9 Thus, in addition to the combinations discussed above, exemplary combinations which
10 render obvious some or all of the Asserted Claims of the '798 Patent are listed in **Appendix A**, Sec.
11 III.

12 **B. Invalidity of the '798 Patent Based on 35 U.S.C. § 112, ¶ 1**

13 Based on Defendant's present understanding of Plaintiff's asserted scope of the claims of
14 the '798 Patent, to the extent such positions can be understood from Plaintiff's Infringement
15 Contentions, Defendant asserts that the Asserted Claims of the '798 Patent are invalid for failure to
16 provide a written description that enables the full scope of the Asserted Claims based on at least the
17 following claim terms, phrases, or limitations:

- 18 • **"synchronization symbol" (claims 11, 19, 22, 28):** To the extent UCT contends that
19 "synchronization symbol" has the same meaning as a "synchronization primitive" or
20 "synchronization information," this lacks written description support in the specification. In
21 addition, the specification of the '798 Patent does not describe a "synchronization primitive"
22 or "synchronization information," so those terms lack written description support and are
23 not enabled by the disclosure in the '798 Patent for at least this additional reason.
- 24 • **"synchronization primitive" (claims 11, 19, 28):** *See above.*
- 25 • **"synchronization information" (claim 19):** *See above.*
- 26 • **"bit-string synchronization symbols" (claim 28):** *See above.*

27 Based on Defendant's present understanding of Plaintiff's Infringement Contentions and
28 apparent interpretation of the scope of the Asserted Claims, at least one or more of these claim terms,

1 phrases, and limitations lack written description support, as they are not identified, described, or
2 explained in the specification of the Asserted Patent, and do not enable a person of ordinary skill in
3 the art to practice the claimed invention without undue experimentation or to distinguish all
4 operative embodiments from all nonoperative embodiments without undue experimentation.

5 **C. Invalidity of the '798 Patent Based on 35 U.S.C. § 112, ¶ 2**

6 Based on Defendant's present understanding of Plaintiff's asserted scope of the claims of
7 the '798 Patent, to the extent such positions can be understood from Plaintiff's Infringement
8 Contentions, the Asserted Claims of the '798 Patent are invalid at least because they fail to inform,
9 with reasonable certainty, those skilled in the art about the scope of the invention. At least the
10 following claim language is indefinite:

- 11 • **“the synchronization symbol permitting properly alignment upon receiving with a**
12 **synchronization primitive to be correctly aligned on a symbol boundary (claim 11):** The
13 scope of this limitation is not clear, as the claim language does not describe what element
14 performs the “receiving” function and does not clearly identify what is “received.”
- 15 • **“synchronization information” (claim 19):** The '798 Patent specification does not describe
16 “synchronization information” or how the synchronization symbol provides such
17 information or how such information is different from the “synchronization symbol” or the
18 “synchronization primitive.”
- 19 • **“the transmitted synchronization symbol . . . permitting an external receiving node to**
20 **properly align with a synchronization primitive to be correctly aligned on a symbol**
21 **boundary” (claim 19):** It is not clear how a “external receiving *node*” aligns with the
22 “synchronization symbol,” “synchronization primitive,” or both, rendering this limitation
23 indefinite.

24 These claim terms, phrases, limitations, and all dependent claims thereof, as apparently
25 construed by Plaintiff, violate the definiteness requirements of 35 U.S.C. § 112 at least because they
26 fail to inform, with reasonable certainty, those skilled in the art about the scope of the invention.

27 Based on Defendant's present understanding of Plaintiff's Infringement Contentions, one or
28 more of these claim terms, phrases, and limitations are indefinite because they are inconsistent with

1 and broader than the alleged invention disclosed in the specification and during prosecution, and
2 given Plaintiff's apparent constructions of the claims, a person of ordinary skill in the art at the time
3 of the invention would not understand what is claimed with reasonable certainty, even when the
4 claims are read in light of the specification and prosecution history.

5 Additionally, at least claim 19 (and its dependent claims) of the '798 Patent recites functions
6 without sufficient structure for performing the function, thus invoking 35 U.S.C. § 112, ¶ 6. Because
7 the specification of the '798 Patent also does not recite structure for performing the function, the
8 following terms are invalid as indefinite under 35 U.S.C. §§ 112, ¶ 2:

- 9 • "an identification component that identifies a packet type of a packet of symbols" (claim
10 19)
- 11 • "a transmission component that transmits a synchronization symbol that corresponds to the
12 identified packet type . . . and that transmits the symbols of the packet" (claim 19)

13 The indefiniteness issues identified above apply to all claims that include the identified
14 limitations, either explicitly or through claim dependency, and are not limited to the exemplary
15 claims listed above.

16 **VI. THE ASSERTED CLAIMS OF THE '265 PATENT ARE INVALID**

17 **A. Identification of Prior Art**

18 Defendant contends that the prior art references charted in Exhibits D-1 through D-12,
19 discussed in any incorporated IPR petition, discussed or cited in the prosecution history and/or any
20 reexamination proceedings, and/or described below anticipate and/or render obvious, alone or in
21 combination, one or more of the Asserted Claims of the '265 Patent.

22 **1. The '265 Patent Claims Are Anticipated by the Prior Art**

23 Some or all of the Asserted Claims of the '265 Patent are invalid as anticipated under 35
24 U.S.C. § 102 in view of each of the prior art references identified in the claim charts included in
25 Exhibit D-1 through D-12 and/or identified in any incorporated IPR petition, which identify specific
26 examples of where each limitation of the Asserted Claims is found in the prior art references. As
27 explained above, the cited portions of prior art references identified in the attached claim charts are
28 exemplary in nature and representative of the content and teaching of the prior art references, and

1 DATED: December 5, 2024

Respectfully submitted,
/s/ Lauren E. Douville
B. Trent Webb (*pro hac vice*)
Lauren E. Douville (*pro hac vice*)
Ryan J. Schletzbaum (*pro hac vice*)
SHOOK, HARDY & BACON LLP
2555 Grand Boulevard
Kansas City, Missouri 64108-2613
T: 816.474.6550
F: 816.421.5547
bwebb@shb.com
ldouville@shb.com
rschletzbaum@shb.com

Amelia Murray (*pro hac vice*)
SHOOK, HARDY & BACON LLP
111 S. Wacker Dr., Suite 4700
Chicago, IL 60606
T: 312.704.7700
F: 312.558.1195
aemurray@shb.com

Megan Whyman Olesek
Keeley I. Vega
TURNER BOYD SERAPHINE LLP
155 Bovet Road, Suite 750
San Mateo, CA 94402
T: 650.521.5930
olesek@turnerboyd.com
vega@turnerboyd.com

Attorneys for Defendant HP Inc

19 **CERTIFICATE OF SERVICE**

20 I hereby certify that on December 5, 2024, I served a copy of the foregoing along with
21 accompanying exhibits and document production via electronic mail on all counsel of record.

23 /s/ Dana E. Hardy
24 On behalf of Lauren E. Douville.

Exhibit C-2

Invalidity of the '798 Patent in View of U.S. Patent No. 5,686,965 ("Auld")

This chart provides examples of where the prior art reference discloses, either expressly or inherently, the claim elements of asserted U.S. Patent No. 7,746,798 ("the '798 patent"). The citations herein are exemplary and are not exhaustive; Defendant has cited the most relevant portions of the identified prior art. For each citation, Defendant hereby incorporates by reference the entire cited prior art reference, and reserves the right to rely on uncited portions of each identified prior art reference. To provide context or identify the state of the art at the time of the alleged invention, Defendant reserves the right to rely on: (1) uncited portions of the identified prior art; (2) other prior art not identified herein; (3) references that show the state of the art (irrespective of whether such references themselves qualify as prior art); (4) factual testimony from the inventors or authors of the prior art references, or purveyors of prior art devices; and (5) expert testimony.

Where Defendant cites a figure, the citation includes any text associated with the figure. Where Defendant cites passages relating to a figure, the citation includes the associated figure any other text associated with that figure. Where Defendant cites particular prior art as disclosing a claim element, that citation equally applies to all identical or similar claim elements in this chart. Defendant reserves the right to further assert that a claim element is inherent, the element would have been obvious to a person of ordinary skill in the art, or the element is disclosed in other prior art that may be combined with the cited prior art. To establish the inherency of certain features of the prior art to invalidate the asserted claims, Defendant may rely on cited or uncited portions of the prior art, other documents, factual testimony, and expert testimony.


This invalidity chart is based on information presently available to Defendant, and the information herein is provided as a non-limiting example.

Auld was filed on September 18, 1995, and published on November 11, 1997; therefore, Auld is prior art under at least 35 U.S.C. § 102(b).

Exhibit C-2

Claim Limitation		Auld
Claim 11		
11[p]	A method for identifying packet types of packets of symbols, the method comprising:	To the extent the preamble is limiting, Auld discloses this element and/or renders it obvious either alone or in combination with other references and/or the knowledge of one of ordinary skill in the art. <i>See, e.g.:</i> “The logical structure of a typical MPEG system bitstream is illustrated in FIG. 5, which represents the highest level MPEG data structure, and is often termed an ISO (International Standards Organization) system stream. An MPEG ISO system bitstream typically includes a sequence of packs 90. A typical pack 90, in turn, includes a sequence pack layer headers 92, a system header packet 94, a sequence of packets 96, 98, and 100, and an end code 102. The pack layer header 92 contains a pack start code, or sync code, used for synchronization purposes, and a system clock value. The system header packet 94 contains a variety of information commonly referred to as system housekeeping data, and a system stream ID which is used to differentiate among multiple system streams. A sequence of one or more packets typically contains either encoded audio or encoded video stream data. The final element of a pack 90 is an end code 102. It should be noted that any one packet 96, 98, 100 may carry either audio or video data, but not both simultaneously. With respect to a packet containing video data, typically a video packet header and one or more video frames is contained within the packet. Further illustrated in FIG. 5 are the hierarchical layers comprising a video packet 96. The various layers comprising the video data stream include a sequence layer, a group of pictures layer, a picture layer, a slice layer, a macroblock layer, and a block layer. The functions associated with each of these layers is also provided in FIG. 5. It should be noted that each packet or layer may contain one or more synchronization codes, referred to in the art as sync code, which are necessary to properly synchronize and align the encoded bitstream.” Auld, 10:16–45.

Exhibit C-2



Auld, Fig. 5.


“A state diagram showing the system synchronization process is shown in FIG. 6. Assuming, for example, that the system synchronizer 62 is not synchronized to the bitstream, the system synchronizer 62 searches for a sync code, or start code, comprising the packet layer 96, as at 120. The length field of the packet layer 96 is then extracted, and the length of the packet data field is read, as at 122 and 124 respectively. The packet data field of the packet layer 96 may contain emulations of the sync codes of other bitstream data, such as those comprising audio packets. The system synchronizer 62, however, analyzes the bitstream to determine whether a legitimate video packet sync code has been located, as at 126. If either the video sync code has not been located or a sync code emulation is located, the system synchronizer 62 will continue to predict the occurrence of a legitimate video sync code in the bitstream. This prediction is based on the length field of the previously occurring header. The system synchronizer 62 then proceeds to repeat the sync code searching process, as at 134. If a sync code is located within the bitstream, as at 126 and 132, the

Exhibit C-2

length of the field in bytes is extracted and read, as at 128 and 130 respectively. When the sync code is located, as at 126 and 132, the bitstream is considered synchronized. The system synchronizer 62 will then attempt to locate the sync code of the next packet, as at 136. If no sync code can be located for the next packet 98, the system synchronizer 62 begins to search for a sync code and the bitstream transitions from a synchronized state to a non-synchronized state, as at 138."

Auld, 10:46-11:6.

FIG. 6



Auld, Fig. 6.

Exhibit C-2

“Turning now to FIG. 7, there is shown a flow diagram illustrating the process of extracting the video bitstream component of a multiplexed system bitstream that is performed by the system stream pre-parser 63. The system stream pre-parser 63 is intended to pass only the video bitstream component of a multiplexed system bitstream. Any non-video information components, such as the audio portion of the system bitstream, are discarded by the system stream pre-parser 63. Essentially, the packet sync code, stream ID, and packet length information at the packet layer are used to identify the type of packet being analyzed, and to skip any non-video packets. When a packet sync code has been located, as at 150, so-called stuffing data is removed, and the length of the packet data field is updated to reflect the decrease in length of the packet word resulting from the removal of the stuffing data, as at 152. Stuffing data, also referred to as stuffing bits or stuffing bytes, are typically code words that may be inserted into a compressed bitstream to increase the bit rate of the stream to that required by the digital storage or transmission medium. The stuffing data may later be discarded by the decoder during the decoding process.”

Auld, 11:31–52.





FIG. 7

Auld, Fig. 7.

Exhibit C-2

		To the extent that Plaintiff contends that Auld does not explicitly or inherently disclose this element, or render obvious this element, this element is obvious in combination with one or more references and/or with the knowledge of a person of ordinary skill in the art, as set forth in Defendant's Invalidity Contentions.
11(a)	receiving a synchronization symbol by a receiver indicating a packet type, each packet type having a different synchronization symbol;	Auld discloses and/or renders obvious this limitation. <i>See, e.g.:</i> "The logical structure of a typical MPEG system bitstream is illustrated in FIG. 5, which represents the highest level MPEG data structure, and is often termed an ISO (International Standards Organization) system stream. An MPEG ISO system bitstream typically includes a sequence of packs 90. A typical pack 90, in turn, includes a sequence pack layer headers 92, a system header packet 94, a sequence of packets 96, 98, and 100, and an end code 102. The pack layer header 92 contains a pack start code, or sync code, used for synchronization purposes, and a system clock value. The system header packet 94 contains a variety of information commonly referred to as system housekeeping data, and a system stream ID which is used to differentiate among multiple system streams. A sequence of one or more packets typically contains either encoded audio or encoded video stream data. The final element of a pack 90 is an end code 102. It should be noted that any one packet 96, 98, 100 may carry either audio or video data, but not both simultaneously. With respect to a packet containing video data, typically a video packet header and one or more video frames is contained within the packet. Further illustrated in FIG. 5 are the hierarchical layers comprising a video packet 96. The various layers comprising the video data stream include a sequence layer, a group of pictures layer, a picture layer, a slice layer, a macroblock layer, and a block layer. The functions associated with each of these layers is also provided in FIG. 5. It should be noted that each packet or layer may contain one or more synchronization codes, referred to in the art as sync code, which are necessary to properly synchronize and align the encoded bitstream." Auld, 10:16–45.

Exhibit C-2



Auld, Fig. 5.


“A state diagram showing the system synchronization process is shown in FIG. 6. Assuming, for example, that the system synchronizer 62 is not synchronized to the bitstream, the system synchronizer 62 searches for a sync code, or start code, comprising the packet layer 96, as at 120. The length field of the packet layer 96 is then extracted, and the length of the packet data field is read, as at 122 and 124 respectively. The packet data field of the packet layer 96 may contain emulations of the sync codes of other bitstream data, such as those comprising audio packets. The system synchronizer 62, however, analyzes the bitstream to determine whether a legitimate video packet sync code has been located, as at 126. If either the video sync code has not been located or a sync code emulation is located, the system synchronizer 62 will continue to predict the occurrence of a legitimate video sync code in the bitstream. This prediction is based on the length field of the previously occurring header. The system synchronizer 62 then proceeds to repeat the sync code searching process, as at 134. If a sync code is located within the bitstream, as at 126 and 132, the

Exhibit C-2

length of the field in bytes is extracted and read, as at 128 and 130 respectively. When the sync code is located, as at 126 and 132, the bitstream is considered synchronized. The system synchronizer 62 will then attempt to locate the sync code of the next packet, as at 136. If no sync code can be located for the next packet 98, the system synchronizer 62 begins to search for a sync code and the bitstream transitions from a synchronized state to a non-synchronized state, as at 138."

Auld, 10:46-11:6.

FIG. 6



Auld, Fig. 6.

Exhibit C-2

“Turning now to FIG. 7, there is shown a flow diagram illustrating the process of extracting the video bitstream component of a multiplexed system bitstream that is performed by the system stream pre-parser 63. The system stream pre-parser 63 is intended to pass only the video bitstream component of a multiplexed system bitstream. Any non-video information components, such as the audio portion of the system bitstream, are discarded by the system stream pre-parser 63. Essentially, the packet sync code, stream ID, and packet length information at the packet layer are used to identify the type of packet being analyzed, and to skip any non-video packets. When a packet sync code has been located, as at 150, so-called stuffing data is removed, and the length of the packet data field is updated to reflect the decrease in length of the packet word resulting from the removal of the stuffing data, as at 152. Stuffing data, also referred to as stuffing bits or stuffing bytes, are typically code words that may be inserted into a compressed bitstream to increase the bit rate of the stream to that required by the digital storage or transmission medium. The stuffing data may later be discarded by the decoder during the decoding process.”

Auld, 11:31–52.




FIG. 7

Auld, Fig. 7.

Exhibit C-2

		“Error propagation due to the emulation of non-video sync codes is generally not a concern with respect to the video bitstream input to the video stream synchronizer 65 and video stream pre-parser 66. Once the system stream pre-parser 63 parses the video bitstream from the multiplexed system bitstream, non-video sync codes generally are not present in the video bitstream. The emulation of non-video sync codes associated with conventional decoding schemes results from mis-interpreting audio and other non-video bitstream sync codes as legitimate video bitstream sync codes. Since these extraneous non-video bitstreams have, in accordance with the present invention, been stripped from the system bitstream by the system stream pre-parser 63, it is highly unlikely that such errant sync codes could remain within the video bitstream. Moreover, in the event that the video stream synchronizer 65 locates an unexpected legitimate sync code within the video bitstream, the video stream synchronizer 65 will synchronize to this unexpected sync code. Thus, synchronization of the video bitstream would simply be established earlier than expected, and would not result in the propagation of a synchronization error into the post-parser 56.” Auld, 12:32–53. To the extent that Plaintiff contends that Auld does not explicitly or inherently disclose this element, or render obvious this element, this element is obvious in combination with one or more references and/or with the knowledge of a person of ordinary skill in the art, as set forth in Defendant’s Invalidity Contentions.
11(b)	receiving a packet of symbols;	Auld discloses and/or renders obvious this limitation. <i>See, e.g.:</i> “Turning now to FIG. 7, there is shown a flow diagram illustrating the process of extracting the video bitstream component of a multiplexed system bitstream that is performed by the system stream pre-parser 63. The system stream pre-parser 63 is intended to pass only the video bitstream component of a multiplexed system bitstream. Any non-video information components, such as the audio portion of the system bitstream, are discarded by the system stream pre-parser 63. Essentially, the packet sync code, stream ID, and packet length information at the packet layer are used to identify the type of packet being analyzed, and to skip any non-video packets. When a packet sync code has been located, as at 150, so-called stuffing data is removed, and the length of the packet data field is updated to reflect the decrease in length of the packet word resulting from the removal

Exhibit C-2

of the stuffing data, as at 152. Stuffing data, also referred to as stuffing bits or stuffing bytes, are typically code words that may be inserted into a compressed bitstream to increase the bit rate of the stream to that required by the digital storage or transmission medium. The stuffing data may later be discarded by the decoder during the decoding process.”

Auld, 11:31–52.




FIG. 7

Auld, Fig. 7.

“After removal of the stuffing data and updating the length of the packet data field, the stream ID of the packet is read to determine whether the bitstream is a video bitstream, as at 154. If the system ID indicates that the bitstream does indeed contain video data, the data representative of the video information is transferred to the channel buffer 54, as shown at 156. If the system ID indicates that the packet does not contain video data, the number of bytes comprising the non-video data are skipped, as at 158, and subsequent packets are searched, as at 162. If a packet sync code cannot be located, the data in the fixed rate channel 50 is considered invalid and is purged, as at 160, and subsequent packet searching continues, as at 162.

Exhibit C-2

		While the system synchronizer 62 may initially synchronize incorrectly to a video bitstream, the probability of incorrect synchronization occurring is very low. Incorrect synchronization may occur as a result of data emulating video start codes, or sync codes, in the packet data field of a packet 96. In order to mis-synchronize, the system synchronizer 62 must find a second sync code after discarding the "length" bits of the packet data field, as at 158 of FIG. 7. Preferably, the "length" field is represented by a 16-bit symbol, and the sync code is represented by a 32-bit symbol. Assuming that the bits in the data field of the packet layer 96 arrive in random order, that is, "0" and "1" bits are equally probable, the mis-synchronization sequence just described has a probability of $0.548 = 3.55 \times 10^{-15}$. With a bitstream rate of 15 megabytes per second, a mis-synchronization failure would occur every 18.7×10^6 seconds, which is approximately 217 days. Note that any mis-synchronization failures can occur only if the system synchronizer 62 is not already synchronized to the bitstream. If such a failure is encountered, the system synchronizer 62 can subsequently reacquire synchronization within a period roughly equivalent to the time required to analyze two packets of data." Auld, 11:53–12:30. To the extent that Plaintiff contends that Auld does not explicitly or inherently disclose this element, or render obvious this element, this element is obvious in combination with one or more references and/or with the knowledge of a person of ordinary skill in the art, as set forth in Defendant's Invalidity Contentions.
11(c)	indicating that the received packet of symbols has the packet type of the received synchronization symbol; and	Auld discloses and/or renders obvious this limitation. <i>See, e.g.:</i> "It should be noted that any one packet 96, 98, 100 may carry either audio or video data, but not both simultaneously. With respect to a packet containing video data, typically a video packet header and one or more video frames is contained within the packet. Further illustrated in FIG. 5 are the hierarchical layers comprising a video packet 96. The various layers comprising the video data stream include a sequence layer, a group of pictures layer, a picture layer, a slice layer, a macroblock layer, and a block layer. The functions associated with each of these layers is also provided in FIG. 5. It should be noted that each packet or layer may contain one or more synchronization codes, referred to in the art as sync code, which are necessary to properly synchronize and align the encoded bitstream."

Exhibit C-2

Auld, 10:32–45.

“Turning now to FIG. 7, there is shown a flow diagram illustrating the process of extracting the video bitstream component of a multiplexed system bitstream that is performed by the system stream pre-parser 63. The system stream pre-parser 63 is intended to pass only the video bitstream component of a multiplexed system bitstream. Any non-video information components, such as the audio portion of the system bitstream, are discarded by the system stream pre-parser 63. Essentially, the packet sync code, stream ID, and packet length information at the packet layer are used to identify the type of packet being analyzed, and to skip any non-video packets. When a packet sync code has been located, as at 150, so-called stuffing data is removed, and the length of the packet data field is updated to reflect the decrease in length of the packet word resulting from the removal of the stuffing data, as at 152. Stuffing data, also referred to as stuffing bits or stuffing bytes, are typically code words that may be inserted into a compressed bitstream to increase the bit rate of the stream to that required by the digital storage or transmission medium. The stuffing data may later be discarded by the decoder during the decoding process.”

Auld, 11:31–52.




FIG. 7

Exhibit C-2

	Auld, Fig. 7. “After removal of the stuffing data and updating the length of the packet data field, the stream ID of the packet is read to determine whether the bitstream is a video bitstream, as at 154. If the system ID indicates that the bitstream does indeed contain video data, the data representative of the video information is transferred to the channel buffer 54, as shown at 156. If the system ID indicates that the packet does not contain video data, the number of bytes comprising the non-video data are skipped, as at 158, and subsequent packets are searched, as at 162. If a packet sync code cannot be located, the data in the fixed rate channel 50 is considered invalid and is purged, as at 160, and subsequent packet searching continues, as at 162.” Auld, 11:53–65. “Error propagation due to the emulation of non-video sync codes is generally not a concern with respect to the video bitstream input to the video stream synchronizer 65 and video stream pre-parser 66. Once the system stream pre-parser 63 parses the video bitstream from the multiplexed system bitstream, non-video sync codes generally are not present in the video bitstream. The emulation of non-video sync codes associated with conventional decoding schemes results from mis-interpreting audio and other non-video bitstream sync codes as legitimate video bitstream sync codes. Since these extraneous non-video bitstreams have, in accordance with the present invention, been stripped from the system bitstream by the system stream pre-parser 63, it is highly unlikely that such errant sync codes could remain within the video bitstream. Moreover, in the event that the video stream synchronizer 65 locates an unexpected legitimate sync code within the video bitstream, the video stream synchronizer 65 will synchronize to this unexpected sync code. Thus, synchronization of the video bitstream would simply be established earlier than expected, and would not result in the propagation of a synchronization error into the post-parser 56.” Auld, 12:32–53. To the extent that Plaintiff contends that Auld does not explicitly or inherently disclose this element, or render obvious this element, this element is obvious in combination with one or more references
--	--

Exhibit C-2


		and/or with the knowledge of a person of ordinary skill in the art, as set forth in Defendant's Invalidity Contentions.
11(d)	the synchronization symbol permitting properly alignment upon receiving with a synchronization primitive to be correctly aligned on a symbol boundary.	Auld discloses and/or renders obvious this limitation. <i>See, e.g.:</i> "It should be noted that any one packet 96, 98, 100 may carry either audio or video data, but not both simultaneously. With respect to a packet containing video data, typically a video packet header and one or more video frames is contained within the packet. Further illustrated in FIG. 5 are the hierarchical layers comprising a video packet 96. The various layers comprising the video data stream include a sequence layer, a group of pictures layer, a picture layer, a slice layer, a macroblock layer, and a block layer. The functions associated with each of these layers is also provided in FIG. 5. It should be noted that each packet or layer may contain one or more synchronization codes, referred to in the art as sync code, which are necessary to properly synchronize and align the encoded bitstream." Auld, 10:32–45. "A state diagram showing the system synchronization process is shown in FIG. 6. Assuming, for example, that the system synchronizer 62 is not synchronized to the bitstream, the system synchronizer 62 searches for a sync code, or start code, comprising the packet layer 96, as at 120. The length field of the packet layer 96 is then extracted, and the length of the packet data field is read, as at 122 and 124 respectively. The packet data field of the packet layer 96 may contain emulations of the sync codes of other bitstream data, such as those comprising audio packets. The system synchronizer 62, however, analyzes the bitstream to determine whether a legitimate video packet sync code has been located, as at 126. If either the video sync code has not been located or a sync code emulation is located, the system synchronizer 62 will continue to predict the occurrence of a legitimate video sync code in the bitstream. This prediction is based on the length field of the previously occurring header. The system synchronizer 62 then proceeds to repeat the sync code searching process, as at 134. If a sync code is located within the bitstream, as at 126 and 132, the length of the field in bytes is extracted and read, as at 128 and 130 respectively. When the sync code is located, as at 126 and 132, the bitstream is considered synchronized. The system synchronizer 62 will then attempt to locate the sync code of the next packet, as at 136. If no sync

Exhibit C-2

code can be located for the next packet 98, the system synchronizer 62 begins to search for a sync code and the bitstream transitions from a synchronized state to a non-synchronized state, as at 138.”

Auld, 10:46–11:6.

FIG. 6



Auld, Fig. 6.

“One important advantage of the present invention concerns the alignment and synchronization of a multiplexed system bitstream received from a fixed rate channel 50 by the pre-parser 52. As


Exhibit C-2

		shown in FIG. 4, the system stream synchronizer 62 aligns the multiplexed system bitstream with respect to sync words embedded in the system bitstream data. Synchronization involves the alignment of data within a packet relative to the position of a sync word within the bitstream. A typical sync word may be represented by a 24-bit or 32-bit word made up of consecutive zeros and ending with a "1" bit. Alignment involves the recognition of word or byte boundaries within a system bitstream in accordance with the syntax of the selected compression algorithm. The location of sync words distributed within the system bitstream delineates the boundaries between informational word groupings within the packet. These word groupings conform to the particular grammar and syntax rules dictated by the compression algorithm employed for data encoding. Once aligned, the organization of the data structure comprising a particular informational word grouping following a sync code is thus of an established format. In a preferred embodiment, the system stream synchronizer 62 performs synchronization and alignment in a serial fashion, one bit at a time. The synchronized system bitstream is then processed by the system stream pre-parser 63.” Auld, 11:7–30. To the extent that Plaintiff contends that Auld does not explicitly or inherently disclose this element, or render obvious this element, this element is obvious in combination with one or more references and/or with the knowledge of a person of ordinary skill in the art, as set forth in Defendant’s Invalidity Contentions.
Claim 19		
19[p]	A communications device for transmitting packet types of packets, comprising:	To the extent the preamble is limiting, Auld discloses this element and/or renders it obvious either alone or in combination with other references and/or the knowledge of one of ordinary skill in the art. <i>See, e.g.:</i> “A typical and conventional video encoding and decoding scheme is shown in FIG. 1. An initial video signal 10 is input to a video assembler 12 wherein pictures or frames are assembled. The assembled pictures are then input into a video encoder 14 which compresses the video data in accordance with a selected compression algorithm. Similarly, an initial audio signal 16 is input into an audio assembler 18 wherein audio frames are assembled. The audio frames are then input to an audio encoder 20 for audio compression. A system clock 15 provides time stamp information which

Exhibit C-2

is incorporated into the encoded audio and video signals for synchronization purposes. The encoded video signal output from video encoder 14 and the encoded audio signal output from audio encoder 20 are then input to a buffer/multiplexer 22 to produce a multiplexed system bitstream. This multiplexed system bitstream may then be stored in a digital storage media 24 or transmitted along an appropriate transmission channel 24 for later reception, decoding, and display.”

Auld, 5:51–6:2.




Auld, Fig. 1.

“In accordance with one aspect of the present invention, a multiplexed system bitstream, comprising at least an audio bitstream portion and a video bitstream portion, is input to a pre-parser through a fixed rate channel. A system bitstream synchronizer then initially aligns and synchronizes to the received multiplexed system bitstream. The system stream pre-parser also extracts the video bitstream portion from the multiplexed system bitstream. A video stream synchronizer then analyzes the video bitstream for the presence of a video sync code. The various layers comprising the video bitstream are subsequently synchronized and aligned by a video stream pre-parser.

Exhibit C-2

	Preferably, the parsed video bitstream is then bit-serially written to a channel write FIFO (First-in-First-Out) buffer prior to being transferred to a channel buffer.” Auld, 7:35–49. “A typical pack 90, in turn, includes a sequence pack layer headers 92, a system header packet 94, a sequence of packets 96, 98, and 100, and an end code 102. The pack layer header 92 contains a pack start code, or sync code, used for synchronization purposes, and a system clock value. The system header packet 94 contains a variety of information commonly referred to as system housekeeping data, and a system stream ID which is used to differentiate among multiple system streams. A sequence of one or more packets typically contains either encoded audio or encoded video stream data. The final element of a pack 90 is an end code 102. It should be noted that any one packet 96, 98, 100 may carry either audio or video data, but not both simultaneously. With respect to a packet containing video data, typically a video packet header and one or more video frames is contained within the packet. Further illustrated in FIG. 5 are the hierarchical layers comprising a video packet 96. The various layers comprising the video data stream include a sequence layer, a group of pictures layer, a picture layer, a slice layer, a macroblock layer, and a block layer. The functions associated with each of these layers is also provided in FIG. 5. It should be noted that each packet or layer may contain one or more synchronization codes, referred to in the art as sync code, which are necessary to properly synchronize and align the encoded bitstream.” Auld, 10:21–45.
--	---

Exhibit C-2




Auld, Fig. 5.

“1. A Motion Picture Experts Group (MPEG) data transmission system, comprising:
a transmitter including an encoder for encoding data in accordance with an MPEG standard for producing and transmitting a multiplexed bitstream, said multiplexed bitstream including:
video bitstream data; and
non-video bitstream data;
a communications channel for carrying the multiplexed bitstream from the transmitter; and
a receiver for receiving the bitstream from the communications channel, the receiver including:
a channel buffer;
a pre-parser for parsing the multiplexed bitstream, separating the video bitstream data from the non-video bitstream data, and storing the separated video bitstream data in the channel buffer; and
a post-parser for reading the video bitstream data from the channel buffer, and parsing and translating the video bitstream data to produce video symbols.”

Exhibit C-2

		Auld, claim 1. To the extent that Plaintiff contends that Auld does not explicitly or inherently disclose this element, or render obvious this element, this element is obvious in combination with one or more references and/or with the knowledge of a person of ordinary skill in the art, as set forth in Defendant's Invalidity Contentions.
19(a)	an identification component that identifies a packet type of a packet of symbols; and	Auld discloses and/or renders obvious this limitation. <i>See, e.g.:</i>

Exhibit C-2



Auld, Fig. 4.


“A state diagram showing the system synchronization process is shown in FIG. 6. Assuming, for example, that the system synchronizer 62 is not synchronized to the bitstream, the system synchronizer 62 searches for a sync code, or start code, comprising the packet layer 96, as at 120.

Exhibit C-2

		The length field of the packet layer 96 is then extracted, and the length of the packet data field is read, as at 122 and 124 respectively. The packet data field of the packet layer 96 may contain emulations of the sync codes of other bitstream data, such as those comprising audio packets. The system synchronizer 62, however, analyzes the bitstream to determine whether a legitimate video packet sync code has been located, as at 126. If either the video sync code has not been located or a sync code emulation is located, the system synchronizer 62 will continue to predict the occurrence of a legitimate video sync code in the bitstream. This prediction is based on the length field of the previously occurring header. The system synchronizer 62 then proceeds to repeat the sync code searching process, as at 134. If a sync code is located within the bitstream, as at 126 and 132, the length of the field in bytes is extracted and read, as at 128 and 130 respectively. When the sync code is located, as at 126 and 132, the bitstream is considered synchronized. The system synchronizer 62 will then attempt to locate the sync code of the next packet, as at 136. If no sync code can be located for the next packet 98, the system synchronizer 62 begins to search for a sync code and the bitstream transitions from a synchronized state to a non-synchronized state, as at 138.” Auld, 10:46–11:6.
--	--	--

Exhibit C-2

FIG. 6



Auld, Fig. 6.

“Turning now to FIG. 7, there is shown a flow diagram illustrating the process of extracting the video bitstream component of a multiplexed system bitstream that is performed by the system stream pre-parser 63. The system stream pre-parser 63 is intended to pass only the video bitstream component of a multiplexed system bitstream. Any non-video information components, such as the audio portion of the system bitstream, are discarded by the system stream pre-parser 63. Essentially, the packet sync code, stream ID, and packet length information at the packet layer are used to identify the type of packet being analyzed, and to skip any non-video packets. When a packet sync

Exhibit C-2

code has been located, as at 150, so-called stuffing data is removed, and the length of the packet data field is updated to reflect the decrease in length of the packet word resulting from the removal of the stuffing data, as at 152. Stuffing data, also referred to as stuffing bits or stuffing bytes, are typically code words that may be inserted into a compressed bitstream to increase the bit rate of the stream to that required by the digital storage or transmission medium. The stuffing data may later be discarded by the decoder during the decoding process.”

Auld, 11:31–52.




FIG. 7

Auld, Fig. 7.

To the extent that Plaintiff contends that Auld does not explicitly or inherently disclose this element, or render obvious this element, this element is obvious in combination with one or more references and/or with the knowledge of a person of ordinary skill in the art, as set forth in Defendant’s Invalidity Contentions.

19(b)

a transmission component that transmits a synchronization symbol


Auld discloses and/or renders obvious this limitation. *See, e.g.:*

Exhibit C-2

	that corresponds to the identified packet type, the transmitted synchronization symbol providing synchronization information and each packet type having a different synchronization symbol and permitting an external receiving node to properly align with a synchronization primitive to be correctly aligned on a symbol boundary, and that transmits the symbols of the packet.	“A state diagram showing the system synchronization process is shown in FIG. 6. Assuming, for example, that the system synchronizer 62 is not synchronized to the bitstream, the system synchronizer 62 searches for a sync code, or start code, comprising the packet layer 96, as at 120. The length field of the packet layer 96 is then extracted, and the length of the packet data field is read, as at 122 and 124 respectively. The packet data field of the packet layer 96 may contain emulations of the sync codes of other bitstream data, such as those comprising audio packets. The system synchronizer 62, however, analyzes the bitstream to determine whether a legitimate video packet sync code has been located, as at 126. If either the video sync code has not been located or a sync code emulation is located, the system synchronizer 62 will continue to predict the occurrence of a legitimate video sync code in the bitstream. This prediction is based on the length field of the previously occurring header. The system synchronizer 62 then proceeds to repeat the sync code searching process, as at 134. If a sync code is located within the bitstream, as at 126 and 132, the length of the field in bytes is extracted and read, as at 128 and 130 respectively. When the sync code is located, as at 126 and 132, the bitstream is considered synchronized. The system synchronizer 62 will then attempt to locate the sync code of the next packet, as at 136. If no sync code can be located for the next packet 98, the system synchronizer 62 begins to search for a sync code and the bitstream transitions from a synchronized state to a non-synchronized state, as at 138.” Auld, 10:46–11:6.
--	---	--

Exhibit C-2

FIG. 6



Auld, Fig. 6.

“Turning now to FIG. 7, there is shown a flow diagram illustrating the process of extracting the video bitstream component of a multiplexed system bitstream that is performed by the system stream pre-parser 63. The system stream pre-parser 63 is intended to pass only the video bitstream component of a multiplexed system bitstream. Any non-video information components, such as the audio portion of the system bitstream, are discarded by the system stream pre-parser 63. Essentially, the packet sync code, stream ID, and packet length information at the packet layer are used to identify the type of packet being analyzed, and to skip any non-video packets. When a packet sync

Exhibit C-2

code has been located, as at 150, so-called stuffing data is removed, and the length of the packet data field is updated to reflect the decrease in length of the packet word resulting from the removal of the stuffing data, as at 152. Stuffing data, also referred to as stuffing bits or stuffing bytes, are typically code words that may be inserted into a compressed bitstream to increase the bit rate of the stream to that required by the digital storage or transmission medium. The stuffing data may later be discarded by the decoder during the decoding process.”

Auld, 11:31–52.




FIG. 7

Auld, Fig. 7.

“Error propagation due to the emulation of non-video sync codes is generally not a concern with respect to the video bitstream input to the video stream synchronizer 65 and video stream pre-parser 66. Once the system stream pre-parser 63 parses the video bitstream from the multiplexed system bitstream, non-video sync codes generally are not present in the video bitstream. The emulation of non-video sync codes associated with conventional decoding schemes results from mis-interpreting audio and other non-video bitstream sync codes as legitimate video bitstream sync codes. Since these extraneous non-video bitstreams have, in accordance with the present invention, been stripped

Exhibit C-2

	from the system bitstream by the system stream pre-parser 63, it is highly unlikely that such errant sync codes could remain within the video bitstream. Moreover, in the event that the video stream synchronizer 65 locates an unexpected legitimate sync code within the video bitstream, the video stream synchronizer 65 will synchronize to this unexpected sync code. Thus, synchronization of the video bitstream would simply be established earlier than expected, and would not result in the propagation of a synchronization error into the post-parser 56.” Auld, 12:32–53. “It should be noted that any one packet 96, 98, 100 may carry either audio or video data, but not both simultaneously. With respect to a packet containing video data, typically a video packet header and one or more video frames is contained within the packet. Further illustrated in FIG. 5 are the hierarchical layers comprising a video packet 96. The various layers comprising the video data stream include a sequence layer, a group of pictures layer, a picture layer, a slice layer, a macroblock layer, and a block layer. The functions associated with each of these layers is also provided in FIG. 5. It should be noted that each packet or layer may contain one or more synchronization codes, referred to in the art as sync code, which are necessary to properly synchronize and align the encoded bitstream.” Auld, 10:32–45. “One important advantage of the present invention concerns the alignment and synchronization of a multiplexed system bitstream received from a fixed rate channel 50 by the pre-parser 52. As shown in FIG. 4, the system stream synchronizer 62 aligns the multiplexed system bitstream with respect to sync words embedded in the system bitstream data. Synchronization involves the alignment of data within a packet relative to the position of a sync word within the bitstream. A typical sync word may be represented by a 24-bit or 32-bit word made up of consecutive zeros and ending with a "1" bit. Alignment involves the recognition of word or byte boundaries within a system bitstream in accordance with the syntax of the selected compression algorithm. The location of sync words distributed within the system bitstream delineates the boundaries between informational word groupings within the packet. These word groupings conform to the particular grammar and syntax rules dictated by the compression algorithm employed for data encoding. Once aligned, the organization of the data structure comprising a particular informational word grouping
--	---

Exhibit C-2


		following a sync code is thus of an established format. In a preferred embodiment, the system stream synchronizer 62 performs synchronization and alignment in a serial fashion, one bit at a time. The synchronized system bitstream is then processed by the system stream pre-parser 63.” Auld, 11:7–30. To the extent that Plaintiff contends that Auld does not explicitly or inherently disclose this element, or render obvious this element, this element is obvious in combination with one or more references and/or with the knowledge of a person of ordinary skill in the art, as set forth in Defendant’s Invalidity Contentions.
Claim 22		
22[p]	The communications device of claim 19	Auld discloses, expressly or inherently, the transmitting device of claim 19 under any proposed construction. <i>See</i> exemplary disclosures and corresponding discussion for claim 19.
22	wherein the synchronization symbol is transmitted before transmitting the symbols of the packet.	Auld discloses and/or renders obvious this limitation. <i>See, e.g.:</i> “A state diagram showing the system synchronization process is shown in FIG. 6. Assuming, for example, that the system synchronizer 62 is not synchronized to the bitstream, the system synchronizer 62 searches for a sync code, or start code, comprising the packet layer 96, as at 120. The length field of the packet layer 96 is then extracted, and the length of the packet data field is read, as at 122 and 124 respectively. The packet data field of the packet layer 96 may contain emulations of the sync codes of other bitstream data, such as those comprising audio packets. The system synchronizer 62, however, analyzes the bitstream to determine whether a legitimate video packet sync code has been located, as at 126. If either the video sync code has not been located or a sync code emulation is located, the system synchronizer 62 will continue to predict the occurrence of a legitimate video sync code in the bitstream. This prediction is based on the length field of the previously occurring header. The system synchronizer 62 then proceeds to repeat the sync code searching process, as at 134. If a sync code is located within the bitstream, as at 126 and 132, the length of the field in bytes is extracted and read, as at 128 and 130 respectively. When the sync code is located, as at 126 and 132, the bitstream is considered synchronized. The system synchronizer 62 will then attempt to locate the sync code of the next packet, as at 136. If no sync

Exhibit C-2

code can be located for the next packet 98, the system synchronizer 62 begins to search for a sync code and the bitstream transitions from a synchronized state to a non-synchronized state, as at 138.”

Auld, 10:46–11:6.

FIG. 6



Auld, Fig. 6.

“Turning now to FIG. 7, there is shown a flow diagram illustrating the process of extracting the video bitstream component of a multiplexed system bitstream that is performed by the system

Exhibit C-2

stream pre-parser 63. The system stream pre-parser 63 is intended to pass only the video bitstream component of a multiplexed system bitstream. Any non-video information components, such as the audio portion of the system bitstream, are discarded by the system stream pre-parser 63. Essentially, the packet sync code, stream ID, and packet length information at the packet layer are used to identify the type of packet being analyzed, and to skip any non-video packets. When a packet sync code has been located, as at 150, so-called stuffing data is removed, and the length of the packet data field is updated to reflect the decrease in length of the packet word resulting from the removal of the stuffing data, as at 152. Stuffing data, also referred to as stuffing bits or stuffing bytes, are typically code words that may be inserted into a compressed bitstream to increase the bit rate of the stream to that required by the digital storage or transmission medium. The stuffing data may later be discarded by the decoder during the decoding process.”

Auld, 11:31–52.




FIG. 7

Auld, Fig. 7.

“After removal of the stuffing data and updating the length of the packet data field, the stream ID of the packet is read to determine whether the bitstream is a video bitstream, as at 154. If the system

Exhibit C-2

		ID indicates that the bitstream does indeed contain video data, the data representative of the video information is transferred to the channel buffer 54, as shown at 156. If the system ID indicates that the packet does not contain video data, the number of bytes comprising the non-video data are skipped, as at 158, and subsequent packets are searched, as at 162. If a packet sync code cannot be located, the data in the fixed rate channel 50 is considered invalid and is purged, as at 160, and subsequent packet searching continues, as at 162.” Auld, 11:53–65. To the extent that Plaintiff contends that Auld does not explicitly or inherently disclose this element, or render obvious this element, this element is obvious in combination with one or more references and/or with the knowledge of a person of ordinary skill in the art, as set forth in Defendant’s Invalidity Contentions.
Claim 25		
25[p]	The communications device of claim 19	Auld discloses, expressly or inherently, the transmitting device of claim 19 under any proposed construction. <i>See exemplary disclosures and corresponding discussion for claim 19.</i>
25	wherein the packet types include a control packet.	One of ordinary skill would find this limitation disclosed either expressly or inherently in the teachings of the reference and its incorporated disclosures taken as a whole, or in combination with the state of the art at the time of the alleged invention, as evidenced, for example, by substantial other references identified in Defendant’s invalidity contentions and accompanying charts. Rather than repeat those disclosures here, they are incorporated by reference into this chart.
Claim 28		
28[p]	The communications device of claim 19	Auld discloses, expressly or inherently, the transmitting device of claim 19 under any proposed construction. <i>See exemplary disclosures and corresponding discussion for claim 19.</i>
28	wherein the synchronization primitive is made up of a sequence	One of ordinary skill would find this limitation disclosed either expressly or inherently in the teachings of the reference and its incorporated disclosures taken as a whole, or in combination with the state of the art at the time of the alleged invention, as evidenced, for example, by substantial

Exhibit C-2

	of bit-string synchronization symbols.	other references identified in Defendant's invalidity contentions and accompanying charts. Rather than repeat those disclosures here, they are incorporated by reference into this chart.
--	--	---

Exhibit C-8

Invalidity of the '798 Patent in View of Additional Prior Art References

This chart provides examples of where the prior art reference discloses, either expressly or inherently, the claim elements of asserted U.S. Patent No. 7,746,798 (“the ’798 patent”). The citations herein are exemplary and are not exhaustive; Defendant has cited the most relevant portions of the identified prior art. For each citation, Defendant hereby incorporates by reference the entire cited prior art reference, and reserves the right to rely on uncited portions of each identified prior art reference. To provide context or identify the state of the art at the time of the alleged invention, Defendant reserves the right to rely on: (1) uncited portions of the identified prior art; (2) other prior art not identified herein; (3) references that show the state of the art (irrespective of whether such references themselves qualify as prior art); (4) factual testimony from the inventors or authors of the prior art references, or purveyors of prior art devices; and (5) expert testimony.

Where Defendant cites a figure, the citation includes any text associated with the figure. Where Defendant cites passages relating to a figure, the citation includes the associated figure any other text associated with that figure. Where Defendant cites particular prior art as disclosing a claim element, that citation equally applies to all identical or similar claim elements in this chart. Defendant reserves the right to further assert that a claim element is inherent, the element would have been obvious to a person of ordinary skill in the art, or the element is disclosed in other prior art that may be combined with the cited prior art. To establish the inherency of certain features of the prior art to invalidate the asserted claims, Defendant may rely on cited or uncited portions of the prior art, other documents, factual testimony, and expert testimony.

This invalidity chart is based on information presently available to Defendant, and the information herein is provided as a non-limiting example.

U.S. Patent No. 5,917,830 (“Chen”) was filed on October 18, 1996, and published on June 29, 1999; therefore, Chen is prior art under at least 35 U.S.C. § 102(b).

U.S. Patent No. 6,430,230 (“Cunningham”) was filed on June 30, 1999, and published on August 6, 2002; therefore, Cunningham is prior art under at least 35 U.S.C. § 102(e)(1).

U.S. Patent No. 5,801,781 (“Hiroshima”) was filed on June 19, 1996, and published on September 1, 1998; therefore, Hiroshima is prior art under at least 35 U.S.C. § 102(b).

WIPO Patent Publication No. 2000/011831 (“Haartsen”) was filed on August 17, 1999, and published on March 2, 2000; therefore, Haartsen is prior art under at least 35 U.S.C. § 102(b).

Exhibit C-8

U.S. Patent No. 5,974,464 (“Shin”) was filed on September 30, 1996, and published on October 26, 1999; therefore, Shin is prior art under at least 35 U.S.C. § 102(b).

U.S. Patent No. 6,498,568 (“Austin”) was filed on September 24, 1997, and published on December 24, 2002; therefore, Austin is prior art under at least 35 U.S.C. 102(e)(1).


This chart refers to the above references collectively as the “Additional Prior Art References.”

Exhibit C-8

Claim Limitation		Chen
Claim 11		
11[p]	A method for identifying packet types of packets of symbols, the method comprising:	To the extent the preamble is limiting, Chen discloses this element and/or renders it obvious either alone or in combination with other references and/or the knowledge of one of ordinary skill in the art. <i>See, e.g.:</i> “In the MPEG-2 system (and the similar DigiCipher® II system proprietary to General Instrument Corporation, the assignee hereof) a transport stream, or transport multiplex is made up of a contiguous set of fixed length packets. Each packet is 188 total bytes in length, with the first four of those bytes being defined as the packet header. The payload portion of each packet is thus normally 184 bytes. However, a variable length adaptation field may be provided to extend the header, when required. When an adaptation field is present, the payload portion of the packet will be correspondingly shorter. Various timing and identification information is provided in different portions of the transport stream. These include a packet identifier (PID) found in the transport header of each transport packet to provide a reference number for identifying the transport packets carrying a specific service component. This number is included in a service definition or "service map" used by the receiver to identify those transport packets required to reconstruct a television program signal. The PID may also be referenced for various grooming and remultiplexing functions. In the case of video, audio or isochronous control data, the stream of packets labeled with a single PID represents a single video, audio or isochronous data service elementary stream, respectively. Each type of packet will have a different PID identifying the packet type.” Chen, 14:45–15:3. “As indicated in FIG. 6b, the transport header of each transport packet includes the PID which identifies the particular service component carried by the transport packet. The PCR will be carried by an adaptation field to provide timing information for a desired service. At a receiver, the PCR for the desired service is detected from the adaptation field. The PIDs of the transport packets will then be monitored in accordance with the timing established by the PCR to recover those transport packets carrying a particular component of the service to be processed.”

Exhibit C-8

Chen, 16:41–50.



Chen, Fig. 6b.

To the extent the preamble is limiting, Cunningham discloses this element and/or renders it obvious either alone or in combination with other references and/or the knowledge of one of ordinary skill in the art. *See, e.g.:*


“According to one aspect of this invention there is provided a method of encoding payload bits for transmission over a communications link such that the running disparity of the transmitted symbol stream is constrained within a predetermined limit, and said payload bits comprising either data or control information”

Cunningham, 1:58–63.

“Referring to FIG. 1, a data frame intended to be communicated over an optical fibre channel 20 typically comprises a binary digital message (payload) to be transferred between stations on the network; this message may comprise user data, or control information for managing operation of the network and its components. A circuit 12 generates an associated cyclic redundancy check (CRC) block containing check data derived from this message in known manner in accordance with a predetermined CRC algorithm, and appends the CRC block to the message.”

Cunningham, 2:65–3:7.

Exhibit C-8




Cunningham, Fig. 1.

“Control information is conveyed by using selected n-bit data block values, together with a distinctive value of 10 for the pair of label bits a and b to indicate that the associated data block represents control information rather than actual data. The data block values used for this purpose

Exhibit C-8

	are selected to be values which inherently have a disparity of zero, such as (for 8-bit blocks) 00001111 or 10101010. It will be noted that the label bit pair value 10 likewise has zero disparity, so the complete encoded block also has zero disparity.” Cunningham, 4:23–31. “The encoded blocks are transmitted over the optical fibre channel 20 (leftmost symbol first), and upon receipt at their destination they are detected by a receiver 22 where compensation is made in known manner for the effects of timing drift and other signal degradation to produce a restored digital data stream.” Cunningham, 4:45–50. To the extent the preamble is limiting, Hiroshima discloses this element and/or renders it obvious either alone or in combination with other references and/or the knowledge of one of ordinary skill in the art. <i>See, e.g.:</i> “FIG. 11 shows the details of the MPEG2-TS. Each packet of the MPEG2-TS has a prefix length of 188 bytes. The header 194 consists of four bytes and the payload 196 consists of 184 bytes. The header 194 is constructed by sync bytes 198, an error indicator 200, a unit start indicator 202, a transport priority 204, a packet ID 206, a scramble control 208, an adaptation field control 210, a continuity counter 212, and an adaptation field 214. The adaptation field 214 is constructed by an adaptation field length 216, a discontinuity indicator 218, a random access indicator 220, an elementary stream priority indicator 222, a 5-flag 224, an optional field 226, and stuffing bytes 228. Among them, the optional field 226 is constructed by a program clock reference 230, an original program clock reference field 232, a splice count down 234, a private data length 236, private data 238, an adaptation field extension 240, 3-flag 241, and an optional field 242.” Hiroshima, 10:51–67.
--	--

Exhibit C-8




Hiroshima, Fig. 11.

“FIG. 15 shows a process for forming the MPEG2-TS from the PES packet 82 by a multiplexing process in the multiplexer 34 in FIG. 5. First in step S1, the MPEG2-PES file is read. The TS header 194 in FIG. 11 is formed in step S2. The TS header 194 is formed in a manner such that the parameters such as error indicator 200 indicative of the presence or absence of the bit error in the packet, unit start indicator 202 indicating whether a new PES packet appears in the TS payload or not, packet ID 206 indicating whether the packet in the payload is video or audio, adaptation field 214, and adaptation field control 210 indicative of the presence or absence of the payload 196 are set in accordance with the state of the payload storage data of the 184-byte unit divided from the MPEG2-PES packet which had been read in step S1.”

Hiroshima, 11:13–27.

Exhibit C-8

FIG. 15



Hiroshima, Fig. 15.

To the extent the preamble is limiting, Haartsen discloses this element and/or renders it obvious either alone or in combination with other references and/or the knowledge of one of ordinary skill in the art. *See, e.g.:*

“The present invention fulfills the above-described and other needs by providing a multi-part digital preamble for use in transmitting digital data packets. Advantageously, a preamble according to the invention substantially reduces the overall length of DC-free sequences that are transmitted with each data packet in a digital FM system. According to exemplary embodiments, a multi-part

Exhibit C-8

	preamble includes a short, substantially DC-free leading part which enables coarse estimation of a prevailing DC level at an FM receiver. The exemplary preamble also includes a synchronization part which is not necessarily substantially DC-free and which carries timing and/or other useful information (e.g. , channel identification, destination address, etc.) that can vary from data packet to data packet (e.g. , depending on sender, recipient, time of day, etc.).” Haartsen, 5:2–15. “To alleviate this problem, the present invention discloses that a digital preamble can be strategically divided into three or more parts so that synchronization and other information can be included within the preamble while still enabling accurate DC estimation. Specifically, a short leading DC-free part of the preamble (e.g., alternating logical 1's and 0's, or any other sequence including an equal number of logical 1's and logical 0's) can be used to provide initial, coarse DC offset estimation. Thereafter, a synchronization part of the preamble, which is not necessarily DC free, but which is protected by code to prevent detection errors, can be used not only to synchronize detection of a data packet, but also to carry timing and/or other useful information (e.g., channel identification, destination address, time of day, user identification, etc.).” Haartsen, 12:13–13:2. “An exemplary multi-part preamble according to the invention is depicted in Figure 5. In the figure, a three-part preamble 500 is shown to include a DC-free leading sequence 510, followed by a multiple-bit synchronization word 520 which is not necessarily DC free and which can be different for each data packet. The synchronization word 520, which can include timing and/or other information and which is coded to protect against detection errors, is in turn followed by a trailing DC-free sequence 530. Although the leading and trailing DC-free sequences 510, 530 are shown in the figure as 4-bit sequences of alternating logical 1's and 0's, those of ordinary skill will appreciate that each DC-free sequence 510, 530 can, in practice, be shorter or longer in length and can include any DC-free sequence having an equal number of logical 1's and logical 0's.” Haartsen, 13:13–14:2.
--	--

Exhibit C-8




Figure 5

Haartsen, Fig. 5.

To the extent the preamble is limiting, Austin discloses this element and/or renders it obvious either alone or in combination with other references and/or the knowledge of one of ordinary skill in the art. *See, e.g.:*

“This system uses up to 5 pre-defined data packet sizes and up to 16 user defined data message types within each packet. Each packet has Forward Error Connection (FEC) and bit interleaving applied to it on transmission. Error checking is performed using a standard CRC. The system uses standard run in and byte synchronisation techniques. The byte synchronisation word used is 32 _bits long and there are 5_different words used, one to identify each of the different packet types.”

Austin, 4:37-45.

To the extent that Plaintiff contends that the Additional Prior Art References do not explicitly or inherently disclose this element, or render obvious this element, this element is obvious in combination with one or more references and/or with the knowledge of a person of ordinary skill in the art, as set forth in Defendant’s Invalidity Contentions.

11(a) receiving a synchronization symbol by a receiver indicating a packet type, each packet type having a

Chen discloses and/or renders obvious this limitation. *See, e.g.:*

“FIG. 1 is a block diagram of a Fiber-to-the-Curb Digital Video Distribution Network in accordance with the present invention. The distribution network includes Video Information Providers (VIPs) shown generally at 100, a gateway, shown generally at 120, an access network, shown generally at

Exhibit C-8

different synchronization symbol;

140, and the home of a Video Information User (VIU), shown generally at 160. The gateway 120 may comprise an LI Gateway manufactured by Teleco. The VIPs may include a video server 102, an interactive data server 104, and an internet gateway 106, all of which communicate with an asynchronous transfer mode (ATM) network 122 via, for example, an OC-3 system, where optical fiber is used to transmit data at 155.52 Mbit/sec. The ATM network 122 may also receive data via a satellite receiving antenna 110, broadcast encoder and server 112, and ATM multiplexer 114. The ATM network 122 communicates with a session manager 123.”

Chen, 4:11–28.




FIG. 1

Chen, Fig. 1.

“FIG. 4 is a block diagram of an Insertion Processing Unit (IPU) of a Digital Ad Insertion Module in accordance with the present invention. The IPU, shown generally at 400, receives the main stream at a packet framer 405, while the insertion stream is received at another packet framer 410. After

Exhibit C-8

processing by the packet framer 405, the main stream is provided to a Main Stream Parser (MSP) 415.”

Chen, 5:62–6:1.




FIG.4

Chen, Fig. 4.

“The main stream parser 415 provides the main stream data to a main stream buffer 480, which communicates with a syntax processor 470. The main stream buffer 480 is part of the output buffer 475. The insertion stream parser 420 provides the insertion stream data to a Program Specific Information (PSI)/Program Identifier (PID) Replacer 435, which replaces the PSI tables and PIDs of the insertion stream with those of the main stream. The PSI/PID replacer 435 retrieves the relevant information from the main stream using a communication path which is not shown. The PSI and PID data of the main stream provide the information to connect the various audio, video, and data packets to a particular programming service. Specifically, this is accomplished using a Program Association Table (PAT) and a Program Map Table (PMT). The insertion stream is then

Exhibit C-8

	provided to an Insertion Stream Buffer (ISB) 490, which is part of the output buffer 475. The ISB 490 communicates with a syntax processor 470. Furthermore, the output buffer 475 receives commands from a buffer controller 465 to provide an output stream with the insertion stream seamlessly spliced into the main stream.” Chen, 6:33–54. “FIG. 5 illustrates the processing flow of the Insertion Processing Unit of FIG. 4 in accordance with the present invention. Note that the steps of FIG. 5 are shown as occurring in a sequential manner for clarity, but many of the steps may occur concurrently and/or sequentially. At block 502, the IPU outputs the main transport stream from the main stream buffer (MSB) 480. The IPU is essentially in a bypass mode at this time since the insertion stream has not yet been inserted. At block 506, the main stream parser parses the main stream to locate the splicing signal T-- in. Note that T-- in need not be carried in the main stream, but may be provided locally, for example, at a cable system headend. When T-- in is detected, the IPU continues to output the main transport stream from the main stream buffer. At block 508, the main stream parser parses the main stream to locate the closest packet with an anchor picture (e.g., I or P picture) start code and passes that packet to the syntax processor 470. This packet will be the last packet of the main stream. To ensure that the resulting data stream is MPEG compliant, no data from the next frame or sequence of frames may be included in this last main stream packet. At block 512, the syntax processor adds a sequence end code field to the last packet of the main stream. The sequence end code is a twenty-four bit code which indicates the end of a sequence of pictures. At block 514, the PSI/PID replacer 435 starts to replace the PSI/PID of the insertion stream.” Chen, 12:28–53.
--	--

Exhibit C-8




FIG.5

Chen, Fig. 5.

“In the MPEG-2 system (and the similar DigiCipher® II system proprietary to General Instrument Corporation, the assignee hereof) a transport stream, or transport multiplex is made up of a


Exhibit C-8

	contiguous set of fixed length packets. Each packet is 188 total bytes in length, with the first four of those bytes being defined as the packet header. The payload portion of each packet is thus normally 184 bytes. However, a variable length adaptation field may be provided to extend the header, when required. When an adaptation field is present, the payload portion of the packet will be correspondingly shorter. Various timing and identification information is provided in different portions of the transport stream. These include a packet identifier (PID) found in the transport header of each transport packet to provide a reference number for identifying the transport packets carrying a specific service component. This number is included in a service definition or "service map" used by the receiver to identify those transport packets required to reconstruct a television program signal. The PID may also be referenced for various grooming and remultiplexing functions. In the case of video, audio or isochronous control data, the stream of packets labeled with a single PID represents a single video, audio or isochronous data service elementary stream, respectively. Each type of packet will have a different PID identifying the packet type. Timing information carried by the transport stream includes a program clock reference (PCR) which effectively represents a sample of the system time clock (STC) time base that underlies the service composed of the PIDs referenced in the service map. The PID carrying the packet with the PCR is also referenced in the service map. The video, audio and isochronous data components of a service are locked through a defined relationship to the system time clock. The PCR serves to define the transport rate, in the sense that between any two successive PCRs in one PID, the transport rate is constant and nominally equal to the system time clock rate times the ratio of the total number of transport bytes between the PCRs divided by the difference in the PCRs in units of system time clock ticks." Chen, 14:45–15:17. "As indicated in FIG. 6b, the transport header of each transport packet includes the PID which identifies the particular service component carried by the transport packet. The PCR will be carried by an adaptation field to provide timing information for a desired service. At a receiver, the PCR for the desired service is detected from the adaptation field. The PIDs of the transport packets will
--	---

Exhibit C-8

then be monitored in accordance with the timing established by the PCR to recover those transport packets carrying a particular component of the service to be processed.”

Chen, 16:41–50.




Chen, Fig. 6b.

Cunningham discloses and/or renders obvious this limitation. *See, e.g.:*

“Referring to FIG. 1, a data frame intended to be communicated over an optical fibre channel 20 typically comprises a binary digital message (payload) to be transferred between stations on the network; this message may comprise user data, or control information for managing operation of the network and its components. A circuit 12 generates an associated cyclic redundancy check (CRC) block containing check data derived from this message in known manner in accordance with a predetermined CRC algorithm, and appends the CRC block to the message.”

Cunningham, 2:65–3:7.

Exhibit C-8



Cunningham, Fig. 1.

“Control information is conveyed by using selected n-bit data block values, together with a distinctive value of 10 for the pair of label bits a and b to indicate that the associated data block represents control information rather than actual data. The data block values used for this purpose are selected to be values which inherently have a disparity of zero, such as (for 8-bit blocks)

Exhibit C-8


	00001111 or 10101010. It will be noted that the label bit pair value 10 likewise has zero disparity, so the complete encoded block also has zero disparity.” Cunningham, 4:23–31. “The encoded blocks are transmitted over the optical fibre channel 20 (leftmost symbol first), and upon receipt at their destination they are detected by a receiver 22 where compensation is made in known manner for the effects of timing drift and other signal degradation to produce a restored digital data stream.” Cunningham, 4:45–50. “There are many potential ways to achieve code block synchronization (i.e. correct alignment in the receiver of the label and payload fields). For example, one technique would be to transmit a ‘comma sequence’ during start-up. As an example of this technique consider a code with n=8 using the simple control payload scheme. A training control state of 1111 1111 would be transmitted as: [10 1111 1111 0000 0000][10 1111 1111 0000 0000][10 1111 1111 0000 0000] The receiver searches for the bit pair 10. This is found in only two locations: the correct location (the label field), and at the end of the run of eight 1's. However, if the receiver (incorrectly) identifies the label field to be the 10 bit pair at the end of the run of 1's, the following 16 bits are not a valid control payload (first 8 bits are not the complement of the second 8 bits). The receiver can therefore detect the mistake. It is also noted that this sequence contains a regular 010 pattern, which aids in phase-locked loop (PLL) clock acquisition.” Cunningham, 7:5–23. Hiroshima discloses and/or renders obvious this limitation. <i>See, e.g.:</i> “FIG. 4 is a block diagram of the first embodiment of the invention in which the moving picture stream conversion in mode 1 is executed. A moving picture stream converting apparatus of FIG. 4 comprises a demultiplexer 30, a packetizer 32, and a multiplexer 34. The demultiplexer 30 receives an MPEG1 system stream as a conversion source and outputs elementary streams which were
--	--

Exhibit C-8

separated into video and audio streams. The packetizer 32 forms packets constructing a packetized elementary stream of MPEG2 with respect to each of the video and audio elementary streams which are supplied from the demultiplexer 30. The packet is usually called a PES packet. The multiplexer 34 divides each of the video and audio PES packets into a TS packet having a fixed length of 188 bytes and subsequently multiplexes, and outputs the resultant data as an MPEG2-TS.”

Hiroshima, 7:59–8:7.

FIG. 4




Hiroshima, Fig. 4.

“FIG. 11 shows the details of the MPEG2-TS. Each packet of the MPEG2-TS has a prefix length of 188 bytes. The header 194 consists of four bytes and the payload 196 consists of 184 bytes. The header 194 is constructed by sync bytes 198, an error indicator 200, a unit start indicator 202, a transport priority 204, a packet ID 206, a scramble control 208, an adaptation field control 210, a continuity counter 212, and an adaptation field 214. The adaptation field 214 is constructed by an adaptation field length 216, a discontinuity indicator 218, a random access indicator 220, an elementary stream priority indicator 222, a 5-flag 224, an optional field 226, and stuffing bytes 228. Among them, the optional field 226 is constructed by a program clock reference 230, an original program clock reference field 232, a splice count down 234, a private data length 236, private data 238, an adaptation field extension 240, 3-flag 241, and an optional field 242.”

Hiroshima, 10:51–67.

Exhibit C-8




Hiroshima, Fig. 11.

“FIG. 15 shows a process for forming the MPEG2-TS from the PES packet 82 by a multiplexing process in the multiplexer 34 in FIG. 5. First in step S1, the MPEG2-PES file is read. The TS header 194 in FIG. 11 is formed in step S2. The TS header 194 is formed in a manner such that the parameters such as error indicator 200 indicative of the presence or absence of the bit error in the packet, unit start indicator 202 indicating whether a new PES packet appears in the TS payload or not, packet ID 206 indicating whether the packet in the payload is video or audio, adaptation field 214, and adaptation field control 210 indicative of the presence or absence of the payload 196 are set in accordance with the state of the payload storage data of the 184-byte unit divided from the MPEG2-PES packet which had been read in step S1.”

Hiroshima, 11:13–27.

Exhibit C-8

FIG. 15



Hiroshima, Fig. 15.

Haartsen discloses and/or renders obvious this limitation. *See, e.g.:*

“The present invention fulfills the above-described and other needs by providing a multi-part digital preamble for use in transmitting digital data packets. Advantageously, a preamble according to the invention substantially reduces the overall length of DC-free sequences that are transmitted with each data packet in a digital FM system. According to exemplary embodiments, a multi-part preamble includes a short, substantially DC-free leading part which enables coarse estimation of a prevailing DC level at an FM receiver. The exemplary preamble also includes a synchronization


Exhibit C-8

part which is not necessarily substantially DC-free and which carries timing and/or other useful information (e.g. , channel identification, destination address, etc.) that can vary from data packet to data packet (e.g. , depending on sender, recipient, time of day, etc.).”

Haartsen, 5:2–15.

“An exemplary DC estimation processor is depicted in Figure 3, wherein an exemplary FM demodulation processor 300 is shown to include a correlation processor 310, a resistor 320, a switch 330 and a capacitor 340, as well as the FM detector 210, the low-pass filter 130 and the slicer 230 of Figure 2.”

Haartsen, 9:11–14.



Haartsen, Fig. 3.

“To alleviate this problem, the present invention discloses that a digital preamble can be strategically divided into three or more parts so that synchronization and other information can be included within the preamble while still enabling accurate DC estimation. Specifically, a short leading DC-free part of the preamble (e.g., alternating logical 1's and 0's, or any other sequence including an equal number of logical 1's and logical 0's) can be used to provide initial, coarse DC offset estimation. Thereafter, a synchronization part of the preamble, which is not necessarily DC free, but which is protected by code to prevent detection errors, can be used not only to synchronize


Exhibit C-8

detection of a data packet, but also to carry timing and/or other useful information (e.g., channel identification, destination address, time of day, user identification, etc.).”

Haartsen, 12:13–13:2.

“An exemplary multi-part preamble according to the invention is depicted in Figure 5. In the figure, a three-part preamble 500 is shown to include a DC-free leading sequence 510, followed by a multiple-bit synchronization word 520 which is not necessarily DC free and which can be different for each data packet. The synchronization word 520, which can include timing and/or other information and which is coded to protect against detection errors, is in turn followed by a trailing DC-free sequence 530. Although the leading and trailing DC-free sequences 510, 530 are shown in the figure as 4-bit sequences of alternating logical 1's and 0's, those of ordinary skill will appreciate that each DC-free sequence 510, 530 can, in practice, be shorter or longer in length and can include any DC-free sequence having an equal number of logical 1's and logical 0's.”

Haartsen, 13:13–14:2.



Haartsen, Fig. 5.

“When the exemplary preamble 500 is received by the system 300 of Figure 3, the leading DC-free sequence 510 induces a coarse estimate of the baseband DC offset on the storage capacitor 340. Although the estimate is coarse, it is sufficient to permit successful detection of the coded synchronization word 520 (e.g., by the correlation processor 310). Methods for coding of a synchronization word to protect against erroneous detection in the presence of some DC-offset are

Exhibit C-8

	known to those of skill in the art. Once the correlation processor 310 has detected the synchronization word 520, timing between received bits is well known. Consequently, the correlation processor 310 can open the switch 330 precisely at the end of the trailing DC-free sequence 530 to thereby store an accurate estimate of the DC offset for reception and detection of the incoming digital data packet.” Haartsen, 14:3–14. Shin discloses and/or renders obvious this limitation. <i>See, e.g.:</i> “It may thus be appreciated that each eight-bit code converted into a 10-bit encoded character during high-transition mode operation includes four or more logical transitions. Similarly, each eight-bit code converted into a 10-bit encoded character during low-transition mode operation includes less than four logical transitions. These sets of 10-bit encoded characters capable of being produced during high-transition and low-transition modes of operation may be characterized as a high-transition set of "in-band" encoded characters, and a low-transition set of in-band encoded characters, respectively. Beyond the 256 in-band characters within the high-transition set and the 256 in-band characters of the low-transition set, there exists a high-transition set of 256 out-of-band 10-bit characters and a low-transition set of 10-bit characters. In accordance with another aspect of the invention, various synchronization and other special characters are defined using the high-transition and low-transition sets of out-of-band characters. Each character corresponding to one of these "out-of-band" characters associated with the high-transition set includes less than four logical transitions, and each of the out-of-band characters associated with the low-transition set of code characters includes more than four logical transitions. The difference in the number of transitions between in-band and out-of-band characters allows selected out-of-band characters to serve as control characters, and to be readily distinguished from in-band characters within the transmitted data stream.” Shin, 9:49–10:8. “In accordance with one aspect of the invention, the synchronization characters associated with both the low-transition and high-transition sets of code characters are selected so as to facilitate rapid synchronization during data recovery. When the low-transition set of code characters is being
--	---

Exhibit C-8

employed, a special group of out-of-band characters is used during synchronization. Each special synchronization character includes a predefined number larger than four (e.g., 7) of logical transitions between binary character values, and a predefined number (e.g., two) of "non-transitions" between character values. As will be discussed below, the special synchronization characters are selected such that random logic may be used to distinguish each special synchronization character from the in-band characters of the low-transition set."

Shin, 10:17–34.

"Referring to FIG. 6, the deserializer 34 receives the sequence of 10-bit characters T[9:0] produced by the encoder and generates 10-bit parallel received data upon bit lines RX9, RX8, . . . , RX0 (i.e., RX[9:0]). This 10-bit parallel received data is provided over bit lines RX[9:0] to the decoder 44, as well as to a decoder synchronization module 114. As is described below in section IV, the synchronization module 114 is operative to ascertain boundaries within the 10-bit parallel received data corresponding to the frame boundaries of the transmitted data (i.e., to T[9:0]). Specifically, synchronization module 114 determines upon which of the bit lines RX[9:0] the deserializer 34 is providing the received bits corresponding to the first bit T[0] of each transmitted byte T[9:0]. Upon making this determination, the synchronization module 114 provides a frame boundary pointer 118 to decoder 44 identifying the one of the bit lines RX[9:0] corresponding to the first bit T[0] of each transmitted 10-bit character T[9:0]. Upon receiving this synchronization information, the decoder 44 is disposed to decode the received data RX[9:0] in the following manner."

Shin, 13:37–57.




FIG. 6

Shin, Fig. 6.

Exhibit C-8

		Austin discloses and/or renders obvious this limitation. <i>See, e.g.:</i> “This system uses up to 5 pre-defined data packet sizes and up to 16 user defined data message types within each packet. Each packet has Forward Error Connection (FEC) and bit interleaving applied to it on transmission. Error checking is performed using a standard CRC. The system uses standard run in and byte synchronisation techniques. The byte synchronisation word used is 32 _bits long and there are 5_different words used, one to identify each of the different packet types.” Austin, 4:37-45. To the extent that Plaintiff contends that the Additional Prior Art References do not explicitly or inherently disclose this element, or render obvious this element, this element is obvious in combination with one or more references and/or with the knowledge of a person of ordinary skill in the art, as set forth in Defendant’s Invalidity Contentions.
--	--	--

Exhibit C-8

11(b)	receiving a packet of symbols;	Chen discloses and/or renders obvious this limitation. <i>See, e.g.:</i> “FIG. 1 is a block diagram of a Fiber-to-the-Curb Digital Video Distribution Network in accordance with the present invention. The distribution network includes Video Information Providers (VIPs) shown generally at 100, a gateway, shown generally at 120, an access network, shown generally at 140, and the home of a Video Information User (VIU), shown generally at 160. The gateway 120 may comprise an LI Gateway manufactured by Teleco. The VIPs may include a video server 102, an interactive data server 104, and an internet gateway 106, all of which communicate with an asynchronous transfer mode (ATM) network 122 via, for example, an OC-3 system, where optical fiber is used to transmit data at 155.52 Mbit/sec. The ATM network 122 may also receive data via a satellite receiving antenna 110, broadcast encoder and server 112, and ATM multiplexer 114. The ATM network 122 communicates with a session manager 123.” Chen, 4:11–28. FIG. 1
--------------	---------------------------------------	---

Exhibit C-8

Chen, Fig. 1.

“FIG. 4 is a block diagram of an Insertion Processing Unit (IPU) of a Digital Ad Insertion Module in accordance with the present invention. The IPU, shown generally at 400, receives the main stream at a packet framer 405, while the insertion stream is received at another packet framer 410. After processing by the packet framer 405, the main stream is provided to a Main Stream Parser (MSP) 415.”

Chen, 5:62–6:1.




FIG.4


Chen, Fig. 4.

“The main stream parser 415 provides the main stream data to a main stream buffer 480, which communicates with a syntax processor 470. The main stream buffer 480 is part of the output buffer 475. The insertion stream parser 420 provides the insertion stream data to a Program Specific Information (PSI)/Program Identifier (PID) Replacer 435, which replaces the PSI tables and PIDs

Exhibit C-8

	of the insertion stream with those of the main stream. The PSI/PID replacer 435 retrieves the relevant information from the main stream using a communication path which is not shown. The PSI and PID data of the main stream provide the information to connect the various audio, video, and data packets to a particular programming service. Specifically, this is accomplished using a Program Association Table (PAT) and a Program Map Table (PMT). The insertion stream is then provided to an Insertion Stream Buffer (ISB) 490, which is part of the output buffer 475. The ISB 490 communicates with a syntax processor 470. Furthermore, the output buffer 475 receives commands from a buffer controller 465 to provide an output stream with the insertion stream seamlessly spliced into the main stream.” Chen, 6:33–54. “FIG. 5 illustrates the processing flow of the Insertion Processing Unit of FIG. 4 in accordance with the present invention. Note that the steps of FIG. 5 are shown as occurring in a sequential manner for clarity, but many of the steps may occur concurrently and/or sequentially. At block 502, the IPU outputs the main transport stream from the main stream buffer (MSB) 480. The IPU is essentially in a bypass mode at this time since the insertion stream has not yet been inserted. At block 506, the main stream parser parses the main stream to locate the splicing signal T-- in. Note that T-- in need not be carried in the main stream, but may be provided locally, for example, at a cable system headend. When T-- in is detected, the IPU continues to output the main transport stream from the main stream buffer. At block 508, the main stream parser parses the main stream to locate the closest packet with an anchor picture (e.g., I or P picture) start code and passes that packet to the syntax processor 470. This packet will be the last packet of the main stream. To ensure that the resulting data stream is MPEG compliant, no data from the next frame or sequence of frames may be included in this last main stream packet. At block 512, the syntax processor adds a sequence end code field to the last packet of the main stream. The sequence end code is a twenty-four bit code which indicates the end of a sequence of pictures. At block 514, the PSI/PID replacer 435 starts to replace the PSI/PID of the insertion stream.” Chen, 12:28–53.
--	---

Exhibit C-8




Chen, Fig. 5.

“In the MPEG-2 system (and the similar DigiCipher® II system proprietary to General Instrument Corporation, the assignee hereof) a transport stream, or transport multiplex is made up of a

Exhibit C-8

	contiguous set of fixed length packets. Each packet is 188 total bytes in length, with the first four of those bytes being defined as the packet header. The payload portion of each packet is thus normally 184 bytes. However, a variable length adaptation field may be provided to extend the header, when required. When an adaptation field is present, the payload portion of the packet will be correspondingly shorter.” Chen, 14:45–55. Cunningham discloses and/or renders obvious this limitation. <i>See, e.g.:</i> “Referring to FIG. 1, a data frame intended to be communicated over an optical fibre channel 20 typically comprises a binary digital message (payload) to be transferred between stations on the network; this message may comprise user data, or control information for managing operation of the network and its components. A circuit 12 generates an associated cyclic redundancy check (CRC) block containing check data derived from this message in known manner in accordance with a predetermined CRC algorithm, and appends the CRC block to the message.” Cunningham, 2:65–3:7.
--	---

Exhibit C-8



Cunningham, Fig. 1.

“The encoded blocks are transmitted over the optical fibre channel 20 (leftmost symbol first), and upon receipt at their destination they are detected by a receiver 22 where compensation is made in known manner for the effects of timing drift and other signal degradation to produce a restored digital data stream.”

Exhibit C-8


Cunningham, 4:45–50.

Hiroshima discloses and/or renders obvious this limitation. *See, e.g.:*

“FIG. 4 is a block diagram of the first embodiment of the invention in which the moving picture stream conversion in mode 1 is executed. A moving picture stream converting apparatus of FIG. 4 comprises a demultiplexer 30, a packetizer 32, and a multiplexer 34. The demultiplexer 30 receives an MPEG1 system stream as a conversion source and outputs elementary streams which were separated into video and audio streams. The packetizer 32 forms packets constructing a packetized elementary stream of MPEG2 with respect to each of the video and audio elementary streams which are supplied from the demultiplexer 30. The packet is usually called a PES packet. The multiplexer 34 divides each of the video and audio PES packets into a TS packet having a fixed length of 188 bytes and subsequently multiplexes, and outputs the resultant data as an MPEG2-TS.”

Hiroshima, 7:59–8:7.

FIG. 4




Hiroshima, Fig. 4.

“FIG. 11 shows the details of the MPEG2-TS. Each packet of the MPEG2-TS has a prefix length of 188 bytes. The header 194 consists of four bytes and the payload 196 consists of 184 bytes. The header 194 is constructed by sync bytes 198, an error indicator 200, a unit start indicator 202, a transport priority 204, a packet ID 206, a scramble control 208, an adaptation field control 210, a

Exhibit C-8

continuity counter 212, and an adaptation field 214. The adaptation field 214 is constructed by an adaptation field length 216, a discontinuity indicator 218, a random access indicator 220, an elementary stream priority indicator 222, a 5-flag 224, an optional field 226, and stuffing bytes 228. Among them, the optional field 226 is constructed by a program clock reference 230, an original program clock reference field 232, a splice count down 234, a private data length 236, private data 238, an adaptation field extension 240, 3-flag 241, and an optional field 242.”

Hiroshima, 10:51–67.



Hiroshima, Fig. 11.


“FIG. 15 shows a process for forming the MPEG2-TS from the PES packet 82 by a multiplexing process in the multiplexer 34 in FIG. 5. First in step S1, the MPEG2-PES file is read. The TS header 194 in FIG. 11 is formed in step S2. The TS header 194 is formed in a manner such that the

Exhibit C-8

parameters such as error indicator 200 indicative of the presence or absence of the bit error in the packet, unit start indicator 202 indicating whether a new PES packet appears in the TS payload or not, packet ID 206 indicating whether the packet in the payload is video or audio, adaptation field 214, and adaptation field control 210 indicative of the presence or absence of the payload 196 are set in accordance with the state of the payload storage data of the 184-byte unit divided from the MPEG2-PES packet which had been read in step S1.”

Hiroshima, 11:13–27.

FIG. 15




Hiroshima, Fig. 15.

Exhibit C-8

Austin discloses and/or renders obvious this limitation. *See, e.g.:*

“A suitable receiver is shown in FIG. 4. The receiver 50 is a non-coherent FFSK matched filter receiver. The receiver 50 includes amplifier 51, low pass filter 52 and automatic gain control 53 and through the band pass filter 54. The filter output is passed via amplifier 55 to detector 56. This allows the analog signal from the pipe to be received and the digital data recovered into streams F1 and F2. Matched filters 57 and 58 and summer 59 provide the recovered (uncloacked) data at its output. Separate matched filter paths are used for each tone and data decisions are made by comparison of the resultant energy output from each filter. Symbol recovery is derived from the data decision stream. Each matched filter is implemented as an integration of the incoming data stream multiplied by a sine wave at the respective sub-carrier frequency (over each successive study symbol period).”

Austin, 2:57-3:5.



Austin, Fig. 4.

Exhibit C-8

“When digital implementation is required a receiver can be constructed as shown in FIG. 7. This uses matched filters and employs a series of correlators. The analog first stage output is received by summing amplifier 81 and converted into digital form via analog to digital converter 82. The dual path filtering 57 and 58 is provided as before now in digital form. As shown in the upper (expanded) path a sampler 83 with delays 84 and adders 85, dividers 86 and summer 87 to path sampler 88 and comparator 89. Hence, the received analog signal is amplified, filtered and fed to an automatic gain control loop. Typically the LPF has a -3 dB cut off at 20 HZ. The AGC uses a bandpass filter with -3 dB points at 7.5 HZ and 17.5 HZ to select the band of interest for detection. The level detector input is used to adjust the gain of the loop so that the output AC peak to peak signal most closely fits the peak to peak full scale (FS) limit of the digital convertor (ADC). The ADC gives typically 10 bit resolution. F1 and F2 paths are concerned respectively with each matched filter.”

Austin, 3:25-43.




Exhibit C-8

		Austin, Fig. 7. “This system uses up to 5 pre-defined data packet sizes and up to 16 user defined data message types within each packet. Each packet has Forward Error Connection (FEC) and bit interleaving applied to it on transmission. Error checking is performed using a standard CRC. The system uses standard run in and byte synchronisation techniques. The byte synchronisation word used is 32 _bits long and there are 5_different words used, one to identify each of the different packet types.” Austin, 4:37-45. To the extent that Plaintiff contends that the Additional Prior Art References do not explicitly or inherently disclose this element, or render obvious this element, this element is obvious in combination with one or more references and/or with the knowledge of a person of ordinary skill in the art, as set forth in Defendant’s Invalidity Contentions.
11(c)	indicating that the received packet of symbols has the packet type of the received synchronization symbol; and	Chen discloses and/or renders obvious this limitation. <i>See, e.g.:</i> “FIG. 4 is a block diagram of an Insertion Processing Unit (IPU) of a Digital Ad Insertion Module in accordance with the present invention. The IPU, shown generally at 400, receives the main stream at a packet framer 405, while the insertion stream is received at another packet framer 410. After processing by the packet framer 405, the main stream is provided to a Main Stream Parser (MSP) 415.” Chen, 5:62–6:1.

Exhibit C-8




FIG.4

Chen, Fig. 4.

“The main stream parser 415 provides the main stream data to a main stream buffer 480, which communicates with a syntax processor 470. The main stream buffer 480 is part of the output buffer 475. The insertion stream parser 420 provides the insertion stream data to a Program Specific Information (PSI)/Program Identifier (PID) Replacer 435, which replaces the PSI tables and PIDs of the insertion stream with those of the main stream. The PSI/PID replacer 435 retrieves the relevant information from the main stream using a communication path which is not shown. The PSI and PID data of the main stream provide the information to connect the various audio, video, and data packets to a particular programming service. Specifically, this is accomplished using a Program Association Table (PAT) and a Program Map Table (PMT). The insertion stream is then provided to an Insertion Stream Buffer (ISB) 490, which is part of the output buffer 475. The ISB 490 communicates with a syntax processor 470. Furthermore, the output buffer 475 receives commands from a buffer controller 465 to provide an output stream with the insertion stream seamlessly spliced into the main stream.”

Exhibit C-8

	Chen, 6:33–54. “FIG. 5 illustrates the processing flow of the Insertion Processing Unit of FIG. 4 in accordance with the present invention. Note that the steps of FIG. 5 are shown as occurring in a sequential manner for clarity, but many of the steps may occur concurrently and/or sequentially. At block 502, the IPU outputs the main transport stream from the main stream buffer (MSB) 480. The IPU is essentially in a bypass mode at this time since the insertion stream has not yet been inserted. At block 506, the main stream parser parses the main stream to locate the splicing signal T-- in. Note that T-- in need not be carried in the main stream, but may be provided locally, for example, at a cable system headend. When T-- in is detected, the IPU continues to output the main transport stream from the main stream buffer. At block 508, the main stream parser parses the main stream to locate the closest packet with an anchor picture (e.g., I or P picture) start code and passes that packet to the syntax processor 470. This packet will be the last packet of the main stream. To ensure that the resulting data stream is MPEG compliant, no data from the next frame or sequence of frames may be included in this last main stream packet. At block 512, the syntax processor adds a sequence end code field to the last packet of the main stream. The sequence end code is a twenty-four bit code which indicates the end of a sequence of pictures. At block 514, the PSI/PID replacer 435 starts to replace the PSI/PID of the insertion stream.” Chen, 12:28–53.
--	---

Exhibit C-8




FIG.5

Chen, Fig. 5.

“In the MPEG-2 system (and the similar DigiCipher® II system proprietary to General Instrument Corporation, the assignee hereof) a transport stream, or transport multiplex is made up of a


Exhibit C-8

	contiguous set of fixed length packets. Each packet is 188 total bytes in length, with the first four of those bytes being defined as the packet header. The payload portion of each packet is thus normally 184 bytes. However, a variable length adaptation field may be provided to extend the header, when required. When an adaptation field is present, the payload portion of the packet will be correspondingly shorter. Various timing and identification information is provided in different portions of the transport stream. These include a packet identifier (PID) found in the transport header of each transport packet to provide a reference number for identifying the transport packets carrying a specific service component. This number is included in a service definition or "service map" used by the receiver to identify those transport packets required to reconstruct a television program signal. The PID may also be referenced for various grooming and remultiplexing functions. In the case of video, audio or isochronous control data, the stream of packets labeled with a single PID represents a single video, audio or isochronous data service elementary stream, respectively. Each type of packet will have a different PID identifying the packet type. Timing information carried by the transport stream includes a program clock reference (PCR) which effectively represents a sample of the system time clock (STC) time base that underlies the service composed of the PIDs referenced in the service map. The PID carrying the packet with the PCR is also referenced in the service map. The video, audio and isochronous data components of a service are locked through a defined relationship to the system time clock. The PCR serves to define the transport rate, in the sense that between any two successive PCRs in one PID, the transport rate is constant and nominally equal to the system time clock rate times the ratio of the total number of transport bytes between the PCRs divided by the difference in the PCRs in units of system time clock ticks." Chen, 14:45–15:17. "As indicated in FIG. 6b, the transport header of each transport packet includes the PID which identifies the particular service component carried by the transport packet. The PCR will be carried by an adaptation field to provide timing information for a desired service. At a receiver, the PCR for the desired service is detected from the adaptation field. The PIDs of the transport packets will
--	---

Exhibit C-8

then be monitored in accordance with the timing established by the PCR to recover those transport packets carrying a particular component of the service to be processed.”

Chen, 16:41–50.




Chen, Fig. 6b.

Cunningham discloses and/or renders obvious this limitation. *See, e.g.:*

“Referring to FIG. 1, a data frame intended to be communicated over an optical fibre channel 20 typically comprises a binary digital message (payload) to be transferred between stations on the network; this message may comprise user data, or control information for managing operation of the network and its components. A circuit 12 generates an associated cyclic redundancy check (CRC) block containing check data derived from this message in known manner in accordance with a predetermined CRC algorithm, and appends the CRC block to the message.”

Cunningham, 2:65–3:7.

Exhibit C-8



Cunningham, Fig. 1.

“Control information is conveyed by using selected n-bit data block values, together with a distinctive value of 10 for the pair of label bits a and b to indicate that the associated data block represents control information rather than actual data. The data block values used for this purpose are selected to be values which inherently have a disparity of zero, such as (for 8-bit blocks)


Exhibit C-8

	00001111 or 10101010. It will be noted that the label bit pair value 10 likewise has zero disparity, so the complete encoded block also has zero disparity.” Cunningham, 4:23–31. “The encoded blocks are transmitted over the optical fibre channel 20 (leftmost symbol first), and upon receipt at their destination they are detected by a receiver 22 where compensation is made in known manner for the effects of timing drift and other signal degradation to produce a restored digital data stream.” Cunningham, 4:45–50. “There are many potential ways to achieve code block synchronization (i.e. correct alignment in the receiver of the label and payload fields). For example, one technique would be to transmit a ‘comma sequence’ during start-up. As an example of this technique consider a code with n=8 using the simple control payload scheme. A training control state of 1111 1111 would be transmitted as: [10 1111 1111 0000 0000][10 1111 1111 0000 0000][10 1111 1111 0000 0000] The receiver searches for the bit pair 10. This is found in only two locations: the correct location (the label field), and at the end of the run of eight 1's. However, if the receiver (incorrectly) identifies the label field to be the 10 bit pair at the end of the run of 1's, the following 16 bits are not a valid control payload (first 8 bits are not the complement of the second 8 bits). The receiver can therefore detect the mistake. It is also noted that this sequence contains a regular 010 pattern, which aids in phase-locked loop (PLL) clock acquisition.” Cunningham, 7:5–23. Hiroshima discloses and/or renders obvious this limitation. <i>See, e.g.:</i> “FIG. 11 shows the details of the MPEG2-TS. Each packet of the MPEG2-TS has a prefix length of 188 bytes. The header 194 consists of four bytes and the payload 196 consists of 184 bytes. The header 194 is constructed by sync bytes 198, an error indicator 200, a unit start indicator 202, a transport priority 204, a packet ID 206, a scramble control 208, an adaptation field control 210, a
--	---

Exhibit C-8

continuity counter 212, and an adaptation field 214. The adaptation field 214 is constructed by an adaptation field length 216, a discontinuity indicator 218, a random access indicator 220, an elementary stream priority indicator 222, a 5-flag 224, an optional field 226, and stuffing bytes 228. Among them, the optional field 226 is constructed by a program clock reference 230, an original program clock reference field 232, a splice count down 234, a private data length 236, private data 238, an adaptation field extension 240, 3-flag 241, and an optional field 242.”

Hiroshima, 10:51–67.



Hiroshima, Fig. 11.


“FIG. 15 shows a process for forming the MPEG2-TS from the PES packet 82 by a multiplexing process in the multiplexer 34 in FIG. 5. First in step S1, the MPEG2-PES file is read. The TS header 194 in FIG. 11 is formed in step S2. The TS header 194 is formed in a manner such that the

Exhibit C-8

parameters such as error indicator 200 indicative of the presence or absence of the bit error in the packet, unit start indicator 202 indicating whether a new PES packet appears in the TS payload or not, packet ID 206 indicating whether the packet in the payload is video or audio, adaptation field 214, and adaptation field control 210 indicative of the presence or absence of the payload 196 are set in accordance with the state of the payload storage data of the 184-byte unit divided from the MPEG2-PES packet which had been read in step S1.”

Hiroshima, 11:13–27.

FIG. 15



Hiroshima, Fig. 15.


Exhibit C-8

		Austin discloses and/or renders obvious this limitation. <i>See, e.g.:</i> “This system uses up to 5 pre-defined data packet sizes and up to 16 user defined data message types within each packet. Each packet has Forward Error Connection (FEC) and bit interleaving applied to it on transmission. Error checking is performed using a standard CRC. The system uses standard run in and byte synchronisation techniques. The byte synchronisation word used is 32 _bits long and there are 5_different words used, one to identify each of the different packet types.” Austin, 4:37-45. To the extent that Plaintiff contends that the Additional Prior Art References do not explicitly or inherently disclose this element, or render obvious this element, this element is obvious in combination with one or more references and/or with the knowledge of a person of ordinary skill in the art, as set forth in Defendant’s Invalidity Contentions.
11(d)	the synchronization symbol permitting properly alignment upon receiving with a synchronization primitive to be correctly aligned on a symbol boundary.	Cunningham discloses and/or renders obvious this limitation. <i>See, e.g.:</i> “There are many potential ways to achieve code block synchronization (i.e. correct alignment in the receiver of the label and payload fields). For example, one technique would be to transmit a ‘comma sequence’ during start-up. As an example of this technique consider a code with n=8 using the simple control payload scheme. A training control state of 1111 1111 would be transmitted as: [10 1111 1111 0000 0000][10 1111 1111 0000 0000][10 1111 1111 0000 0000]” Cunningham, 7:5–13. Haartsen discloses and/or renders obvious this limitation. <i>See, e.g.:</i> “The present invention fulfills the above-described and other needs by providing a multi-part digital preamble for use in transmitting digital data packets. Advantageously, a preamble according to the invention substantially reduces the overall length of DC-free sequences that are transmitted with each data packet in a digital FM system. According to exemplary embodiments, a multi-part preamble includes a short, substantially DC-free leading part which enables coarse estimation of a

Exhibit C-8

	prevailing DC level at an FM receiver. The exemplary preamble also includes a synchronization part which is not necessarily substantially DC-free and which carries timing and/or other useful information (e.g. , channel identification, destination address, etc.) that can vary from data packet to data packet (e.g. , depending on sender, recipient, time of day, etc.).” Haartsen, 5:2–15. “To alleviate this problem, the present invention discloses that a digital preamble can be strategically divided into three or more parts so that synchronization and other information can be included within the preamble while still enabling accurate DC estimation. Specifically, a short leading DC-free part of the preamble (e.g., alternating logical 1 's and 0's, or any other sequence including an equal number of logical 1's and logical 0's) can be used to provide initial, coarse DC offset estimation. Thereafter, a synchronization part of the preamble, which is not necessarily DC free, but which is protected by code to prevent detection errors, can be used not only to synchronize detection of a data packet, but also to carry timing and/or other useful information (e.g., channel identification, destination address, time of day, user identification, etc.).” Haartsen, 12:13–13:2. “An exemplary multi-part preamble according to the invention is depicted in Figure 5. In the figure, a three-part preamble 500 is shown to include a DC-free leading sequence 510, followed by a multiple-bit synchronization word 520 which is not necessarily DC free and which can be different for each data packet. The synchronization word 520, which can include timing and/or other information and which is coded to protect against detection errors, is in turn followed by a trailing DC-free sequence 530. Although the leading and trailing DC-free sequences 510, 530 are shown in the figure as 4-bit sequences of alternating logical 1 's and 0's, those of ordinary skill will appreciate that each DC-free sequence 510, 530 can, in practice, be shorter or longer in length and can include any DC-free sequence having an equal number of logical 1's and logical 0's.” Haartsen, 13:13–14:2.
--	---

Exhibit C-8



Haartsen, Fig. 5.

“When the exemplary preamble 500 is received by the system 300 of Figure 3, the leading DC-free sequence 510 induces a coarse estimate of the baseband DC offset on the storage capacitor 340. Although the estimate is coarse, it is sufficient to permit successful detection of the coded synchronization word 520 (e.g., by the correlation processor 310). Methods for coding of a synchronization word to protect against erroneous detection in the presence of some DC-offset are known to those of skill in the art. Once the correlation processor 310 has detected the synchronization word 520, timing between received bits is well known. Consequently, the correlation processor 310 can open the switch 330 precisely at the end of the trailing DC-free sequence 530 to thereby store an accurate estimate of the DC offset for reception and detection of the incoming digital data packet.”

Haartsen, 14:3–14.

Shin discloses and/or renders obvious this limitation. *See, e.g.:*

“In accordance with one aspect of the invention, the synchronization characters associated with both the low-transition and high-transition sets of code characters are selected so as to facilitate rapid synchronization during data recovery. When the low-transition set of code characters is being employed, a special group of out-of-band characters is used during synchronization. Each special synchronization character includes a predefined number larger than four (e.g., 7) of logical transitions between binary character values, and a predefined number (e.g., two) of "non-transitions" between character values. As will be discussed below, the special synchronization


Exhibit C-8

characters are selected such that random logic may be used to distinguish each special synchronization character from the in-band characters of the low-transition set.”

Shin, 10:17–34.

“Referring to FIG. 6, the deserializer 34 receives the sequence of 10-bit characters T[9:0] produced by the encoder and generates 10-bit parallel received data upon bit lines RX9, RX8, . . . , RX0 (i.e., RX[9:0]). This 10-bit parallel received data is provided over bit lines RX[9:0] to the decoder 44, as well as to a decoder synchronization module 114. As is described below in section IV, the synchronization module 114 is operative to ascertain boundaries within the 10-bit parallel received data corresponding to the frame boundaries of the transmitted data (i.e., to T[9:0]). Specifically, synchronization module 114 determines upon which of the bit lines RX[9:0] the deserializer 34 is providing the received bits corresponding to the first bit T[0] of each transmitted byte T[9:0]. Upon making this determination, the synchronization module 114 provides a frame boundary pointer 118 to decoder 44 identifying the one of the bit lines RX[9:0] corresponding to the first bit T[0] of each transmitted 10-bit character T[9:0]. Upon receiving this synchronization information, the decoder 44 is disposed to decode the received data RX[9:0] in the following manner.”

Shin, 13:37–57.




Shin, Fig. 6.

To the extent that Plaintiff contends that the Additional Prior Art References do not explicitly or inherently disclose this element, or render obvious this element, this element is obvious in

Exhibit C-8

		combination with one or more references and/or with the knowledge of a person of ordinary skill in the art, as set forth in Defendant's Invalidity Contentions.
Claim 19		
19[p]	A communications device for transmitting packet types of packets, comprising:	To the extent the preamble is limiting, Chen discloses this element and/or renders it obvious either alone or in combination with other references and/or the knowledge of one of ordinary skill in the art. <i>See, e.g.:</i> “FIG. 1 is a block diagram of a Fiber-to-the-Curb Digital Video Distribution Network in accordance with the present invention. The distribution network includes Video Information Providers (VIPs) shown generally at 100, a gateway, shown generally at 120, an access network, shown generally at 140, and the home of a Video Information User (VIU), shown generally at 160. The gateway 120 may comprise an L1 Gateway manufactured by Teleco. The VIPs may include a video server 102, an interactive data server 104, and an internet gateway 106, all of which communicate with an asynchronous transfer mode (ATM) network 122 via, for example, an OC-3 system, where optical fiber is used to transmit data at 155.52 Mbit/sec. The ATM network 122 may also receive data via a satellite receiving antenna 110, broadcast encoder and server 112, and ATM multiplexer 114. The ATM network 122 communicates with a session manager 123.” Chen, 4:11–28.

Exhibit C-8



Chen, Fig. 1.

“FIG. 4 is a block diagram of an Insertion Processing Unit (IPU) of a Digital Ad Insertion Module in accordance with the present invention. The IPU, shown generally at 400, receives the main stream at a packet framer 405, while the insertion stream is received at another packet framer 410. After processing by the packet framer 405, the main stream is provided to a Main Stream Parser (MSP) 415.”

Chen, 5:62–6:1.

Exhibit C-8




FIG. 4


Chen, Fig. 4.

“The main stream parser 415 provides the main stream data to a main stream buffer 480, which communicates with a syntax processor 470. The main stream buffer 480 is part of the output buffer 475. The insertion stream parser 420 provides the insertion stream data to a Program Specific Information (PSI)/Program Identifier (PID) Replacer 435, which replaces the PSI tables and PIDs of the insertion stream with those of the main stream. The PSI/PID replacer 435 retrieves the relevant information from the main stream using a communication path which is not shown. The PSI and PID data of the main stream provide the information to connect the various audio, video, and data packets to a particular programming service. Specifically, this is accomplished using a Program Association Table (PAT) and a Program Map Table (PMT). The insertion stream is then provided to an Insertion Stream Buffer (ISB) 490, which is part of the output buffer 475. The ISB 490 communicates with a syntax processor 470. Furthermore, the output buffer 475 receives commands from a buffer controller 465 to provide an output stream with the insertion stream seamlessly spliced into the main stream.”

Exhibit C-8

	Chen, 6:33–54. “FIG. 5 illustrates the processing flow of the Insertion Processing Unit of FIG. 4 in accordance with the present invention. Note that the steps of FIG. 5 are shown as occurring in a sequential manner for clarity, but many of the steps may occur concurrently and/or sequentially. At block 502, the IPU outputs the main transport stream from the main stream buffer (MSB) 480. The IPU is essentially in a bypass mode at this time since the insertion stream has not yet been inserted. At block 506, the main stream parser parses the main stream to locate the splicing signal T-- in. Note that T-- in need not be carried in the main stream, but may be provided locally, for example, at a cable system headend. When T-- in is detected, the IPU continues to output the main transport stream from the main stream buffer. At block 508, the main stream parser parses the main stream to locate the closest packet with an anchor picture (e.g., I or P picture) start code and passes that packet to the syntax processor 470. This packet will be the last packet of the main stream. To ensure that the resulting data stream is MPEG compliant, no data from the next frame or sequence of frames may be included in this last main stream packet. At block 512, the syntax processor adds a sequence end code field to the last packet of the main stream. The sequence end code is a twenty-four bit code which indicates the end of a sequence of pictures. At block 514, the PSI/PID replacer 435 starts to replace the PSI/PID of the insertion stream.” Chen, 12:28–53.
--	---

Exhibit C-8




Chen, Fig. 5.

“In the MPEG-2 system (and the similar DigiCipher® II system proprietary to General Instrument Corporation, the assignee hereof) a transport stream, or transport multiplex is made up of a

Exhibit C-8

	contiguous set of fixed length packets. Each packet is 188 total bytes in length, with the first four of those bytes being defined as the packet header. The payload portion of each packet is thus normally 184 bytes. However, a variable length adaptation field may be provided to extend the header, when required. When an adaptation field is present, the payload portion of the packet will be correspondingly shorter. Various timing and identification information is provided in different portions of the transport stream. These include a packet identifier (PID) found in the transport header of each transport packet to provide a reference number for identifying the transport packets carrying a specific service component. This number is included in a service definition or "service map" used by the receiver to identify those transport packets required to reconstruct a television program signal. The PID may also be referenced for various grooming and remultiplexing functions. In the case of video, audio or isochronous control data, the stream of packets labeled with a single PID represents a single video, audio or isochronous data service elementary stream, respectively. Each type of packet will have a different PID identifying the packet type." Chen, 14:45–15:3. To the extent the preamble is limiting, Cunningham discloses this element and/or renders it obvious either alone or in combination with other references and/or the knowledge of one of ordinary skill in the art. <i>See, e.g.:</i> "Referring to FIG. 1, a data frame intended to be communicated over an optical fibre channel 20 typically comprises a binary digital message (payload) to be transferred between stations on the network; this message may comprise user data, or control information for managing operation of the network and its components. A circuit 12 generates an associated cyclic redundancy check (CRC) block containing check data derived from this message in known manner in accordance with a predetermined CRC algorithm, and appends the CRC block to the message." Cunningham, 2:65–3:7.
--	--

Exhibit C-8



Cunningham, Fig. 1.


“Control information is conveyed by using selected n-bit data block values, together with a distinctive value of 10 for the pair of label bits a and b to indicate that the associated data block represents control information rather than actual data. The data block values used for this purpose are selected to be values which inherently have a disparity of zero, such as (for 8-bit blocks)

Exhibit C-8

	00001111 or 10101010. It will be noted that the label bit pair value 10 likewise has zero disparity, so the complete encoded block also has zero disparity.” Cunningham, 4:23–31. “The encoded blocks are transmitted over the optical fibre channel 20 (leftmost symbol first), and upon receipt at their destination they are detected by a receiver 22 where compensation is made in known manner for the effects of timing drift and other signal degradation to produce a restored digital data stream.” Cunningham, 4:45–50. To the extent the preamble is limiting, Hiroshima discloses this element and/or renders it obvious either alone or in combination with other references and/or the knowledge of one of ordinary skill in the art. <i>See, e.g.:</i> “FIG. 4 is a block diagram of the first embodiment of the invention in which the moving picture stream conversion in mode 1 is executed. A moving picture stream converting apparatus of FIG. 4 comprises a demultiplexer 30, a packetizer 32, and a multiplexer 34. The demultiplexer 30 receives an MPEG1 system stream as a conversion source and outputs elementary streams which were separated into video and audio streams. The packetizer 32 forms packets constructing a packetized elementary stream of MPEG2 with respect to each of the video and audio elementary streams which are supplied from the demultiplexer 30. The packet is usually called a PES packet. The multiplexer 34 divides each of the video and audio PES packets into a TS packet having a fixed length of 188 bytes and subsequently multiplexes, and outputs the resultant data as an MPEG2-TS.” Hiroshima, 7:59–8:7.
--	---

Exhibit C-8

FIG. 4




Hiroshima, Fig. 4.

“FIG. 11 shows the details of the MPEG2-TS. Each packet of the MPEG2-TS has a prefix length of 188 bytes. The header 194 consists of four bytes and the payload 196 consists of 184 bytes. The header 194 is constructed by sync bytes 198, an error indicator 200, a unit start indicator 202, a transport priority 204, a packet ID 206, a scramble control 208, an adaptation field control 210, a continuity counter 212, and an adaptation field 214. The adaptation field 214 is constructed by an adaptation field length 216, a discontinuity indicator 218, a random access indicator 220, an elementary stream priority indicator 222, a 5-flag 224, an optional field 226, and stuffing bytes 228. Among them, the optional field 226 is constructed by a program clock reference 230, an original program clock reference field 232, a splice count down 234, a private data length 236, private data 238, an adaptation field extension 240, 3-flag 241, and an optional field 242.”

Hiroshima, 10:51–67.

Exhibit C-8




Hiroshima, Fig. 11.

“FIG. 15 shows a process for forming the MPEG2-TS from the PES packet 82 by a multiplexing process in the multiplexer 34 in FIG. 5. First in step S1, the MPEG2-PES file is read. The TS header 194 in FIG. 11 is formed in step S2. The TS header 194 is formed in a manner such that the parameters such as error indicator 200 indicative of the presence or absence of the bit error in the packet, unit start indicator 202 indicating whether a new PES packet appears in the TS payload or not, packet ID 206 indicating whether the packet in the payload is video or audio, adaptation field 214, and adaptation field control 210 indicative of the presence or absence of the payload 196 are set in accordance with the state of the payload storage data of the 184-byte unit divided from the MPEG2-PES packet which had been read in step S1.”

Hiroshima, 11:13–27.

Exhibit C-8

FIG. 15



Hiroshima, Fig. 15.

To the extent the preamble is limiting, Haartsen discloses this element and/or renders it obvious either alone or in combination with other references and/or the knowledge of one of ordinary skill in the art. *See, e.g.:*

“The present invention fulfills the above-described and other needs by providing a multi-part digital preamble for use in transmitting digital data packets. Advantageously, a preamble according to the invention substantially reduces the overall length of DC-free sequences that are transmitted with each data packet in a digital FM system. According to exemplary embodiments, a multi-part


Exhibit C-8

preamble includes a short, substantially DC-free leading part which enables coarse estimation of a prevailing DC level at an FM receiver. The exemplary preamble also includes a synchronization part which is not necessarily substantially DC-free and which carries timing and/or other useful information (e.g. , channel identification, destination address, etc.) that can vary from data packet to data packet (e.g. , depending on sender, recipient, time of day, etc.).”

Haartsen, 5:2–15.

“An exemplary DC estimation processor is depicted in Figure 3, wherein an exemplary FM demodulation processor 300 is shown to include a correlation processor 310, a resistor 320, a switch 330 and a capacitor 340, as well as the FM detector 210, the low-pass filter 130 and the slicer 230 of Figure 2.”

Haartsen, 9:11–14.



Haartsen, Fig. 3.

“To alleviate this problem, the present invention discloses that a digital preamble can be strategically divided into three or more parts so that synchronization and other information can be included within the preamble while still enabling accurate DC estimation. Specifically, a short leading DC-free part of the preamble (e.g., alternating logical 1's and 0's, or any other sequence including an equal number of logical 1's and logical 0's) can be used to provide initial, coarse DC offset estimation. Thereafter, a synchronization part of the preamble, which is not necessarily DC


Exhibit C-8

free, but which is protected by code to prevent detection errors, can be used not only to synchronize detection of a data packet, but also to carry timing and/or other useful information (e.g., channel identification, destination address, time of day, user identification, etc.).”

Haartsen, 12:13–13:2.

“An exemplary multi-part preamble according to the invention is depicted in Figure 5. In the figure, a three-part preamble 500 is shown to include a DC-free leading sequence 510, followed by a multiple-bit synchronization word 520 which is not necessarily DC free and which can be different for each data packet. The synchronization word 520, which can include timing and/or other information and which is coded to protect against detection errors, is in turn followed by a trailing DC-free sequence 530. Although the leading and trailing DC-free sequences 510, 530 are shown in the figure as 4-bit sequences of alternating logical 1's and 0's, those of ordinary skill will appreciate that each DC-free sequence 510, 530 can, in practice, be shorter or longer in length and can include any DC-free sequence having an equal number of logical 1's and logical 0's.”

Haartsen, 13:13–14:2.



Haartsen, Fig. 5.

“When the exemplary preamble 500 is received by the system 300 of Figure 3, the leading DC-free sequence 510 induces a coarse estimate of the baseband DC offset on the storage capacitor 340. Although the estimate is coarse, it is sufficient to permit successful detection of the coded synchronization word 520 (e.g., by the correlation processor 310). Methods for coding of a

Exhibit C-8

	synchronization word to protect against erroneous detection in the presence of some DC-offset are known to those of skill in the art. Once the correlation processor 310 has detected the synchronization word 520, timing between received bits is well known. Consequently, the correlation processor 310 can open the switch 330 precisely at the end of the trailing DC-free sequence 530 to thereby store an accurate estimate of the DC offset for reception and detection of the incoming digital data packet.” Haartsen, 14:3–14. To the extent the preamble is limiting, Shin discloses this element and/or renders it obvious either alone or in combination with other references and/or the knowledge of one of ordinary skill in the art. <i>See, e.g.:</i> “Referring to FIG. 6, the deserializer 34 receives the sequence of 10-bit characters T[9:0] produced by the encoder and generates 10-bit parallel received data upon bit lines RX9, RX8, . . . , RX0 (i.e., RX[9:0]). This 10-bit parallel received data is provided over bit lines RX[9:0] to the decoder 44, as well as to a decoder synchronization module 114. As is described below in section IV, the synchronization module 114 is operative to ascertain boundaries within the 10-bit parallel received data corresponding to the frame boundaries of the transmitted data (i.e., to T[9:0]). Specifically, synchronization module 114 determines upon which of the bit lines RX[9:0] the deserializer 34 is providing the received bits corresponding to the first bit T[0] of each transmitted byte T[9:0]. Upon making this determination, the synchronization module 114 provides a frame boundary pointer 118 to decoder 44 identifying the one of the bit lines RX[9:0] corresponding to the first bit T[0] of each transmitted 10-bit character T[9:0]. Upon receiving this synchronization information, the decoder 44 is disposed to decode the received data RX[9:0] in the following manner.” Shin, 13:37–57.
--	--

Exhibit C-8




FIG. 6

Shin, Fig. 6.

To the extent the preamble is limiting, Austin discloses this element and/or renders it obvious either alone or in combination with other references and/or the knowledge of one of ordinary skill in the art. *See, e.g.:*

“FIG. 2 shows the structure of a transmitter suitable for transformer/rectifier sites. The transmitter 30 comprises waveform generators 31-34 which transmitter receives the data at input 35 which selects 10 Hz tone as an indication of data ‘1’ and a 15 Hz tone as a data ‘0’ via selector 36. Using Fast Frequency Shift Keying (FFSK) at 10 baud, the modulated data signal is applied by modulating part of the current applied to the pipe (current switch 37) which provides cathodic protection (CP) as a result of the pipeline resistance. The signals are in digital form with multistage representations of the two modulating tones. Each state of these tones represents the switching levels for the transformer/rectifier pipeline devices. The bandwidth of the pipe is such that the switching harmonics are considerably reduced through their passage in the pipe so are of no great concern at the receiver.”

Austin, 2:32-47.

Exhibit C-8

		FIG. 2. Austin, Fig. 2. “This system uses up to 5 pre-defined data packet sizes and up to 16 user defined data message types within each packet. Each packet has Forward Error Connection (FEC) and bit interleaving applied to it on transmission. Error checking is performed using a standard CRC. The system uses standard run in and byte synchronisation techniques. The byte synchronisation word used is 32 _bits long and there are 5_different words used, one to identify each of the different packet types.” Austin, 4:37-45. To the extent that Plaintiff contends that the Additional Prior Art References do not explicitly or inherently disclose this element, or render obvious this element, this element is obvious in combination with one or more references and/or with the knowledge of a person of ordinary skill in the art, as set forth in Defendant’s Invalidity Contentions.
19(a)	an identification component that identifies a	Chen discloses and/or renders obvious this limitation. <i>See, e.g.:</i>

Exhibit C-8

packet type of a packet of symbols; and

“FIG. 4 is a block diagram of an Insertion Processing Unit (IPU) of a Digital Ad Insertion Module in accordance with the present invention. The IPU, shown generally at 400, receives the main stream at a packet framer 405, while the insertion stream is received at another packet framer 410. After processing by the packet framer 405, the main stream is provided to a Main Stream Parser (MSP) 415.”

Chen, 5:62–6:1.




FIG.4

Chen, Fig. 4.

“The main stream parser 415 provides the main stream data to a main stream buffer 480, which communicates with a syntax processor 470. The main stream buffer 480 is part of the output buffer 475. The insertion stream parser 420 provides the insertion stream data to a Program Specific Information (PSI)/Program Identifier (PID) Replacer 435, which replaces the PSI tables and PIDs of the insertion stream with those of the main stream. The PSI/PID replacer 435 retrieves the relevant information from the main stream using a communication path which is not shown. The

Exhibit C-8

	PSI and PID data of the main stream provide the information to connect the various audio, video, and data packets to a particular programming service. Specifically, this is accomplished using a Program Association Table (PAT) and a Program Map Table (PMT). The insertion stream is then provided to an Insertion Stream Buffer (ISB) 490, which is part of the output buffer 475. The ISB 490 communicates with a syntax processor 470. Furthermore, the output buffer 475 receives commands from a buffer controller 465 to provide an output stream with the insertion stream seamlessly spliced into the main stream.” Chen, 6:33–54. “FIG. 5 illustrates the processing flow of the Insertion Processing Unit of FIG. 4 in accordance with the present invention. Note that the steps of FIG. 5 are shown as occurring in a sequential manner for clarity, but many of the steps may occur concurrently and/or sequentially. At block 502, the IPU outputs the main transport stream from the main stream buffer (MSB) 480. The IPU is essentially in a bypass mode at this time since the insertion stream has not yet been inserted. At block 506, the main stream parser parses the main stream to locate the splicing signal T-- in. Note that T-- in need not be carried in the main stream, but may be provided locally, for example, at a cable system headend. When T-- in is detected, the IPU continues to output the main transport stream from the main stream buffer. At block 508, the main stream parser parses the main stream to locate the closest packet with an anchor picture (e.g., I or P picture) start code and passes that packet to the syntax processor 470. This packet will be the last packet of the main stream. To ensure that the resulting data stream is MPEG compliant, no data from the next frame or sequence of frames may be included in this last main stream packet. At block 512, the syntax processor adds a sequence end code field to the last packet of the main stream. The sequence end code is a twenty-four bit code which indicates the end of a sequence of pictures. At block 514, the PSI/PID replacer 435 starts to replace the PSI/PID of the insertion stream.” Chen, 12:28–53.
--	--

Exhibit C-8




FIG.5


Chen, Fig. 5.

“In the MPEG-2 system (and the similar DigiCipher® II system proprietary to General Instrument Corporation, the assignee hereof) a transport stream, or transport multiplex is made up of a

Exhibit C-8

	contiguous set of fixed length packets. Each packet is 188 total bytes in length, with the first four of those bytes being defined as the packet header. The payload portion of each packet is thus normally 184 bytes. However, a variable length adaptation field may be provided to extend the header, when required. When an adaptation field is present, the payload portion of the packet will be correspondingly shorter. Various timing and identification information is provided in different portions of the transport stream. These include a packet identifier (PID) found in the transport header of each transport packet to provide a reference number for identifying the transport packets carrying a specific service component. This number is included in a service definition or "service map" used by the receiver to identify those transport packets required to reconstruct a television program signal. The PID may also be referenced for various grooming and remultiplexing functions. In the case of video, audio or isochronous control data, the stream of packets labeled with a single PID represents a single video, audio or isochronous data service elementary stream, respectively. Each type of packet will have a different PID identifying the packet type." Chen, 14:45–15:3. Cunningham discloses and/or renders obvious this limitation. <i>See, e.g.:</i> "Referring to FIG. 1, a data frame intended to be communicated over an optical fibre channel 20 typically comprises a binary digital message (payload) to be transferred between stations on the network; this message may comprise user data, or control information for managing operation of the network and its components. A circuit 12 generates an associated cyclic redundancy check (CRC) block containing check data derived from this message in known manner in accordance with a predetermined CRC algorithm, and appends the CRC block to the message." Cunningham, 2:65–3:7.
--	--

Exhibit C-8



Cunningham, Fig. 1.


“Control information is conveyed by using selected n-bit data block values, together with a distinctive value of 10 for the pair of label bits a and b to indicate that the associated data block represents control information rather than actual data. The data block values used for this purpose are selected to be values which inherently have a disparity of zero, such as (for 8-bit blocks)

Exhibit C-8

	00001111 or 10101010. It will be noted that the label bit pair value 10 likewise has zero disparity, so the complete encoded block also has zero disparity.” Cunningham, 4:23–31. “The encoded blocks are transmitted over the optical fibre channel 20 (leftmost symbol first), and upon receipt at their destination they are detected by a receiver 22 where compensation is made in known manner for the effects of timing drift and other signal degradation to produce a restored digital data stream.” Cunningham, 4:45–50. Hiroshima discloses and/or renders obvious this limitation. <i>See, e.g.</i>: “FIG. 4 is a block diagram of the first embodiment of the invention in which the moving picture stream conversion in mode 1 is executed. A moving picture stream converting apparatus of FIG. 4 comprises a demultiplexer 30, a packetizer 32, and a multiplexer 34. The demultiplexer 30 receives an MPEG1 system stream as a conversion source and outputs elementary streams which were separated into video and audio streams. The packetizer 32 forms packets constructing a packetized elementary stream of MPEG2 with respect to each of the video and audio elementary streams which are supplied from the demultiplexer 30. The packet is usually called a PES packet. The multiplexer 34 divides each of the video and audio PES packets into a TS packet having a fixed length of 188 bytes and subsequently multiplexes, and outputs the resultant data as an MPEG2-TS.” Hiroshima, 7:59–8:7.
--	---

Exhibit C-8

FIG. 4




Hiroshima, Fig. 4.

“FIG. 11 shows the details of the MPEG2-TS. Each packet of the MPEG2-TS has a prefix length of 188 bytes. The header 194 consists of four bytes and the payload 196 consists of 184 bytes. The header 194 is constructed by sync bytes 198, an error indicator 200, a unit start indicator 202, a transport priority 204, a packet ID 206, a scramble control 208, an adaptation field control 210, a continuity counter 212, and an adaptation field 214. The adaptation field 214 is constructed by an adaptation field length 216, a discontinuity indicator 218, a random access indicator 220, an elementary stream priority indicator 222, a 5-flag 224, an optional field 226, and stuffing bytes 228. Among them, the optional field 226 is constructed by a program clock reference 230, an original program clock reference field 232, a splice count down 234, a private data length 236, private data 238, an adaptation field extension 240, 3-flag 241, and an optional field 242.”

Hiroshima, 10:51–67.

Exhibit C-8




Hiroshima, Fig. 11.

“FIG. 15 shows a process for forming the MPEG2-TS from the PES packet 82 by a multiplexing process in the multiplexer 34 in FIG. 5. First in step S1, the MPEG2-PES file is read. The TS header 194 in FIG. 11 is formed in step S2. The TS header 194 is formed in a manner such that the parameters such as error indicator 200 indicative of the presence or absence of the bit error in the packet, unit start indicator 202 indicating whether a new PES packet appears in the TS payload or not, packet ID 206 indicating whether the packet in the payload is video or audio, adaptation field 214, and adaptation field control 210 indicative of the presence or absence of the payload 196 are set in accordance with the state of the payload storage data of the 184-byte unit divided from the MPEG2-PES packet which had been read in step S1.”

Hiroshima, 11:13–27.

Exhibit C-8

FIG. 15



Hiroshima, Fig. 15.

Austin discloses and/or renders obvious this limitation. *See, e.g.:*

“This system uses up to 5 pre-defined data packet sizes and up to 16 user defined data message types within each packet. Each packet has Forward Error Connection (FEC) and bit interleaving applied to it on transmission. Error checking is performed using a standard CRC. The system uses standard run in and byte synchronisation techniques. The byte synchronisation word used is 32_bits long and there are 5_different words used, one to identify each of the different packet types.”

Exhibit C-8

		Austin, 4:37-45. To the extent that Plaintiff contends that the Additional Prior Art References do not explicitly or inherently disclose this element, or render obvious this element, this element is obvious in combination with one or more references and/or with the knowledge of a person of ordinary skill in the art, as set forth in Defendant's Invalidity Contentions.
19(b)	a transmission component that transmits a synchronization symbol that corresponds to the identified packet type, the transmitted synchronization symbol providing synchronization information and each packet type having a different synchronization symbol and permitting an external receiving node to properly align with a synchronization primitive to be correctly aligned on a symbol boundary, and that transmits the symbols of the packet.	Chen discloses and/or renders obvious this limitation. <i>See, e.g.:</i> "FIG. 4 is a block diagram of an Insertion Processing Unit (IPU) of a Digital Ad Insertion Module in accordance with the present invention. The IPU, shown generally at 400, receives the main stream at a packet framer 405, while the insertion stream is received at another packet framer 410. After processing by the packet framer 405, the main stream is provided to a Main Stream Parser (MSP) 415." Chen, 5:62-6:1.

Exhibit C-8




FIG.4

Chen, Fig. 4.

“The main stream parser 415 provides the main stream data to a main stream buffer 480, which communicates with a syntax processor 470. The main stream buffer 480 is part of the output buffer 475. The insertion stream parser 420 provides the insertion stream data to a Program Specific Information (PSI)/Program Identifier (PID) Replacer 435, which replaces the PSI tables and PIDs of the insertion stream with those of the main stream. The PSI/PID replacer 435 retrieves the relevant information from the main stream using a communication path which is not shown. The PSI and PID data of the main stream provide the information to connect the various audio, video, and data packets to a particular programming service. Specifically, this is accomplished using a Program Association Table (PAT) and a Program Map Table (PMT). The insertion stream is then provided to an Insertion Stream Buffer (ISB) 490, which is part of the output buffer 475. The ISB 490 communicates with a syntax processor 470. Furthermore, the output buffer 475 receives commands from a buffer controller 465 to provide an output stream with the insertion stream seamlessly spliced into the main stream.”

Exhibit C-8

	Chen, 6:33–54. “FIG. 5 illustrates the processing flow of the Insertion Processing Unit of FIG. 4 in accordance with the present invention. Note that the steps of FIG. 5 are shown as occurring in a sequential manner for clarity, but many of the steps may occur concurrently and/or sequentially. At block 502, the IPU outputs the main transport stream from the main stream buffer (MSB) 480. The IPU is essentially in a bypass mode at this time since the insertion stream has not yet been inserted. At block 506, the main stream parser parses the main stream to locate the splicing signal T-- in. Note that T-- in need not be carried in the main stream, but may be provided locally, for example, at a cable system headend. When T-- in is detected, the IPU continues to output the main transport stream from the main stream buffer. At block 508, the main stream parser parses the main stream to locate the closest packet with an anchor picture (e.g., I or P picture) start code and passes that packet to the syntax processor 470. This packet will be the last packet of the main stream. To ensure that the resulting data stream is MPEG compliant, no data from the next frame or sequence of frames may be included in this last main stream packet. At block 512, the syntax processor adds a sequence end code field to the last packet of the main stream. The sequence end code is a twenty-four bit code which indicates the end of a sequence of pictures. At block 514, the PSI/PID replacer 435 starts to replace the PSI/PID of the insertion stream.” Chen, 12:28–53.
--	---

Exhibit C-8




FIG.5

Chen, Fig. 5.

“In the MPEG-2 system (and the similar DigiCipher® II system proprietary to General Instrument Corporation, the assignee hereof) a transport stream, or transport multiplex is made up of a


Exhibit C-8

	contiguous set of fixed length packets. Each packet is 188 total bytes in length, with the first four of those bytes being defined as the packet header. The payload portion of each packet is thus normally 184 bytes. However, a variable length adaptation field may be provided to extend the header, when required. When an adaptation field is present, the payload portion of the packet will be correspondingly shorter. Various timing and identification information is provided in different portions of the transport stream. These include a packet identifier (PID) found in the transport header of each transport packet to provide a reference number for identifying the transport packets carrying a specific service component. This number is included in a service definition or "service map" used by the receiver to identify those transport packets required to reconstruct a television program signal. The PID may also be referenced for various grooming and remultiplexing functions. In the case of video, audio or isochronous control data, the stream of packets labeled with a single PID represents a single video, audio or isochronous data service elementary stream, respectively. Each type of packet will have a different PID identifying the packet type. Timing information carried by the transport stream includes a program clock reference (PCR) which effectively represents a sample of the system time clock (STC) time base that underlies the service composed of the PIDs referenced in the service map. The PID carrying the packet with the PCR is also referenced in the service map. The video, audio and isochronous data components of a service are locked through a defined relationship to the system time clock. The PCR serves to define the transport rate, in the sense that between any two successive PCRs in one PID, the transport rate is constant and nominally equal to the system time clock rate times the ratio of the total number of transport bytes between the PCRs divided by the difference in the PCRs in units of system time clock ticks." Chen, 14:45–15:17. "As indicated in FIG. 6b, the transport header of each transport packet includes the PID which identifies the particular service component carried by the transport packet. The PCR will be carried by an adaptation field to provide timing information for a desired service. At a receiver, the PCR for the desired service is detected from the adaptation field. The PIDs of the transport packets will
--	---

Exhibit C-8

then be monitored in accordance with the timing established by the PCR to recover those transport packets carrying a particular component of the service to be processed.”

Chen, 16:41–50.




Chen, Fig. 6b.

Cunningham discloses and/or renders obvious this limitation. *See, e.g.:*

“Referring to FIG. 1, a data frame intended to be communicated over an optical fibre channel 20 typically comprises a binary digital message (payload) to be transferred between stations on the network; this message may comprise user data, or control information for managing operation of the network and its components. A circuit 12 generates an associated cyclic redundancy check (CRC) block containing check data derived from this message in known manner in accordance with a predetermined CRC algorithm, and appends the CRC block to the message.”

Cunningham, 2:65–3:7.

Exhibit C-8



Cunningham, Fig. 1.

“Control information is conveyed by using selected n-bit data block values, together with a distinctive value of 10 for the pair of label bits a and b to indicate that the associated data block represents control information rather than actual data. The data block values used for this purpose are selected to be values which inherently have a disparity of zero, such as (for 8-bit blocks)

Exhibit C-8


	00001111 or 10101010. It will be noted that the label bit pair value 10 likewise has zero disparity, so the complete encoded block also has zero disparity.” Cunningham, 4:23–31. “The encoded blocks are transmitted over the optical fibre channel 20 (leftmost symbol first), and upon receipt at their destination they are detected by a receiver 22 where compensation is made in known manner for the effects of timing drift and other signal degradation to produce a restored digital data stream.” Cunningham, 4:45–50. “There are many potential ways to achieve code block synchronization (i.e. correct alignment in the receiver of the label and payload fields). For example, one technique would be to transmit a ‘comma sequence’ during start-up. As an example of this technique consider a code with n=8 using the simple control payload scheme. A training control state of 1111 1111 would be transmitted as: [10 1111 1111 0000 0000][10 1111 1111 0000 0000][10 1111 1111 0000 0000] The receiver searches for the bit pair 10. This is found in only two locations: the correct location (the label field), and at the end of the run of eight 1's. However, if the receiver (incorrectly) identifies the label field to be the 10 bit pair at the end of the run of 1's, the following 16 bits are not a valid control payload (first 8 bits are not the complement of the second 8 bits). The receiver can therefore detect the mistake. It is also noted that this sequence contains a regular 010 pattern, which aids in phase-locked loop (PLL) clock acquisition.” Cunningham, 7:5–23. Hiroshima discloses and/or renders obvious this limitation. <i>See, e.g.</i> “FIG. 4 is a block diagram of the first embodiment of the invention in which the moving picture stream conversion in mode 1 is executed. A moving picture stream converting apparatus of FIG. 4 comprises a demultiplexer 30, a packetizer 32, and a multiplexer 34. The demultiplexer 30 receives an MPEG1 system stream as a conversion source and outputs elementary streams which were
--	---

Exhibit C-8

separated into video and audio streams. The packetizer 32 forms packets constructing a packetized elementary stream of MPEG2 with respect to each of the video and audio elementary streams which are supplied from the demultiplexer 30. The packet is usually called a PES packet. The multiplexer 34 divides each of the video and audio PES packets into a TS packet having a fixed length of 188 bytes and subsequently multiplexes, and outputs the resultant data as an MPEG2-TS.”

Hiroshima, 7:59–8:7.

FIG. 4




Hiroshima, Fig. 4.

“FIG. 11 shows the details of the MPEG2-TS. Each packet of the MPEG2-TS has a prefix length of 188 bytes. The header 194 consists of four bytes and the payload 196 consists of 184 bytes. The header 194 is constructed by sync bytes 198, an error indicator 200, a unit start indicator 202, a transport priority 204, a packet ID 206, a scramble control 208, an adaptation field control 210, a continuity counter 212, and an adaptation field 214. The adaptation field 214 is constructed by an adaptation field length 216, a discontinuity indicator 218, a random access indicator 220, an elementary stream priority indicator 222, a 5-flag 224, an optional field 226, and stuffing bytes 228. Among them, the optional field 226 is constructed by a program clock reference 230, an original program clock reference field 232, a splice count down 234, a private data length 236, private data 238, an adaptation field extension 240, 3-flag 241, and an optional field 242.”

Hiroshima, 10:51–67.

Exhibit C-8




Hiroshima, Fig. 11.

“FIG. 15 shows a process for forming the MPEG2-TS from the PES packet 82 by a multiplexing process in the multiplexer 34 in FIG. 5. First in step S1, the MPEG2-PES file is read. The TS header 194 in FIG. 11 is formed in step S2. The TS header 194 is formed in a manner such that the parameters such as error indicator 200 indicative of the presence or absence of the bit error in the packet, unit start indicator 202 indicating whether a new PES packet appears in the TS payload or not, packet ID 206 indicating whether the packet in the payload is video or audio, adaptation field 214, and adaptation field control 210 indicative of the presence or absence of the payload 196 are set in accordance with the state of the payload storage data of the 184-byte unit divided from the MPEG2-PES packet which had been read in step S1.”

Hiroshima, 11:13–27.

Exhibit C-8

FIG. 15



Hiroshima, Fig. 15.

Haartsen discloses and/or renders obvious this limitation. *See, e.g.:*

“The present invention fulfills the above-described and other needs by providing a multi-part digital preamble for use in transmitting digital data packets. Advantageously, a preamble according to the invention substantially reduces the overall length of DC-free sequences that are transmitted with each data packet in a digital FM system. According to exemplary embodiments, a multi-part preamble includes a short, substantially DC-free leading part which enables coarse estimation of a prevailing DC level at an FM receiver. The exemplary preamble also includes a synchronization


Exhibit C-8

part which is not necessarily substantially DC-free and which carries timing and/or other useful information (e.g. , channel identification, destination address, etc.) that can vary from data packet to data packet (e.g. , depending on sender, recipient, time of day, etc.).”

Haartsen, 5:2–15.

“An exemplary DC estimation processor is depicted in Figure 3, wherein an exemplary FM demodulation processor 300 is shown to include a correlation processor 310, a resistor 320, a switch 330 and a capacitor 340, as well as the FM detector 210, the low-pass filter 130 and the slicer 230 of Figure 2.”

Haartsen, 9:11–14.



Haartsen, Fig. 3.

“To alleviate this problem, the present invention discloses that a digital preamble can be strategically divided into three or more parts so that synchronization and other information can be included within the preamble while still enabling accurate DC estimation. Specifically, a short leading DC-free part of the preamble (e.g., alternating logical 1's and 0's, or any other sequence including an equal number of logical 1's and logical 0's) can be used to provide initial, coarse DC offset estimation. Thereafter, a synchronization part of the preamble, which is not necessarily DC free, but which is protected by code to prevent detection errors, can be used not only to synchronize


Exhibit C-8

detection of a data packet, but also to carry timing and/or other useful information (e.g., channel identification, destination address, time of day, user identification, etc.).”

Haartsen, 12:13–13:2.

“An exemplary multi-part preamble according to the invention is depicted in Figure 5. In the figure, a three-part preamble 500 is shown to include a DC-free leading sequence 510, followed by a multiple-bit synchronization word 520 which is not necessarily DC free and which can be different for each data packet. The synchronization word 520, which can include timing and/or other information and which is coded to protect against detection errors, is in turn followed by a trailing DC-free sequence 530. Although the leading and trailing DC-free sequences 510, 530 are shown in the figure as 4-bit sequences of alternating logical 1's and 0's, those of ordinary skill will appreciate that each DC-free sequence 510, 530 can, in practice, be shorter or longer in length and can include any DC-free sequence having an equal number of logical 1's and logical 0's.”

Haartsen, 13:13–14:2.



Haartsen, Fig. 5.

“When the exemplary preamble 500 is received by the system 300 of Figure 3, the leading DC-free sequence 510 induces a coarse estimate of the baseband DC offset on the storage capacitor 340. Although the estimate is coarse, it is sufficient to permit successful detection of the coded synchronization word 520 (e.g., by the correlation processor 310). Methods for coding of a synchronization word to protect against erroneous detection in the presence of some DC-offset are


Exhibit C-8

	known to those of skill in the art. Once the correlation processor 310 has detected the synchronization word 520, timing between received bits is well known. Consequently, the correlation processor 310 can open the switch 330 precisely at the end of the trailing DC-free sequence 530 to thereby store an accurate estimate of the DC offset for reception and detection of the incoming digital data packet. As is described above, the fact that the DC-free sequences 510, 530 can be very short, combined with the fact that the synchronization word 520 can carry useful information, means that the overhead associated with the exemplary three- part preamble 500 is considerably reduced as compared to a conventional DC-free preamble. Those of skill in the art will appreciate that the precise lengths of the three parts 510, 520, 530 of the preamble 500 shown in Figure 5 can vary and are a matter of design choice. Additionally, the exemplary preamble 500 can include additional trailing DC-free sequences which can, for example, be embedded within the synchronization word 520.” Haartsen, 14:3–15:2. Shin discloses and/or renders obvious this limitation. <i>See, e.g.:</i> “It may thus be appreciated that each eight-bit code converted into a 10-bit encoded character during high-transition mode operation includes four or more logical transitions. Similarly, each eight-bit code converted into a 10-bit encoded character during low-transition mode operation includes less than four logical transitions. These sets of 10-bit encoded characters capable of being produced during high-transition and low-transition modes of operation may be characterized as a high-transition set of "in-band" encoded characters, and a low-transition set of in-band encoded characters, respectively. Beyond the 256 in-band characters within the high-transition set and the 256 in-band characters of the low-transition set, there exists a high-transition set of 256 out-of-band 10-bit characters and a low-transition set of 10-bit characters. In accordance with another aspect of the invention, various synchronization and other special characters are defined using the high-transition and low-transition sets of out-of-band characters. Each character corresponding to one of these "out-of-band" characters associated with the high-transition set includes less than four logical transitions, and each of the out-of-band characters associated with the low-transition set of code characters includes more than four logical transitions. The difference in the number of transitions
--	--

Exhibit C-8

	between in-band and out-of-band characters allows selected out-of-band characters to serve as control characters, and to be readily distinguished from in-band characters within the transmitted data stream.” Shin, 9:49–10:8. “In accordance with one aspect of the invention, the synchronization characters associated with both the low-transition and high-transition sets of code characters are selected so as to facilitate rapid synchronization during data recovery. When the low-transition set of code characters is being employed, a special group of out-of-band characters is used during synchronization. Each special synchronization character includes a predefined number larger than four (e.g., 7) of logical transitions between binary character values, and a predefined number (e.g., two) of "non-transitions" between character values. As will be discussed below, the special synchronization characters are selected such that random logic may be used to distinguish each special synchronization character from the in-band characters of the low-transition set.” Shin, 10:17–34. “Referring to FIG. 6, the deserializer 34 receives the sequence of 10-bit characters T[9:0] produced by the encoder and generates 10-bit parallel received data upon bit lines RX9, RX8, . . . , RX0 (i.e., RX[9:0]). This 10-bit parallel received data is provided over bit lines RX[9:0] to the decoder 44, as well as to a decoder synchronization module 114. As is described below in section IV, the synchronization module 114 is operative to ascertain boundaries within the 10-bit parallel received data corresponding to the frame boundaries of the transmitted data (i.e., to T[9:0]). Specifically, synchronization module 114 determines upon which of the bit lines RX[9:0] the deserializer 34 is providing the received bits corresponding to the first bit T[0] of each transmitted byte T[9:0]. Upon making this determination, the synchronization module 114 provides a frame boundary pointer 118 to decoder 44 identifying the one of the bit lines RX[9:0] corresponding to the first bit T[0] of each transmitted 10-bit character T[9:0]. Upon receiving this synchronization information, the decoder 44 is disposed to decode the received data RX[9:0] in the following manner.” Shin, 13:37–57.
--	--

Exhibit C-8




Shin, Fig. 6.

Austin discloses and/or renders obvious this limitation. *See, e.g.:*

“FIG. 2 shows the structure of a transmitter suitable for transformer/rectifier sites. The transmitter 30 comprises waveform generators 31-34 which transmitter receives the data at input 35 which selects 10 Hz tone as an indication of data ‘1’ and a 15 Hz tone as a data ‘0’ via selector 36. Using Fast Frequency Shift Keying (FFSK) at 10 baud, the modulated data signal is applied by modulating part of the current applied to the pipe (current switch 37) which provides cathodic protection (CP) as a result of the pipeline resistance. The signals are in digital form with multistage representations of the two modulating tones. Each state of these tones represents the switching levels for the transformer/rectifier pipeline devices. The bandwidth of the pipe is such that the switching harmonics are considerably reduced through their passage in the pipe so are of no great concern at the receiver.”

Austin, 2:32-47.

Exhibit C-8



Austin, Fig. 2.

“This system uses up to 5 pre-defined data packet sizes and up to 16 user defined data message types within each packet. Each packet has Forward Error Connection (FEC) and bit interleaving applied to it on transmission. Error checking is performed using a standard CRC. The system uses standard run in and byte synchronisation techniques. The byte synchronisation word used is 32_bits long and there are 5_different words used, one to identify each of the different packet types.”

Austin, 4:37-45.

To the extent that Plaintiff contends that the Additional Prior Art References do not explicitly or inherently disclose this element, or render obvious this element, this element is obvious in combination with one or more references and/or with the knowledge of a person of ordinary skill in the art, as set forth in Defendant’s Invalidity Contentions.

Claim 22

Exhibit C-8

22[p]	The communications device of claim 19	The Additional Prior Art References disclose, expressly or inherently, the transmitting device of claim 19 under any proposed construction. See exemplary disclosures and corresponding discussion for claim 19.
22	wherein the synchronization symbol is transmitted before transmitting the symbols of the packet.	Chen discloses and/or renders obvious this limitation. <i>See, e.g.:</i> “FIG. 5 illustrates the processing flow of the Insertion Processing Unit of FIG. 4 in accordance with the present invention. Note that the steps of FIG. 5 are shown as occurring in a sequential manner for clarity, but many of the steps may occur concurrently and/or sequentially. At block 502, the IPU outputs the main transport stream from the main stream buffer (MSB) 480. The IPU is essentially in a bypass mode at this time since the insertion stream has not yet been inserted. At block 506, the main stream parser parses the main stream to locate the splicing signal T-- in. Note that T-- in need not be carried in the main stream, but may be provided locally, for example, at a cable system headend. When T-- in is detected, the IPU continues to output the main transport stream from the main stream buffer. At block 508, the main stream parser parses the main stream to locate the closest packet with an anchor picture (e.g., I or P picture) start code and passes that packet to the syntax processor 470. This packet will be the last packet of the main stream. To ensure that the resulting data stream is MPEG compliant, no data from the next frame or sequence of frames may be included in this last main stream packet. At block 512, the syntax processor adds a sequence end code field to the last packet of the main stream. The sequence end code is a twenty-four bit code which indicates the end of a sequence of pictures. At block 514, the PSI/PID replacer 435 starts to replace the PSI/PID of the insertion stream.” Chen, 12:28–53.

Exhibit C-8




FIG.5

Chen, Fig. 5.

“In the MPEG-2 system (and the similar DigiCipher® II system proprietary to General Instrument Corporation, the assignee hereof) a transport stream, or transport multiplex is made up of a


Exhibit C-8

	contiguous set of fixed length packets. Each packet is 188 total bytes in length, with the first four of those bytes being defined as the packet header. The payload portion of each packet is thus normally 184 bytes. However, a variable length adaptation field may be provided to extend the header, when required. When an adaptation field is present, the payload portion of the packet will be correspondingly shorter. Various timing and identification information is provided in different portions of the transport stream. These include a packet identifier (PID) found in the transport header of each transport packet to provide a reference number for identifying the transport packets carrying a specific service component. This number is included in a service definition or "service map" used by the receiver to identify those transport packets required to reconstruct a television program signal. The PID may also be referenced for various grooming and remultiplexing functions. In the case of video, audio or isochronous control data, the stream of packets labeled with a single PID represents a single video, audio or isochronous data service elementary stream, respectively. Each type of packet will have a different PID identifying the packet type. Timing information carried by the transport stream includes a program clock reference (PCR) which effectively represents a sample of the system time clock (STC) time base that underlies the service composed of the PIDs referenced in the service map. The PID carrying the packet with the PCR is also referenced in the service map. The video, audio and isochronous data components of a service are locked through a defined relationship to the system time clock. The PCR serves to define the transport rate, in the sense that between any two successive PCRs in one PID, the transport rate is constant and nominally equal to the system time clock rate times the ratio of the total number of transport bytes between the PCRs divided by the difference in the PCRs in units of system time clock ticks." Chen, 14:45–15:17. "As indicated in FIG. 6b, the transport header of each transport packet includes the PID which identifies the particular service component carried by the transport packet. The PCR will be carried by an adaptation field to provide timing information for a desired service. At a receiver, the PCR for the desired service is detected from the adaptation field. The PIDs of the transport packets will
--	---

Exhibit C-8

then be monitored in accordance with the timing established by the PCR to recover those transport packets carrying a particular component of the service to be processed.”

Chen, 16:41–50.




Chen, Fig. 6b.

Cunningham discloses and/or renders obvious this limitation. *See, e.g.:*

“Referring to FIG. 1, a data frame intended to be communicated over an optical fibre channel 20 typically comprises a binary digital message (payload) to be transferred between stations on the network; this message may comprise user data, or control information for managing operation of the network and its components. A circuit 12 generates an associated cyclic redundancy check (CRC) block containing check data derived from this message in known manner in accordance with a predetermined CRC algorithm, and appends the CRC block to the message.”

Cunningham, 2:65–3:7.

Exhibit C-8



Cunningham, Fig. 1.

“Control information is conveyed by using selected n-bit data block values, together with a distinctive value of 10 for the pair of label bits a and b to indicate that the associated data block represents control information rather than actual data. The data block values used for this purpose are selected to be values which inherently have a disparity of zero, such as (for 8-bit blocks)


Exhibit C-8

	00001111 or 10101010. It will be noted that the label bit pair value 10 likewise has zero disparity, so the complete encoded block also has zero disparity.” Cunningham, 4:23–31. “The encoded blocks are transmitted over the optical fibre channel 20 (leftmost symbol first), and upon receipt at their destination they are detected by a receiver 22 where compensation is made in known manner for the effects of timing drift and other signal degradation to produce a restored digital data stream.” Cunningham, 4:45–50. “There are many potential ways to achieve code block synchronization (i.e. correct alignment in the receiver of the label and payload fields). For example, one technique would be to transmit a ‘comma sequence’ during start-up. As an example of this technique consider a code with n=8 using the simple control payload scheme. A training control state of 1111 1111 would be transmitted as: [10 1111 1111 0000 0000][10 1111 1111 0000 0000][10 1111 1111 0000 0000] The receiver searches for the bit pair 10. This is found in only two locations: the correct location (the label field), and at the end of the run of eight 1's. However, if the receiver (incorrectly) identifies the label field to be the 10 bit pair at the end of the run of 1's, the following 16 bits are not a valid control payload (first 8 bits are not the complement of the second 8 bits). The receiver can therefore detect the mistake. It is also noted that this sequence contains a regular 010 pattern, which aids in phase-locked loop (PLL) clock acquisition.” Cunningham, 7:5–23. Hiroshima discloses and/or renders obvious this limitation. <i>See, e.g.:</i> “FIG. 11 shows the details of the MPEG2-TS. Each packet of the MPEG2-TS has a prefix length of 188 bytes. The header 194 consists of four bytes and the payload 196 consists of 184 bytes. The header 194 is constructed by sync bytes 198, an error indicator 200, a unit start indicator 202, a transport priority 204, a packet ID 206, a scramble control 208, an adaptation field control 210, a
--	---

Exhibit C-8

continuity counter 212, and an adaptation field 214. The adaptation field 214 is constructed by an adaptation field length 216, a discontinuity indicator 218, a random access indicator 220, an elementary stream priority indicator 222, a 5-flag 224, an optional field 226, and stuffing bytes 228. Among them, the optional field 226 is constructed by a program clock reference 230, an original program clock reference field 232, a splice count down 234, a private data length 236, private data 238, an adaptation field extension 240, 3-flag 241, and an optional field 242.”

Hiroshima, 10:51–67.



Hiroshima, Fig. 11.

Haartsen discloses and/or renders obvious this limitation. *See, e.g.:*

Exhibit C-8

“To alleviate this problem, the present invention discloses that a digital preamble can be strategically divided into three or more parts so that synchronization and other information can be included within the preamble while still enabling accurate DC estimation. Specifically, a short leading DC-free part of the preamble (e.g., alternating logical 1's and 0's, or any other sequence including an equal number of logical 1's and logical 0's) can be used to provide initial, coarse DC offset estimation. Thereafter, a synchronization part of the preamble, which is not necessarily DC free, but which is protected by code to prevent detection errors, can be used not only to synchronize detection of a data packet, but also to carry timing and/or other useful information (e.g., channel identification, destination address, time of day, user identification, etc.).”

Haartsen, 12:13–13:2.

“An exemplary multi-part preamble according to the invention is depicted in Figure 5. In the figure, a three-part preamble 500 is shown to include a DC-free leading sequence 510, followed by a multiple-bit synchronization word 520 which is not necessarily DC free and which can be different for each data packet. The synchronization word 520, which can include timing and/or other information and which is coded to protect against detection errors, is in turn followed by a trailing DC-free sequence 530. Although the leading and trailing DC-free sequences 510, 530 are shown in the figure as 4-bit sequences of alternating logical 1's and 0's, those of ordinary skill will appreciate that each DC-free sequence 510, 530 can, in practice, be shorter or longer in length and can include any DC-free sequence having an equal number of logical 1's and logical 0's.”

Haartsen, 13:13–14:2.




Exhibit C-8

		Haartsen, Fig. 5. “When the exemplary preamble 500 is received by the system 300 of Figure 3, the leading DC-free sequence 510 induces a coarse estimate of the baseband DC offset on the storage capacitor 340. Although the estimate is coarse, it is sufficient to permit successful detection of the coded synchronization word 520 (e.g., by the correlation processor 310). Methods for coding of a synchronization word to protect against erroneous detection in the presence of some DC-offset are known to those of skill in the art. Once the correlation processor 310 has detected the synchronization word 520, timing between received bits is well known. Consequently, the correlation processor 310 can open the switch 330 precisely at the end of the trailing DC-free sequence 530 to thereby store an accurate estimate of the DC offset for reception and detection of the incoming digital data packet.” Haartsen, 14:3–14. To the extent that Plaintiff contends that the Additional Prior Art References do not explicitly or inherently disclose this element, or render obvious this element, this element is obvious in combination with one or more references and/or with the knowledge of a person of ordinary skill in the art, as set forth in Defendant’s Invalidity Contentions.
Claim 25		
25[p]	The communications device of claim 19	Chen discloses, expressly or inherently, the transmitting device of claim 19 under any proposed construction. <i>See exemplary disclosures and corresponding discussion for claim 19.</i>
25	wherein the packet types include a control packet.	Chen discloses and/or renders obvious this limitation. <i>See, e.g.:</i> “Various timing and identification information is provided in different portions of the transport stream. These include a packet identifier (PID) found in the transport header of each transport packet to provide a reference number for identifying the transport packets carrying a specific service component. This number is included in a service definition or "service map" used by the receiver to identify those transport packets required to reconstruct a television program signal. The PID may also be referenced for various grooming and remultiplexing functions. In the case of video, audio

Exhibit C-8

	or isochronous control data, the stream of packets labeled with a single PID represents a single video, audio or isochronous data service elementary stream, respectively. Each type of packet will have a different PID identifying the packet type.” Chen, 14:56–15:3. “The PES payload contains the information data that is desired to be transmitted to a receiver. Thus, the payload includes all of the video, audio and control information necessary for the receiver to decode and reconstruct, e.g., a digital television signal). In order to meet the requirements of robustness and simplicity, a fixed packet length approach is preferred to the variable length PES packets. Thus, as illustrated in FIG. 6b, the packet elementary stream containing the PES packets 670 is reconfigured into a stream of fixed length transport packets 680. The transport packets illustrated in FIG. 6b all correspond to the same service component, such as the video component of a digital television transmission. In the MPEG-2 and DigiCipher® II embodiments, each packet is 188 total bytes in length, with the first four bytes comprising a transport packet header (TP HDR) 682. The payload portion 684 of each packet 680 is thus normally 184 bytes. However, an adaptation field mechanism is present, as illustrated by transport packet 680', to extend the header when required. The adaptation field 686 provides additional information which is not required for every transport packet. The adaptation field (ADPT FIELD) 686 extends the regular transport header 682 at the expense of payload 684, which will be less than 184 bytes whenever the adaptation is provided. The adaptation field 686 is of variable length, depending on the information it contains. Typically, the adaptation field will support additional information for time base recovery and other functions, and also provides a mechanism for padding the payload when it does not occupy the full 184 bytes. Such padding can be used, for example, to make a variable rate video packetized elementary stream into a constant rate transport stream.” Chen, 16:9–40.
--	--

Exhibit C-8




FIG.6b

Chen, Fig. 6b.

Cunningham discloses and/or renders obvious this limitation. *See, e.g.:*

“Control information is conveyed by using selected n-bit data block values, together with a distinctive value of 10 for the pair of label bits a and b to indicate that the associated data block represents control information rather than actual data. The data block values used for this purpose are selected to be values which inherently have a disparity of zero, such as (for 8-bit blocks) 00001111 or 10101010. It will be noted that the label bit pair value 10 likewise has zero disparity, so the complete encoded block also has zero disparity.”

Cunningham, 4:23–31.

To the extent that Plaintiff contends that the Additional Prior Art References do not explicitly or inherently disclose this element, or render obvious this element, this element is obvious in combination with one or more references and/or with the knowledge of a person of ordinary skill in the art, as set forth in Defendant’s Invalidity Contentions.

Claim 28

28[p]

The communications device of claim 19

Chen discloses, expressly or inherently, the transmitting device of claim 19 under any proposed construction.

See exemplary disclosures and corresponding discussion for claim 19.

Exhibit C-8

28	wherein the synchronization primitive is made up of a sequence of bit-string synchronization symbols.	Cunningham discloses and/or renders obvious this limitation. <i>See, e.g.:</i> “There are many potential ways to achieve code block synchronization (i.e. correct alignment in the receiver of the label and payload fields). For example, one technique would be to transmit a ‘comma sequence’ during start-up. As an example of this technique consider a code with $n=8$ using the simple control payload scheme. A training control state of 1111 1111 would be transmitted as: [10 1111 1111 0000 0000][10 1111 1111 0000 0000][10 1111 1111 0000 0000] The receiver searches for the bit pair 10. This is found in only two locations: the correct location (the label field), and at the end of the run of eight 1's. However, if the receiver (incorrectly) identifies the label field to be the 10 bit pair at the end of the run of 1's, the following 16 bits are not a valid control payload (first 8 bits are not the complement of the second 8 bits). The receiver can therefore detect the mistake. It is also noted that this sequence contains a regular 010 pattern, which aids in phase-locked loop (PLL) clock acquisition.” Cunningham, 7:5–23. Shin discloses and/or renders obvious this limitation. <i>See, e.g.:</i> “During each preamble period, three repetitions of the same synchronization character are produced by the encoder 18. As is described herein, by processing the 21 bits most recently produced by the encoder 14, synchronization module 114 is able to detect at least one of the three repetitions of the synchronization character transmitted during a given preamble period. This advantageously allows synchronization to be achieved within a relatively short preamble period.” Shin, 15:20–28. “In the exemplary embodiment a preamble sequence (i.e., three repetitions of the same out-of-band SYNC character) is sent upon system power-up as well as during lapses in data transmission over the serial link 30. This allows timing synchronization to be maintained between the encoder 18 and decoder 44 even in the absence of data transmission for extended periods.”
----	---	--

Exhibit C-8

	Shin, 18:19–25. Austin discloses and/or renders obvious this limitation. <i>See, e.g.:</i> “This system uses up to 5 pre-defined data packet sizes and up to 16 user defined data message types within each packet. Each packet has Forward Error Connection (FEC) and bit interleaving applied to it on transmission. Error checking is performed using a standard CRC. The system uses standard run in and byte synchronisation techniques. The byte synchronisation word used is 32 _bits long and there are 5_different words used, one to identify each of the different packet types.” Austin, 4:37-45. To the extent that Plaintiff contends that Chen does not explicitly or inherently disclose this element, or render obvious this element, this element is obvious in combination with one or more references and/or with the knowledge of a person of ordinary skill in the art, as set forth in Defendant’s Invalidity Contentions.
--	--

Exhibit C-9

Invalidity of the '798 Patent in View of IEEE Std 802.3, 1998 Edition (802.3-1998)

This chart provides examples of where the prior art reference discloses, either expressly or inherently, the claim elements of asserted U.S. Patent No. 7,746,798 (“the ’798 patent”). The citations herein are exemplary and are not exhaustive; Defendant has cited the most relevant portions of the identified prior art. For each citation, Defendant hereby incorporates by reference the entire cited prior art reference, and reserves the right to rely on uncited portions of each identified prior art reference. To provide context or identify the state of the art at the time of the alleged invention, Defendant reserves the right to rely on: (1) uncited portions of the identified prior art; (2) other prior art not identified herein; (3) references that show the state of the art (irrespective of whether such references themselves qualify as prior art); (4) factual testimony from the inventors or authors of the prior art references, or purveyors of prior art devices; and (5) expert testimony.

Where Defendant cites a figure, the citation includes any text associated with the figure. Where Defendant cites passages relating to a figure, the citation includes the associated figure any other text associated with that figure. Where Defendant cites particular prior art as disclosing a claim element, that citation equally applies to all identical or similar claim elements in this chart. Defendant reserves the right to further assert that a claim element is inherent, the element would have been obvious to a person of ordinary skill in the art, or the element is disclosed in other prior art that may be combined with the cited prior art. To establish the inherency of certain features of the prior art to invalidate the asserted claims, Defendant may rely on cited or uncited portions of the prior art, other documents, factual testimony, and expert testimony.

This invalidity chart is based on information presently available to Defendant, and the information herein is provided as a non-limiting example.

IEEE Std 802.3 is a working group of standard specifications for Ethernet, a method of packet-based physical communication in a local area network maintained by the Institute of Electrical and Electronics Engineers (IEEE). Its 1998 Edition was first implemented in systems in the United States by at least 1998. 802.3-1998 accordingly is prior art to the ’798 patent under pre-AIA 35 U.S.C. 102(a)(1), and for at least the reasons described below, invalidates the asserted claim 11, 19, 22, 25, 28 of the ’798 patent. Additionally, devices employing the IEEE Std 802.3 were offered for sale before the ’798 patent. For example, the 1000BASE-T Gigabit Ethernet Connection employing the IEEE Std 802.3 was included in Apple’s Power Mac G4 no later than October 2000.

Exhibit C-11

Claim Limitation		802.3-1998
Claim 11		
11[p]	A method for identifying packet types of packets of symbols, the method comprising:	To the extent the preamble is limiting, 802.3-1998 discloses this element and/or renders it obvious either alone or in combination with other references and/or the knowledge of one of ordinary skill in the art. <i>See, e.g.:</i> “24.2.2 Functional requirements The PCS comprises the Transmit, Receive, and Carrier Sense functions for 100BASE-T. In addition, the collisionDetect signal required by the MAC (COL on the MII) is derived from the PMA code-bit stream. The PCS shields the Reconciliation sublayer (and MAC) from the specific nature of the underlying channel. Specifically for receiving, the 100BASE-X PCS passes to the MII a sequence of data nibbles derived from incoming code-groups, each comprised of five code-bits, received from the medium. Code-group alignment and MAC packet delimiting is performed by embedding special non-data code-groups. The MII uses a nibble- wide, synchronous data path, with packet delimiting being provided by separate TX_EN and RX_DV signals. The PCS provides the functions necessary to map these two views of the exchanged data. The process is reversed for transmit. The following provides a detailed specification of the functions performed by the PCS, which comprise five parallel processes (Transmit, Transmit Bits, Receive, Receive Bits, and Carrier Sense). Figure 24-4 includes a functional block diagram of the PCS. The Receive Bits process accepts continuous code-bits via the PMA_UNITDATA.indicate primitive. Receive monitors these bits and generates RXD <3:0>, RX_DV and RX_ER on the MII, and the internal flag, receiving, used by the Carrier Sense and Transmit processes. The Transmit process generates continuous code-groups based upon the TXD <3:0>, TX_EN, and TX_ER signals on the MII. These code-groups are transmitted by Transmit Bits via the PMA_UNITDATA.request primitive. The Transmit process generates the MII signal COL based on whether a reception is occurring simultaneously with transmission. Additionally, it generates the internal flag, transmitting, for use by the Carrier Sense process. The Carrier Sense process asserts the MII signal CRS when either transmitting or receiving is TRUE. Both the Transmit and Receive processes monitor link_status via the PMA_LINK.indicate primitive, to account for potential link failure conditions.

Exhibit C-11

		24.2.2.1 Code-groups The PCS maps four-bit nibbles from the MII into five-bit code-groups, and vice versa, using a 4B/5B block coding scheme. A code-group is a consecutive sequence of five code-bits interpreted and mapped by the PCS. Implicit in the definition of a code-group is an establishment of code-group boundaries by an alignment function within the PCS Receive process. It is important to note that, with the sole exception of the SSD, which is used to achieve alignment, code-groups are undetectable and have no meaning outside the 100BASE-X physical protocol data unit, called a “stream.” The coding method used, derived from ISO/IEC 9314-1, provides a) Adequate codes (32) to provide for all Data code-groups (16) plus necessary control code-groups; b) Appropriate coding efficiency (4 data bits per 5 code-bits; 80%) to effect a 100 Mb/s Physical Layer interface on a 125 Mb/s physical channel as provided by FDDI PMDs; and c) Sufficient transition density to facilitate clock recovery (when not scrambled). Table 24-1 specifies the interpretation assigned to each five bit code-group, including the mapping to the nibble- wide (TXD or RXD) Data signals on the MII. The 32 code-groups are divided into four categories, as shown. For clarity in the remainder of this clause, code-group names are shown between /slashes/. Code-group sequences are shown in succession, e.g.: /1/2/.... The indicated code-group mapping is identical to ISO/IEC 9314-1: 1989, with four exceptions: a) The FDDI term symbol is avoided in order to prevent confusion with other 100BASE-T terminology. In general, the term code-group is used in its place. b) The /S/ and /Q/ code-groups are not used by 100BASE-X and are interpreted as INVALID. c) The /R/ code-group is used in 100BASE-X as the second code-group of the End-of-Stream delimiter rather than to indicate a Reset condition. d) The /H/ code-group is used to propagate receive errors rather than to indicate the Halt Line State. 24.2.2.1.1 Data code-groups A Data code-group conveys one nibble of arbitrary data between the MII and the PCS. The sequence of Data code-groups is arbitrary, where any Data code-group can be followed by any other Data code-group. Data code-groups are coded and decoded but not interpreted by the PCS. Successful decoding of Data code-groups depends on proper receipt of the Start-of-Stream delimiter sequence, as defined in Table 24-1. 24.2.2.1.2 Idle code-groups The Idle code-group (/I/) is transferred between streams. It provides a continuous fill pattern to establish and maintain clock synchronization. Idle code-groups are emitted from, and interpreted by, the PCS. 24.2.2.1.3 Control code-groups
--	--	---

Exhibit C-11

		The Control code-groups are used in pairs (/J/K/, /T/R/) to delimit MAC packets. Control code-groups are emitted from, and interpreted by, the PCS. 24.2.2.1.4 Start-of-Stream Delimiter (/J/K/) A Start-of-Stream Delimiter (SSD) is used to delineate the boundary of a data transmission sequence and to authenticate carrier events. The SSD is unique in that it may be recognized independently of previously established code-group boundaries. The Receive function within the PCS uses the SSD to establish code-group boundaries. A SSD consists of the sequence /J/K/. On transmission, the first 8 bits of the MAC preamble are replaced by the SSD, a replacement that is reversed on reception. 24.2.2.1.5 End-of-Stream delimiter (/T/R/) An End-of-Stream delimiter (ESD) terminates all normal data transmissions. Unlike the SSD, an ESD cannot be recognized independent of previously established code-group boundaries. An ESD consists of the sequence /T/R/.
--	--	--

Exhibit C-11

Table 24-1—4B/5B code-groups

	PCS code-group [4:0] 4 3 2 1 0	Name	MII (TXD/RXD) <3:0> 3 2 1 0	Interpretation
D A T A	1 1 1 1 0	0	0 0 0 0	Data 0
	0 1 0 0 1	1	0 0 0 1	Data 1
	1 0 1 0 0	2	0 0 1 0	Data 2
	1 0 1 0 1	3	0 0 1 1	Data 3
	0 1 0 1 0	4	0 1 0 0	Data 4
	0 1 0 1 1	5	0 1 0 1	Data 5
	0 1 1 1 0	6	0 1 1 0	Data 6
	0 1 1 1 1	7	0 1 1 1	Data 7
	1 0 0 1 0	8	1 0 0 0	Data 8
	1 0 0 1 1	9	1 0 0 1	Data 9
	1 0 1 1 0	A	1 0 1 0	Data A
	1 0 1 1 1	B	1 0 1 1	Data B
	1 1 0 1 0	C	1 1 0 0	Data C
	1 1 0 1 1	D	1 1 0 1	Data D
	1 1 1 0 0	E	1 1 1 0	Data E
	1 1 1 0 1	F	1 1 1 1	Data F
C O N T R O L	1 1 1 1 1	I	undefined	IDLE; used as inter-stream fill code
	1 1 0 0 0	J	0 1 0 1	Start-of-Stream Delimiter, Part 1 of 2; always used in pairs with K
	1 0 0 0 1	K	0 1 0 1	Start-of-Stream Delimiter, Part 2 of 2; always used in pairs with J
	0 1 1 0 1	T	undefined	End-of-Stream Delimiter, Part 1 of 2; always used in pairs with R
I N V A L I D	0 0 1 1 1	R	undefined	End-of-Stream Delimiter, Part 2 of 2; always used in pairs with T
	0 0 1 0 0	H	Undefined	Transmit Error; used to force signaling errors
	0 0 0 0 0	V	Undefined	Invalid code
	0 0 0 0 1	V	Undefined	Invalid code
	0 0 0 1 0	V	Undefined	Invalid code
	0 0 0 1 1	V	Undefined	Invalid code
	0 0 1 0 1	V	Undefined	Invalid code
	0 0 1 1 0	V	Undefined	Invalid code
	0 1 0 0 0	V	Undefined	Invalid code
	0 1 1 0 0	V	Undefined	Invalid code
	1 0 0 0 0	V	Undefined	Invalid code
	1 1 0 0 1	V	Undefined	Invalid code

24.2.2.1.6 Invalid code-groups

Exhibit C-11

		The /H/ code-group indicates that the PCS's client wishes to indicate a Transmit Error to its peer entity. The normal use of this indicator is for repeaters to propagate received errors. Transmit Error code-groups are emitted from the PCS, at the request of the PCS's client through the use of the TX_ER signal, as described in 24.2.4.2. The presence of any invalid code-group on the medium, including /H/, denotes a collision artifact or an error condition. Invalid code-groups are not intentionally transmitted onto the medium by DTE's. The PCS indicates the reception of an Invalid code-group on the MII through the use of the RX_ER signal, as described in 24.2.4.4.” 802.3-1998, p616-619.
--	--	--

Exhibit C-11




Figure 24-11—Receive state diagram

802.3-1998, p628.

“24.3.1.1 PMA_TYPE.indicate

Exhibit C-11

	This primitive is generated by the PMA to indicate the nature of the PMA instantiation. The purpose of this primitive is to allow clients to support connections to the various types of 100BASE-T PMA entities in a generalized manner. 24.3.1.1.1 Semantics of the service primitive PMA_TYPE.indicate (pma_type) The pma_type parameter for use with a 100BASE-X PMA is “X”. 802.3-1998, p629. “36.2.2 Functions within the PCS The PCS comprises the PCS Transmit, Carrier Sense, Synchronization, PCS Receive, and Auto-Negotiation processes for 1000BASE-X. The PCS shields the Reconciliation sublayer (and MAC) from the specific nature of the underlying channel. When communicating with the GMII, the PCS uses an octet-wide, synchronous data path, with packet delimiting being provided by separate transmit control signals (TX_EN and TX_ER) and receive control signals (RX_DV and RX_ER). When communicating with the PMA, the PCS uses a ten-bit wide, synchronous data path, which conveys ten-bit code-groups. At the PMA Service Interface, code-group alignment and MAC packet delimiting are made possible by embedding special non-data code-groups in the transmitted code-group stream. The PCS provides the functions necessary to map packets between the GMII format and the PMA Service Interface format. The PCS Transmit process continuously generates code-groups based upon the TXD <7:0>, TX_EN, and TX_ER signals on the GMII, sending them immediately to the PMA Service Interface via the PMA_UNITDATA.request primitive. The PCS Transmit process generates the GMII signal COL based on whether a reception is occurring simultaneously with transmission. Additionally, it generates the internal flag, transmitting, for use by the Carrier Sense process. The PCS Transmit process monitors the Auto-Negotiation process xmit flag to determine whether to transmit data or reconfigure the link. The Carrier Sense process controls the GMII signal CRS (see Figure 36-8). The PCS Synchronization process continuously accepts code-groups via the PMA_UNITDATA.indicate primitive and conveys received code-groups to the PCS Receive process via the SYNC_UNITDATA.indicate primitive. The PCS Synchronization process sets the sync_status flag to indicate whether the PMA is functioning dependably (as well as can be determined without exhaustive error-rate analysis). The PCS Receive process continuously accepts code-groups via the SYNC_UNITDATA.indicate primitive. The PCS Receive process monitors these code-groups and generates RXD <7:0>, RX_DV, and RX_ER on the GMII, and the internal flag, receiving, used by the Carrier Sense and Transmit processes. The
--	---

Exhibit C-11

	PCS Auto-Negotiation process sets the xmit flag to inform the PCS Transmit process to either transmit normal idles interspersed with packets as requested by the GMII or to reconfigure the link. The PCS Auto- Negotiation process is specified in Clause 37. 36.2.3 Use of code-groups The PCS maps GMII signals into ten-bit code groups, and vice versa, using an 8B/10B block coding scheme. Implicit in the definition of a code-group is an establishment of code-group boundaries by a PMA code-group alignment function as specified in 36.3.2.4. Code-groups are unobservable and have no meaning outside the PCS. The PCS functions ENCODE and DECODE generate, manipulate, and interpret codegroups as provided by the rules in 36.2.4. 36.2.4 8B/10B transmission code The PCS uses a transmission code to improve the transmission characteristics of information to be transferred across the link. The encodings defined by the transmission code ensure that sufficient transitions are present in the PHY bit stream to make clock recovery possible at the receiver. Such encoding also greatly increases the likelihood of detecting any single or multiple bit errors that may occur during transmission and reception of information. In addition, some of the special code-groups of the transmission code contain a distinct and easily recognizable bit pattern that assists a receiver in achieving code-group alignment on the incoming PHY bit stream. The 8B/10B transmission code specified for use in this standard has a high transition density, is a run-length-limited code, and is dc-balanced. The transition density of the 8B/10B symbols ranges from 3 to 8 transitions per symbol. The definition of the 8B/10B transmission code in this standard is identical to that specified in ANSI X3.230-1994 (FC-PH), Clause 11. The relationship of code-group bit positions to PMA and other PCS constructs is illustrated in Figure 36-3.
--	--

Exhibit C-11




Figure 36-3—PCS reference diagram

36.2.4.1 Notation conventions


8B/10B transmission code uses letter notation for describing the bits of an unencoded information octet and a single control variable. Each bit of the unencoded information octet contains either a binary zero or a binary one. A control variable, Z, has either the value D or the value K. When the control variable associated with an unencoded information octet contains the value D, the associated encoded code-group is referred to as a data code-group. When the control variable associated with an unencoded information octet contains the value K, the associated encoded code-group is referred to as a special code-group. The bit notation of A,B,C,D,E,F,G,H for an unencoded information octet is used in the description of the 8B/10B transmission code. The bits A,B,C,D,E,F,G,H are translated to bits a,b,c,d,e,i,f,g,h,j of 10-bit transmission code-groups. 8B/10B code-group bit assignments are illustrated in Figure 36-3. Each valid codegroup has been given a name using the following convention: /Dx.y/ for the 256 valid data code-groups, and /Kx.y/ for special control code-groups, where x is the decimal value of bits EDCBA, and y is the decimal value of bits HGF.”

802.3-1998, p927-928.

“36.2.5.2.2 Receive The PCS shall implement its Receive process as depicted in Figures 36-7a and 36-7b, including compliance with the associated state variables as specified in 36.2.5.1. The PCS

Exhibit C-11


Receive process continuously passes RXD<7:0> and sets the RX_DV and RX_ER signals to the GMII based on the received code-group from the PMA.



NOTE—Outgoing arcs leading to labeled polygons flow offpage to corresponding incoming arcs leading from labeled circles on Figure 36-7b, and vice versa.

Figure 36-7a—PCS receive state diagram, part a

Exhibit C-11



NOTE—Outgoing arcs leading to labeled polygons flow offpage to corresponding incoming arcs leading from labeled circles on Figure 36-7a, and vice versa.

Figure 36-7b—PCS receive state diagram, part b

When the Auto-Negotiation process xmit flag indicates CONFIGURATION or IDLE, the PCS Receive process continuously passes /C/ and /I/ ordered sets and rx_Config_Reg<D15:D0> to the Auto-Negotiation process.

Exhibit C-11

	36.2.5.2.3 State variable function carrier_detect(x) The detection of carrier on the underlying channel is used both by the MAC (via the GMII CRS signal and the Reconciliation sublayer) for deferral purposes, and by the PCS Transmit process for collision detection. A carrier event, signaled by the assertion of receiving, is indicated by the detection of a difference between the received code-group and /K28.5/ as specified in 36.2.5.1.4. A carrier event is in error if it does not start with an SPD. The PCS Receive process performs this function by continuously monitoring incoming code-groups for specific patterns that indicate non-/I/ activity such as SPD. The detection of an SPD carrier event causes the PCS to substitute the value (01010101) for the SPD, set RXD<7:0> to this value, and assert RX_DV. The pattern substituted for the SPD is consistent with the preamble pattern expected by the MAC. The detection of a non-SPD carrier event (false carrier) causes the PCS to substitute the value (00001110) for the code-group received, set RXD<7:0> to this value, and assert RX_ER. 36.2.5.2.4 Code-group stream decoding Subsequent to the detection of an SPD carrier event, the PCS Receive process performs the DECODE function on the incoming code-groups, passing decoded data to the GMII, including those corresponding to the remainder of the MAC preamble and SFD. The GMII signal RX_ER is asserted upon decoding any codegroup following the SPD that neither is a valid /D/ code-group nor follows the EPD rules in 36.2.4.14.1. Packets are terminated with an EPD as specified in 36.2.4.14. The PCS Receive process performs the check_end function to preserve the ability of the MAC to properly delimit the FCS at the end of a packet. Detection of /T/R/R/ or /T/R/K28.5/ by the check_end function denotes normal (i.e. non-error) packet termination. Detection of /R/R/R/ by the check_end function denotes packet termination with error and Carrier_Extend processing. Detection of /K28.5/D/K28.5/ by the check_end function denotes packet termination with error. Detection of /K28.5/(D21.5 or D2.2)/D0.0 by the check_end function denotes packet termination with error.” 802.3-1998, p947-952. To the extent that Plaintiff contends that 802.3-1998 does not explicitly or inherently disclose this element, or render obvious this element, this element is obvious in combination with one or more references and/or with the knowledge of a person of ordinary skill in the art, as set forth in Defendant’s Invalidity Contentions.
--	---

Exhibit C-11

11(a)	receiving a synchronization symbol by a receiver indicating a packet type, each packet type having a different synchronization symbol;	802.3-1998 discloses and/or renders obvious this limitation. <i>See, e.g.:</i> “24.2.2 Functional requirements The PCS comprises the Transmit, Receive, and Carrier Sense functions for 100BASE-T. In addition, the collisionDetect signal required by the MAC (COL on the MII) is derived from the PMA code-bit stream. The PCS shields the Reconciliation sublayer (and MAC) from the specific nature of the underlying channel. Specifically for receiving, the 100BASE-X PCS passes to the MII a sequence of data nibbles derived from incoming code-groups, each comprised of five code-bits, received from the medium. Code-group alignment and MAC packet delimiting is performed by embedding special non-data code-groups. The MII uses a nibble- wide, synchronous data path, with packet delimiting being provided by separate TX_EN and RX_DV signals. The PCS provides the functions necessary to map these two views of the exchanged data. The process is reversed for transmit. The following provides a detailed specification of the functions performed by the PCS, which comprise five parallel processes (Transmit, Transmit Bits, Receive, Receive Bits, and Carrier Sense). Figure 24-4 includes a functional block diagram of the PCS. The Receive Bits process accepts continuous code-bits via the PMA_UNITDATA.indicate primitive. Receive monitors these bits and generates RXD <3:0>, RX_DV and RX_ER on the MII, and the internal flag, receiving, used by the Carrier Sense and Transmit processes. The Transmit process generates continuous code-groups based upon the TXD <3:0>, TX_EN, and TX_ER signals on the MII. These code-groups are transmitted by Transmit Bits via the PMA_UNITDATA.request primitive. The Transmit process generates the MII signal COL based on whether a reception is occurring simultaneously with transmission. Additionally, it generates the internal flag, transmitting, for use by the Carrier Sense process. The Carrier Sense process asserts the MII signal CRS when either transmitting or receiving is TRUE. Both the Transmit and Receive processes monitor link_status via the PMA_LINK.indicate primitive, to account for potential link failure conditions. 24.2.2.1 Code-groups The PCS maps four-bit nibbles from the MII into five-bit code-groups, and vice versa, using a 4B/5B block coding scheme. A code-group is a consecutive sequence of five code-bits interpreted and mapped by the PCS. Implicit in the definition of a code-group is an establishment of code-group boundaries by an alignment function within the PCS Receive process. It is important to note that, with the sole exception of the SSD, which is used to achieve alignment,
-------	--	--

Exhibit C-11

		code-groups are undetectable and have no meaning outside the 100BASE-X physical protocol data unit, called a “stream.” The coding method used, derived from ISO/IEC 9314-1, provides a) Adequate codes (32) to provide for all Data code-groups (16) plus necessary control code-groups; b) Appropriate coding efficiency (4 data bits per 5 code-bits; 80%) to effect a 100 Mb/s Physical Layer interface on a 125 Mb/s physical channel as provided by FDDI PMDs; and c) Sufficient transition density to facilitate clock recovery (when not scrambled). Table 24-1 specifies the interpretation assigned to each five bit code-group, including the mapping to the nibble- wide (TXD or RXD) Data signals on the MII. The 32 code-groups are divided into four categories, as shown. For clarity in the remainder of this clause, code-group names are shown between slashes/. Code-group sequences are shown in succession, e.g.: /1/2/.... The indicated code-group mapping is identical to ISO/IEC 9314-1: 1989, with four exceptions: a) The FDDI term symbol is avoided in order to prevent confusion with other 100BASE-T terminology. In general, the term code-group is used in its place. b) The /S/ and /Q/ code-groups are not used by 100BASE-X and are interpreted as INVALID. c) The /R/ code-group is used in 100BASE-X as the second code-group of the End-of-Stream delimiter rather than to indicate a Reset condition. d) The /H/ code-group is used to propagate receive errors rather than to indicate the Halt Line State. 24.2.2.1.1 Data code-groups A Data code-group conveys one nibble of arbitrary data between the MII and the PCS. The sequence of Data code-groups is arbitrary, where any Data code-group can be followed by any other Data code-group. Data code-groups are coded and decoded but not interpreted by the PCS. Successful decoding of Data code-groups depends on proper receipt of the Start-of-Stream delimiter sequence, as defined in Table 24-1. 24.2.2.1.2 Idle code-groups The Idle code-group (/I/) is transferred between streams. It provides a continuous fill pattern to establish and maintain clock synchronization. Idle code-groups are emitted from, and interpreted by, the PCS. 24.2.2.1.3 Control code-groups The Control code-groups are used in pairs (/J/K/, /T/R/) to delimit MAC packets. Control code-groups are emitted from, and interpreted by, the PCS. 24.2.2.1.4 Start-of-Stream Delimiter (/J/K/) A Start-of-Stream Delimiter (SSD) is used to delineate the boundary of a data transmission sequence and to authenticate carrier events. The SSD is unique in that it may be recognized
--	--	--

Exhibit C-11

		independently of previously established code-group boundaries. The Receive function within the PCS uses the SSD to establish code-group boundaries. A SSD consists of the sequence /J/K/. On transmission, the first 8 bits of the MAC preamble are replaced by the SSD, a replacement that is reversed on reception. 24.2.2.1.5 End-of-Stream delimiter (/T/R/) An End-of-Stream delimiter (ESD) terminates all normal data transmissions. Unlike the SSD, an ESD cannot be recognized independent of previously established code-group boundaries. An ESD consists of the sequence /T/R/.
--	--	--

Exhibit C-11

Table 24-1—4B/5B code-groups

	PCS code-group [4:0] 4 3 2 1 0	Name	MH (TXD/RXD) <3:0> 3 2 1 0	Interpretation
D A T A	1 1 1 1 0	0	0 0 0 0	Data 0
	0 1 0 0 1	1	0 0 0 1	Data 1
	1 0 1 0 0	2	0 0 1 0	Data 2
	1 0 1 0 1	3	0 0 1 1	Data 3
	0 1 0 1 0	4	0 1 0 0	Data 4
	0 1 0 1 1	5	0 1 0 1	Data 5
	0 1 1 1 0	6	0 1 1 0	Data 6
	0 1 1 1 1	7	0 1 1 1	Data 7
	1 0 0 1 0	8	1 0 0 0	Data 8
	1 0 0 1 1	9	1 0 0 1	Data 9
	1 0 1 1 0	A	1 0 1 0	Data A
	1 0 1 1 1	B	1 0 1 1	Data B
	1 1 0 1 0	C	1 1 0 0	Data C
	1 1 0 1 1	D	1 1 0 1	Data D
	1 1 1 0 0	E	1 1 1 0	Data E
	1 1 1 0 1	F	1 1 1 1	Data F
C O N T R O L	1 1 1 1 1	I	undefined	IDLE; used as inter-stream fill code
	1 1 0 0 0	J	0 1 0 1	Start-of-Stream Delimiter, Part 1 of 2; always used in pairs with K
	1 0 0 0 1	K	0 1 0 1	Start-of-Stream Delimiter, Part 2 of 2; always used in pairs with J
	0 1 1 0 1	T	undefined	End-of-Stream Delimiter, Part 1 of 2; always used in pairs with R
I N V A L I D	0 0 1 1 1	R	undefined	End-of-Stream Delimiter, Part 2 of 2; always used in pairs with T
	0 0 1 0 0	H	Undefined	Transmit Error; used to force signaling errors
	0 0 0 0 0	V	Undefined	Invalid code
	0 0 0 0 1	V	Undefined	Invalid code
	0 0 0 1 0	V	Undefined	Invalid code
	0 0 0 1 1	V	Undefined	Invalid code
	0 0 1 0 1	V	Undefined	Invalid code
	0 0 1 1 0	V	Undefined	Invalid code
	0 1 0 0 0	V	Undefined	Invalid code
	0 1 1 0 0	V	Undefined	Invalid code
	1 0 0 0 0	V	Undefined	Invalid code
	1 1 0 0 1	V	Undefined	Invalid code

24.2.2.1.6 Invalid code-groups

Exhibit C-11

The /H/ code-group indicates that the PCS's client wishes to indicate a Transmit Error to its peer entity. The normal use of this indicator is for repeaters to propagate received errors. Transmit Error code-groups are emitted from the PCS, at the request of the PCS's client through the use of the TX_ER signal, as described in 24.2.4.2. The presence of any invalid code-group on the medium, including /H/, denotes a collision artifact or an error condition. Invalid code-groups are not intentionally transmitted onto the medium by DTE's. The PCS indicates the reception of an Invalid code-group on the MII through the use of the RX_ER signal, as described in 24.2.4.4.”

802.3-1998, p616-619.

“24.2.2.4 Mapping between MII and PMA

Figure 24-6 depicts the mapping of the nibble-wide data path of the MII to the five-bit-wide code-groups (internal to the PCS) and the code-bit path of the PMA interface.




Figure 24-6—PCS reference diagram

802.3-1998, p620.

“24.2.4.3 Receive Bits The Receive Bits process collects code-bits from the PMA interface passing them to the Receive process via rx_bits. rx_bits [9:0] represents a sliding, 10-bit window on the PMA code-bits, with newly received code-bits from the PMA (rx_code-bit) being shifted into rx_bits [0]. This is depicted in Figure 24-9. Bits are collected serially until Receive indicates alignment by asserting RX_DV, after which Receive Bits signals Receive for every five code-bits accumulated. Serial processing resumes with the de-assertion of RX_DV. The PCS shall implement the Receive Bits process as depicted in Figure 24-10 including compliance with the associated state variables as specified in 24.2.3.

Exhibit C-11

	24.2.4.4 Receive The Receive process state machine can be viewed as comprising two sections: prealigned and aligned. In the prealigned states, IDLE, CARRIER DETECT, and CONFIRM K, the Receive process is waiting for an indication of channel activity followed by a SSD. After successful alignment, the incoming code-groups are decoded while waiting for stream termination.” 802.3-1998, p624. “24.3.1.1 PMA_TYPE.indicate This primitive is generated by the PMA to indicate the nature of the PMA instantiation. The purpose of this primitive is to allow clients to support connections to the various types of 100BASE-T PMA entities in a generalized manner. 24.3.1.1.1 Semantics of the service primitive PMA_TYPE.indicate (pma_type) The pma_type parameter for use with a 100BASE-X PMA is “X”. 802.3-1998, p629. “36.2.2 Functions within the PCS The PCS comprises the PCS Transmit, Carrier Sense, Synchronization, PCS Receive, and Auto-Negotiation processes for 1000BASE-X. The PCS shields the Reconciliation sublayer (and MAC) from the specific nature of the underlying channel. When communicating with the GMII, the PCS uses an octet-wide, synchronous data path, with packet delimiting being provided by separate transmit control signals (TX_EN and TX_ER) and receive control signals (RX_DV and RX_ER). When communicating with the PMA, the PCS uses a ten-bit wide, synchronous data path, which conveys ten-bit code-groups. At the PMA Service Interface, code-group alignment and MAC packet delimiting are made possible by embedding special non-data code-groups in the transmitted code-group stream. The PCS provides the functions necessary to map packets between the GMII format and the PMA Service Interface format. The PCS Transmit process continuously generates code-groups based upon the TXD <7:0>, TX_EN, and TX_ER signals on the GMII, sending them immediately to the PMA Service Interface via the PMA_UNITDATA.request primitive. The PCS Transmit process generates the GMII signal COL based on whether a reception is occurring simultaneously with transmission. Additionally, it generates the internal flag, transmitting, for use by the Carrier Sense process. The PCS Transmit process monitors the Auto-Negotiation process xmit flag to determine whether to transmit data or reconfigure the link. The Carrier Sense process controls the GMII signal CRS (see Figure 36-8). The PCS Synchronization process continuously
--	---

Exhibit C-11

	accepts code-groups via the PMA_UNITDATA.indicate primitive and conveys received code-groups to the PCS Receive process via the SYNC_UNITDATA.indicate primitive. The PCS Synchronization process sets the sync_status flag to indicate whether the PMA is functioning dependably (as well as can be determined without exhaustive error-rate analysis). The PCS Receive process continuously accepts code-groups via the SYNC_UNITDATA.indicate primitive. The PCS Receive process monitors these code-groups and generates RXD <7:0>, RX_DV, and RX_ER on the GMII, and the internal flag, receiving, used by the Carrier Sense and Transmit processes. The PCS Auto-Negotiation process sets the xmit flag to inform the PCS Transmit process to either transmit normal idles interspersed with packets as requested by the GMII or to reconfigure the link. The PCS Auto- Negotiation process is specified in Clause 37. 36.2.3 Use of code-groups The PCS maps GMII signals into ten-bit code groups, and vice versa, using an 8B/10B block coding scheme. Implicit in the definition of a code-group is an establishment of code-group boundaries by a PMA code-group alignment function as specified in 36.3.2.4. Code-groups are unobservable and have no meaning outside the PCS. The PCS functions ENCODE and DECODE generate, manipulate, and interpret codegroups as provided by the rules in 36.2.4. 36.2.4 8B/10B transmission code The PCS uses a transmission code to improve the transmission characteristics of information to be transferred across the link. The encodings defined by the transmission code ensure that sufficient transitions are present in the PHY bit stream to make clock recovery possible at the receiver. Such encoding also greatly increases the likelihood of detecting any single or multiple bit errors that may occur during transmission and reception of information. In addition, some of the special code-groups of the transmission code contain a distinct and easily recognizable bit pattern that assists a receiver in achieving code-group alignment on the incoming PHY bit stream. The 8B/10B transmission code specified for use in this standard has a high transition density, is a run-length-limited code, and is dc-balanced. The transition density of the 8B/10B symbols ranges from 3 to 8 transitions per symbol. The definition of the 8B/10B transmission code in this standard is identical to that specified in ANSI X3.230-1994 (FC-PH), Clause 11. The relationship of code-group bit positions to PMA and other PCS constructs is illustrated in Figure 36-3.
--	--

Exhibit C-11




Figure 36-3—PCS reference diagram

36.2.4.1 Notation conventions

8B/10B transmission code uses letter notation for describing the bits of an unencoded information octet and a single control variable. Each bit of the unencoded information octet contains either a binary zero or a binary one. A control variable, Z, has either the value D or the value K. When the control variable associated with an unencoded information octet contains the value D, the associated encoded code-group is referred to as a data code-group. When the control variable associated with an unencoded information octet contains the value K, the associated encoded code-group is referred to as a special code-group. The bit notation of A,B,C,D,E,F,G,H for an unencoded information octet is used in the description of the 8B/10B transmission code. The bits A,B,C,D,E,F,G,H are translated to bits a,b,c,d,e,i,f,g,h,j of 10-bit transmission code-groups. 8B/10B code-group bit assignments are illustrated in Figure 36-3. Each valid codegroup has been given a name using the following convention: /Dx.y/ for the 256 valid data code-groups, and /Kx.y/ for special control code-groups, where x is the decimal value of bits EDCBA, and y is the decimal value of bits HGF.

36.2.4.2 Transmission order

Code-group bit transmission order is illustrated in Figure 36-3. Code-groups within multi-code-group ordered_sets (as specified in Table 36-3) are transmitted sequentially beginning with the special code-group used to distinguish the ordered_set (e.g., /K28.5/) and proceeding code-group

Exhibit C-11

by code-group from left to right within the definition of the ordered_set until all code-groups of the ordered_set are transmitted.

The first code-group of every multi-code-group ordered_set is transmitted in an even-numbered code-group position counting from the first code-group after a reset or power-on. Subsequent code-groups continuously alternate as odd and even-numbered code-groups. The contents of a packet are transmitted sequentially beginning with the ordered_set used to denote the Start_of_Packet (the SPD delimiter) and proceeding code-group by code-group from left to right within the definition of the packet until the ordered_set used to denote the End_of_Packet (the EPD delimiter) is transmitted.

36.2.4.3 Valid and invalid code-groups

Table 36-1a defines the valid data code-groups (D code-groups) of the 8B/10B transmission code. Table 36- 2 defines the valid special code-groups (K code-groups) of the code. The tables are used for both generating valid code-groups (encoding) and checking the validity of received code-groups (decoding). In the tables, each octet entry has two columns that represent two (not necessarily different) code-groups. The two columns correspond to the valid code-group based on the current value of the running disparity (Current RD – or Current RD +). Running disparity is a binary parameter with either the value negative (–) or the value positive (+). Annex 36B provides several 8B/10B transmission code running disparity calculation examples.”

802.3-1998, p927-929.

Table 36-1a—Valid data code-groups

Code Group Name	Octet Value	Octet Bits HGFEDCBA	Current RD –	Current RD +
			abcdei fghj	abcdei fghj
D0.0	00	000 00000	100111 0100	011000 1011
D1.0	01	000 00001	011101 0100	100010 1011
D2.0	02	000 00010	101101 0100	010010 1011
D3.0	03	000 00011	110001 1011	110001 0100
D4.0	04	000 00100	110101 0100	001010 1011
D5.0	05	000 00101	101001 1011	101001 0100
D6.0	06	000 00110	011001 1011	011001 0100
D7.0	07	000 00111	111000 1011	000111 0100
D8.0	08	000 01000	111001 0100	000110 1011
D9.0	09	000 01001	100101 1011	100101 0100
D10.0	0A	000 01010	010101 1011	010101 0100
D11.0	0B	000 01011	110100 1011	110100 0100
D12.0	0C	000 01100	001101 1011	001101 0100
D13.0	0D	000 01101	101100 1011	101100 0100
D14.0	0E	000 01110	011100 1011	011100 0100
D15.0	0F	000 01111	010111 0100	101000 1011
D16.0	10	000 10000	011011 0100	100100 1011

“

Exhibit C-11

...

Table 36-1e—Valid data code-groups

Code Group Name	Octet Value	Octet Bits HGF EDCBA	Current RD –	Current RD +
			abedei fghj	abedei fghj
D16.7	F0	111 10000	011011 0001	100100 1110
D17.7	F1	111 10001	100011 0111	100011 0001
D18.7	F2	111 10010	010011 0111	010011 0001
D19.7	F3	111 10011	110010 1110	110010 0001
D20.7	F4	111 10100	001011 0111	001011 0001
D21.7	F5	111 10101	101010 1110	101010 0001
D22.7	F6	111 10110	011010 1110	011010 0001
D23.7	F7	111 10111	111010 0001	000101 1110
D24.7	F8	111 11000	110011 0001	001100 1110
D25.7	F9	111 11001	100110 1110	100110 0001
D26.7	FA	111 11010	010110 1110	010110 0001
D27.7	FB	111 11011	110110 0001	001001 1110
D28.7	FC	111 11100	001110 1110	001110 0001
D29.7	FD	111 11101	101110 0001	010001 1110
D30.7	FE	111 11110	011110 0001	100001 1110
D31.7	FF	111 11111	101011 0001	010100 1110
(concluded)				

Table 36-2—Valid special code-groups

Code Group Name	Octet Value	Octet Bits HGF EDCBA	Current RD –	Current RD +	Notes
			abedei fghj	abedei fghj	
K28.0	1C	000 11100	001111 0100	110000 1011	1
K28.1	3C	001 11100	001111 1001	110000 0110	1,2
K28.2	5C	010 11100	001111 0101	110000 1010	1
K28.3	7C	011 11100	001111 0011	110000 1100	1
K28.4	9C	100 11100	001111 0010	110000 1101	1
K28.5	BC	101 11100	001111 1010	110000 0101	2
K28.6	DC	110 11100	001111 0110	110000 1001	1
K28.7	FC	111 11100	001111 1000	110000 0111	1,2
K23.7	F7	111 10111	111010 1000	000101 0111	
K27.7	FB	111 11011	110110 1000	001001 0111	
K29.7	FD	111 11101	101110 1000	010001 0111	
K30.7	FE	111 11110	011110 1000	100001 0111	
NOTE 1—Reserved.					
NOTE 2—Contains a comma.					

36.2.4.7 Ordered_sets Eight ordered_sets, consisting of a single special code-group or combinations of special and data codegroups are specifically defined. Ordered_sets which include /K28.5/ provide the ability to obtain bit and code-group synchronization and establish ordered_set

Exhibit C-11

		alignment (see 36.2.4.9 and 36.3.2.4). Ordered_sets provide for the delineation of a packet and synchronization between the transmitter and receiver circuits at opposite ends of a link. Table 36-3 lists the defined ordered_sets. Table 36-3—Defined ordered_sets <table><tr><th>Code</th><th>Ordered_Set</th><th>Number of Code-Groups</th><th>Encoding</th></tr><tr><td>/C/</td><td>Configuration</td><td></td><td>Alternating /C1/ and /C2/</td></tr><tr><td>/C1/</td><td>Configuration 1</td><td>4</td><td>/K28.5/D21.5/Config_Reg^a</td></tr><tr><td>/C2/</td><td>Configuration 2</td><td>4</td><td>/K28.5/D2.2/Config_Reg^a</td></tr><tr><td>/I/</td><td>IDLE</td><td></td><td>Correcting /I1/, Preserving /I2/</td></tr><tr><td>/I1/</td><td>IDLE 1</td><td>2</td><td>/K28.5/D5.6/</td></tr><tr><td>/I2/</td><td>IDLE 2</td><td>2</td><td>/K28.5/D16.2/</td></tr><tr><td></td><td>Encapsulation</td><td></td><td></td></tr><tr><td>/R/</td><td>Carrier_Extend</td><td>1</td><td>/K23.7/</td></tr><tr><td>/S/</td><td>Start_of_Packet</td><td>1</td><td>/K27.7/</td></tr><tr><td>/T/</td><td>End_of_Packet</td><td>1</td><td>/K29.7/</td></tr><tr><td>/V/</td><td>Error_Propagation</td><td>1</td><td>/K30.7/</td></tr></table> ^aTwo data code-groups representing the Config_Reg value. 802.3-1998, p931-936. To the extent that Plaintiff contends that 802.3-1998 does not explicitly or inherently disclose this element, or render obvious this element, this element is obvious in combination with one or more references and/or with the knowledge of a person of ordinary skill in the art, as set forth in Defendant’s Invalidity Contentions.	Code	Ordered_Set	Number of Code-Groups	Encoding	/C/	Configuration		Alternating /C1/ and /C2/	/C1/	Configuration 1	4	/K28.5/D21.5/Config_Reg ^a	/C2/	Configuration 2	4	/K28.5/D2.2/Config_Reg ^a	/I/	IDLE		Correcting /I1/, Preserving /I2/	/I1/	IDLE 1	2	/K28.5/D5.6/	/I2/	IDLE 2	2	/K28.5/D16.2/		Encapsulation			/R/	Carrier_Extend	1	/K23.7/	/S/	Start_of_Packet	1	/K27.7/	/T/	End_of_Packet	1	/K29.7/	/V/	Error_Propagation	1	/K30.7/
Code	Ordered_Set	Number of Code-Groups	Encoding																																															
/C/	Configuration		Alternating /C1/ and /C2/																																															
/C1/	Configuration 1	4	/K28.5/D21.5/Config_Reg ^a																																															
/C2/	Configuration 2	4	/K28.5/D2.2/Config_Reg ^a																																															
/I/	IDLE		Correcting /I1/, Preserving /I2/																																															
/I1/	IDLE 1	2	/K28.5/D5.6/																																															
/I2/	IDLE 2	2	/K28.5/D16.2/																																															
	Encapsulation																																																	
/R/	Carrier_Extend	1	/K23.7/																																															
/S/	Start_of_Packet	1	/K27.7/																																															
/T/	End_of_Packet	1	/K29.7/																																															
/V/	Error_Propagation	1	/K30.7/																																															
11(b)	receiving a packet of symbols;	802.3-1998 discloses and/or renders obvious this limitation. <i>See, e.g.:</i> “24.2.2.2 Encapsulation The 100BASE-X PCS accepts frames from the MAC through the Reconciliation sublayer and MII. Due to the continuously signaled nature of the underlying PMA, and the encoding performed by the PCS, the 100BASE-X PCS encapsulates the MAC frame (100BASE-X Service Data Unit, SDU) into a Physical Layer stream (100BASE-X Protocol Data Unit, PDU).																																																

Exhibit C-11

Except for the two code-group SSD, data nibbles within the SDU (including the non-SSD portions of the MAC preamble and SFD) are not interpreted by the 100BASE-X PHY. The conversion from a MAC frame to a Physical Layer stream and back to a MAC frame is transparent to the MAC. Figure 24-5 depicts the mapping between MAC frames and Physical Layer streams.




Figure 24-5—PCS encapsulation

802.3-1998, p619.

“24.2.2.4 Mapping between MII and PMA

Figure 24-6 depicts the mapping of the nibble-wide data path of the MII to the five-bit-wide code-groups (internal to the PCS) and the code-bit path of the PMA interface.




Figure 24-6—PCS reference diagram

802.3-1998, p620.

Exhibit C-11

		“24.2.4.3 Receive Bits The Receive Bits process collects code-bits from the PMA interface passing them to the Receive process via rx_bits. rx_bits [9:0] represents a sliding, 10-bit window on the PMA code-bits, with newly received code-bits from the PMA (rx_code-bit) being shifted into rx_bits [0]. This is depicted in Figure 24-9. Bits are collected serially until Receive indicates alignment by asserting RX_DV, after which Receive Bits signals Receive for every five code-bits accumulated. Serial processing resumes with the de-assertion of RX_DV. The PCS shall implement the Receive Bits process as depicted in Figure 24-10 including compliance with the associated state variables as specified in 24.2.3. 24.2.4.4 Receive The Receive process state machine can be viewed as comprising two sections: prealigned and aligned. In the prealigned states, IDLE, CARRIER DETECT, and CONFIRM K, the Receive process is waiting for an indication of channel activity followed by a SSD. After successful alignment, the incoming code-groups are decoded while waiting for stream termination.” 802.3-1998, p624.
--	--	---

Exhibit C-11





Figure 24-11—Receive state diagram

802.3-1998, p628.

“36.2.5.2.2 Receive The PCS shall implement its Receive process as depicted in Figures 36-7a and 36-7b, including compliance with the associated state variables as specified in 36.2.5.1. The PCS

Exhibit C-11


Receive process continuously passes RXD<7:0> and sets the RX_DV and RX_ER signals to the GMII based on the received code-group from the PMA.



NOTE—Outgoing arcs leading to labeled polygons flow offpage to corresponding incoming arcs leading from labeled circles on Figure 36-7b, and vice versa.

Figure 36-7a—PCS receive state diagram, part a

Exhibit C-11



NOTE—Outgoing arcs leading to labeled polygons flow offpage to corresponding incoming arcs leading from labeled circles on Figure 36-7a, and vice versa.

Figure 36-7b—PCS receive state diagram, part b

When the Auto-Negotiation process xmit flag indicates CONFIGURATION or IDLE, the PCS Receive process continuously passes /C/ and /I/ ordered sets and rx_Config_Reg<D15:D0> to the Auto-Negotiation process.”

Exhibit C-11

		802.3-1998, p947-952. To the extent that Plaintiff contends that 802.3-1998 does not explicitly or inherently disclose this element, or render obvious this element, this element is obvious in combination with one or more references and/or with the knowledge of a person of ordinary skill in the art, as set forth in Defendant's Invalidity Contentions.
11(c)	indicating that the received packet of symbols has the packet type of the received synchronization symbol; and	802.3-1998 discloses and/or renders obvious this limitation. <i>See, e.g.:</i> “24.2.2 Functional requirements The PCS comprises the Transmit, Receive, and Carrier Sense functions for 100BASE-T. In addition, the collisionDetect signal required by the MAC (COL on the MII) is derived from the PMA code-bit stream. The PCS shields the Reconciliation sublayer (and MAC) from the specific nature of the underlying channel. Specifically for receiving, the 100BASE-X PCS passes to the MII a sequence of data nibbles derived from incoming code-groups, each comprised of five code-bits, received from the medium. Code-group alignment and MAC packet delimiting is performed by embedding special non-data code-groups. The MII uses a nibble- wide, synchronous data path, with packet delimiting being provided by separate TX_EN and RX_DV signals. The PCS provides the functions necessary to map these two views of the exchanged data. The process is reversed for transmit. The following provides a detailed specification of the functions performed by the PCS, which comprise five parallel processes (Transmit, Transmit Bits, Receive, Receive Bits, and Carrier Sense). Figure 24-4 includes a functional block diagram of the PCS. The Receive Bits process accepts continuous code-bits via the PMA_UNITDATA.indicate primitive. Receive monitors these bits and generates RXD <3:0>, RX_DV and RX_ER on the MII, and the internal flag, receiving, used by the Carrier Sense and Transmit processes. The Transmit process generates continuous code-groups based upon the TXD <3:0>, TX_EN, and TX_ER signals on the MII. These code-groups are transmitted by Transmit Bits via the PMA_UNITDATA.request primitive. The Transmit process generates the MII signal COL based on whether a reception is occurring simultaneously with transmission. Additionally, it generates the internal flag, transmitting, for use by the Carrier Sense process.

Exhibit C-11

		The Carrier Sense process asserts the MII signal CRS when either transmitting or receiving is TRUE. Both the Transmit and Receive processes monitor link_status via the PMA_LINK.indicate primitive, to account for potential link failure conditions. 24.2.2.1 Code-groups The PCS maps four-bit nibbles from the MII into five-bit code-groups, and vice versa, using a 4B/5B block coding scheme. A code-group is a consecutive sequence of five code-bits interpreted and mapped by the PCS. Implicit in the definition of a code-group is an establishment of code-group boundaries by an alignment function within the PCS Receive process. It is important to note that, with the sole exception of the SSD, which is used to achieve alignment, code-groups are undetectable and have no meaning outside the 100BASE-X physical protocol data unit, called a “stream.” The coding method used, derived from ISO/IEC 9314-1, provides a) Adequate codes (32) to provide for all Data code-groups (16) plus necessary control code-groups; b) Appropriate coding efficiency (4 data bits per 5 code-bits; 80%) to effect a 100 Mb/s Physical Layer interface on a 125 Mb/s physical channel as provided by FDDI PMDs; and c) Sufficient transition density to facilitate clock recovery (when not scrambled). Table 24-1 specifies the interpretation assigned to each five bit code-group, including the mapping to the nibble- wide (TXD or RXD) Data signals on the MII. The 32 code-groups are divided into four categories, as shown. For clarity in the remainder of this clause, code-group names are shown between /slashes/. Code-group sequences are shown in succession, e.g.: /1/2/.... The indicated code-group mapping is identical to ISO/IEC 9314-1: 1989, with four exceptions: a) The FDDI term symbol is avoided in order to prevent confusion with other 100BASE-T terminology. In general, the term code-group is used in its place. b) The /S/ and /Q/ code-groups are not used by 100BASE-X and are interpreted as INVALID. c) The /R/ code-group is used in 100BASE-X as the second code-group of the End-of-Stream delimiter rather than to indicate a Reset condition. d) The /H/ code-group is used to propagate receive errors rather than to indicate the Halt Line State. 24.2.2.1.1 Data code-groups A Data code-group conveys one nibble of arbitrary data between the MII and the PCS. The sequence of Data code-groups is arbitrary, where any Data code-group can be followed by any other Data code-group. Data code-groups are coded and decoded but not interpreted by the PCS. Successful decoding of Data code-groups depends on proper receipt of the Start-of-Stream delimiter sequence, as defined in Table 24-1. 24.2.2.1.2 Idle code-groups
--	--	---

Exhibit C-11

	The Idle code-group (/I/) is transferred between streams. It provides a continuous fill pattern to establish and maintain clock synchronization. Idle code-groups are emitted from, and interpreted by, the PCS. 24.2.2.1.3 Control code-groups The Control code-groups are used in pairs (/J/K/, /T/R/) to delimit MAC packets. Control code-groups are emitted from, and interpreted by, the PCS. 24.2.2.1.4 Start-of-Stream Delimiter (/J/K/) A Start-of-Stream Delimiter (SSD) is used to delineate the boundary of a data transmission sequence and to authenticate carrier events. The SSD is unique in that it may be recognized independently of previously established code-group boundaries. The Receive function within the PCS uses the SSD to establish code-group boundaries. A SSD consists of the sequence /J/K/. On transmission, the first 8 bits of the MAC preamble are replaced by the SSD, a replacement that is reversed on reception. 24.2.2.1.5 End-of-Stream delimiter (/T/R/) An End-of-Stream delimiter (ESD) terminates all normal data transmissions. Unlike the SSD, an ESD cannot be recognized independent of previously established code-group boundaries. An ESD consists of the sequence /T/R/.
--	--

Exhibit C-11

Table 24-1—4B/5B code-groups

	PCS code-group [4:0] 4 3 2 1 0	Name	MII (TXD/RXD) <3:0> 3 2 1 0	Interpretation
D A T A	1 1 1 1 0	0	0 0 0 0	Data 0
	0 1 0 0 1	1	0 0 0 1	Data 1
	1 0 1 0 0	2	0 0 1 0	Data 2
	1 0 1 0 1	3	0 0 1 1	Data 3
	0 1 0 1 0	4	0 1 0 0	Data 4
	0 1 0 1 1	5	0 1 0 1	Data 5
	0 1 1 1 0	6	0 1 1 0	Data 6
	0 1 1 1 1	7	0 1 1 1	Data 7
	1 0 0 1 0	8	1 0 0 0	Data 8
	1 0 0 1 1	9	1 0 0 1	Data 9
	1 0 1 1 0	A	1 0 1 0	Data A
	1 0 1 1 1	B	1 0 1 1	Data B
	1 1 0 1 0	C	1 1 0 0	Data C
	1 1 0 1 1	D	1 1 0 1	Data D
	1 1 1 0 0	E	1 1 1 0	Data E
	1 1 1 0 1	F	1 1 1 1	Data F
C O N T R O L	1 1 1 1 1	I	undefined	IDLE; used as inter-stream fill code
	1 1 0 0 0	J	0 1 0 1	Start-of-Stream Delimiter, Part 1 of 2; always used in pairs with K
	1 0 0 0 1	K	0 1 0 1	Start-of-Stream Delimiter, Part 2 of 2; always used in pairs with J
	0 1 1 0 1	T	undefined	End-of-Stream Delimiter, Part 1 of 2; always used in pairs with R
I N V A L I D	0 0 1 1 1	R	undefined	End-of-Stream Delimiter, Part 2 of 2; always used in pairs with T
	0 0 1 0 0	H	Undefined	Transmit Error; used to force signaling errors
	0 0 0 0 0	V	Undefined	Invalid code
	0 0 0 0 1	V	Undefined	Invalid code
	0 0 0 1 0	V	Undefined	Invalid code
	0 0 0 1 1	V	Undefined	Invalid code
	0 0 1 0 1	V	Undefined	Invalid code
	0 0 1 1 0	V	Undefined	Invalid code
	0 1 0 0 0	V	Undefined	Invalid code
	0 1 1 0 0	V	Undefined	Invalid code
	1 0 0 0 0	V	Undefined	Invalid code
	1 1 0 0 1	V	Undefined	Invalid code

24.2.2.1.6 Invalid code-groups

Exhibit C-11

		The /H/ code-group indicates that the PCS's client wishes to indicate a Transmit Error to its peer entity. The normal use of this indicator is for repeaters to propagate received errors. Transmit Error code-groups are emitted from the PCS, at the request of the PCS's client through the use of the TX_ER signal, as described in 24.2.4.2. The presence of any invalid code-group on the medium, including /H/, denotes a collision artifact or an error condition. Invalid code-groups are not intentionally transmitted onto the medium by DTE's. The PCS indicates the reception of an Invalid code-group on the MII through the use of the RX_ER signal, as described in 24.2.4.4.” 802.3-1998, p616-619.
--	--	--

Exhibit C-11




Figure 24-11—Receive state diagram


802.3-1998, p628.

“24.3.1.1 PMA_TYPE.indicate

Exhibit C-11

		This primitive is generated by the PMA to indicate the nature of the PMA instantiation. The purpose of this primitive is to allow clients to support connections to the various types of 100BASE-T PMA entities in a generalized manner. 24.3.1.1.1 Semantics of the service primitive PMA_TYPE.indicate (pma_type) The pma_type parameter for use with a 100BASE-X PMA is “X”. 802.3-1998, p629. “36.2.5.2.2 Receive The PCS shall implement its Receive process as depicted in Figures 36-7a and 36-7b, including compliance with the associated state variables as specified in 36.2.5.1. The PCS Receive process continuously passes RXD<7:0> and sets the RX_DV and RX_ER signals to the GMII based on the received code-group from the PMA.
--	--	--


Exhibit C-11



NOTE—Outgoing arcs leading to labeled polygons flow off-page to corresponding incoming arcs leading from labeled circles on Figure 36-7b, and vice versa.

Figure 36-7a—PCS receive state diagram, part a

Exhibit C-11



NOTE—Outgoing arcs leading to labeled polygons flow offpage to corresponding incoming arcs leading from labeled circles on Figure 36-7a, and vice versa.

Figure 36-7b—PCS receive state diagram, part b

When the Auto-Negotiation process xmit flag indicates CONFIGURATION or IDLE, the PCS Receive process continuously passes /C/ and /I/ ordered sets and rx_Config_Reg<D15:D0> to the Auto-Negotiation process.

Exhibit C-11

	36.2.5.2.3 State variable function carrier_detect(x) The detection of carrier on the underlying channel is used both by the MAC (via the GMII CRS signal and the Reconciliation sublayer) for deferral purposes, and by the PCS Transmit process for collision detection. A carrier event, signaled by the assertion of receiving, is indicated by the detection of a difference between the received code-group and /K28.5/ as specified in 36.2.5.1.4. A carrier event is in error if it does not start with an SPD. The PCS Receive process performs this function by continuously monitoring incoming code-groups for specific patterns that indicate non-/I/ activity such as SPD. The detection of an SPD carrier event causes the PCS to substitute the value (01010101) for the SPD, set RXD<7:0> to this value, and assert RX_DV. The pattern substituted for the SPD is consistent with the preamble pattern expected by the MAC. The detection of a non-SPD carrier event (false carrier) causes the PCS to substitute the value (00001110) for the code-group received, set RXD<7:0> to this value, and assert RX_ER. 36.2.5.2.4 Code-group stream decoding Subsequent to the detection of an SPD carrier event, the PCS Receive process performs the DECODE function on the incoming code-groups, passing decoded data to the GMII, including those corresponding to the remainder of the MAC preamble and SFD. The GMII signal RX_ER is asserted upon decoding any codegroup following the SPD that neither is a valid /D/ code-group nor follows the EPD rules in 36.2.4.14.1. Packets are terminated with an EPD as specified in 36.2.4.14. The PCS Receive process performs the check_end function to preserve the ability of the MAC to properly delimit the FCS at the end of a packet. Detection of /T/R/R/ or /T/R/K28.5/ by the check_end function denotes normal (i.e. non-error) packet termination. Detection of /R/R/R/ by the check_end function denotes packet termination with error and Carrier_Extend processing. Detection of /K28.5/D/K28.5/ by the check_end function denotes packet termination with error. Detection of /K28.5/(D21.5 or D2.2)/D0.0 by the check_end function denotes packet termination with error.” 802.3-1998, p947-952. To the extent that Plaintiff contends that 802.3-1998 does not explicitly or inherently disclose this element, or render obvious this element, this element is obvious in combination with one or more references and/or with the knowledge of a person of ordinary skill in the art, as set forth in Defendant’s Invalidity Contentions.
--	---

Exhibit C-11

11(d)	the synchronization symbol permitting properly alignment upon receiving with a synchronization primitive to be correctly aligned on a symbol boundary.	802.3-1998 discloses and/or renders obvious this limitation. <i>See, e.g.:</i> “24.2.4.4.2 Code-group alignment After channel activity is detected, the Receive process first aligns the incoming code-bits on code-group boundaries for subsequent data decoding. This is achieved by scanning the rx_bits vector for a SSD (/J/K/). The MII RX_DV signal remains deasserted during this time, which ensures that the Reconciliation sublayer will ignore any signals on RXD <3:0>. Detection of the SSD causes the Receive process to enter the START OF STREAM J state. Figure 24-9—Receive Bits reference diagram Well-formed streams contain SSD (/J/K/) in place of the first eight preamble bits. In the event that something else is sensed immediately following detection of carrier, a False Carrier Indication is signaled to the MII by asserting RX_ER and setting RXD to 1110 while RX_DV remains deasserted. The associated carrier event, as terminated by 10 ONEs, is otherwise ignored.” 802.3-1998, p625-626. “36.2.5.2.6 Synchronization The PCS shall implement the Synchronization process as depicted in Figure 36-9 including compliance with the associated state variables as specified in 36.2.5.1. The Synchronization process is responsible for determining whether the underlying receive channel is ready for operation. Failure of the underlying channel typically causes the PMA client to suspend normal actions. A receiver that is in the LOSS_OF_SYNC state and that has acquired bit synchronization attempts to acquire code-group synchronization via the Synchronization process. Code-group synchronization is acquired by the detection of three ordered_sets containing commas in their leftmost bit positions without intervening invalid code-group errors. Upon acquisition of code-group synchronization, the receiver enters the SYNC_ACQUIRED_1 state. Acquisition of synchronization ensures the alignment of multi-code-group ordered_sets to even-numbered code-group boundaries. Once synchronization is acquired, the Synchronization process tests received code-groups in sets of four code-groups and employs multiple sub-states, effecting hysteresis, to move between the
-------	---	--

Exhibit C-11

	SYNC_ACQUIRED_1 and LOSS_OF_SYNC states. The condition sync_status=FAIL existing for ten ms or more causes the PCS Auto-Negotiation process to begin and the PCS Transmit process to begin transmission of /C/. Upon reception of three matching /C/s from the link partner, the PCS Auto-Negotiation process begins. The internal signal receiving is de-asserted in the PCS Receive process LINK_FAILED state when sync_status=FAIL and a code-group is received.” 802.3-1998, p953. “36.3.2.4 Code-group alignment In the event the PMA sublayer detects a comma+ within the incoming rx_bit stream, it may realign its current code-group boundary, if necessary, to that of the received comma+ as shown in Figure 36-3. This process is referred to in this document as code-group alignment. The code-group alignment function shall be operational when the EN_CDET signal is active (see 36.3.3.1). During the code-group alignment process, the PMA sublayer may delete or modify up to four, but shall delete or modify no more than four, ten-bit code-groups in order to align the correct receive clock and code-group containing the comma+. This process is referred to as code-group slipping. In addition, the PMA sublayer is permitted to realign the current code-group boundary upon receipt of a comma- pattern.” 802.3-1998, p956. “Code-group alignment occurs in the PMA Receive function, if enabled by EN_CDET, when a comma pattern occurs in the PHY bit stream. Upon recognition of the comma pattern, the PMA Receive function outputs the ten-bit code-group containing the comma on rx_code-group<9:0> with the alignment specified in Figure 36-3, and clocked on the rising edge of PMA_RX_CLK<1>.” 802.3-1998, p957. To the extent that Plaintiff contends that 802.3-1998 does not explicitly or inherently disclose this element, or render obvious this element, this element is obvious in combination with one or more references and/or with the knowledge of a person of ordinary skill in the art, as set forth in Defendant’s Invalidity Contentions.
Claim 19	

Exhibit C-11

19[p]	A communications device for transmitting packet types of packets, comprising:	To the extent the preamble is limiting, 802.3-1998 discloses this element and/or renders it obvious either alone or in combination with other references and/or the knowledge of one of ordinary skill in the art. <i>See, e.g.:</i> “24.2.2 Functional requirements The PCS comprises the Transmit, Receive, and Carrier Sense functions for 100BASE-T. In addition, the collisionDetect signal required by the MAC (COL on the MII) is derived from the PMA code-bit stream. The PCS shields the Reconciliation sublayer (and MAC) from the specific nature of the underlying channel. Specifically for receiving, the 100BASE-X PCS passes to the MII a sequence of data nibbles derived from incoming code-groups, each comprised of five code-bits, received from the medium. Code-group alignment and MAC packet delimiting is performed by embedding special non-data code-groups. The MII uses a nibble- wide, synchronous data path, with packet delimiting being provided by separate TX_EN and RX_DV signals. The PCS provides the functions necessary to map these two views of the exchanged data. The process is reversed for transmit. The following provides a detailed specification of the functions performed by the PCS, which comprise five parallel processes (Transmit, Transmit Bits, Receive, Receive Bits, and Carrier Sense). Figure 24-4 includes a functional block diagram of the PCS.
-------	--	--

Exhibit C-11




Figure 24-4—Functional block diagram

802.3-1998, p615-616.

“24.3.1.1 PMA_TYPE.indicate

Exhibit C-11

This primitive is generated by the PMA to indicate the nature of the PMA instantiation. The purpose of this primitive is to allow clients to support connections to the various types of 100BASE-T PMA entities in a generalized manner.

24.3.1.1.1 Semantics of the service primitive

PMA_TYPE.indicate (pma_type)

The pma_type parameter for use with a 100BASE-X PMA is “X”.

802.3-1998, p629.

“24.2.2.4 Mapping between MII and PMA

Figure 24-6 depicts the mapping of the nibble-wide data path of the MII to the five-bit-wide code-groups (internal to the PCS) and the code-bit path of the PMA interface.




Figure 24-6—PCS reference diagram

802.3-1998, p620.

“36.2.2 Functions within the PCS

The PCS comprises the PCS Transmit, Carrier Sense, Synchronization, PCS Receive, and Auto-Negotiation processes for 1000BASE-X. The PCS shields the Reconciliation sublayer (and MAC) from the specific nature of the underlying channel. When communicating with the GMII, the PCS uses an octet-wide, synchronous data path, with packet delimiting being provided by separate transmit control signals (TX_EN and TX_ER) and receive control signals (RX_DV and RX_ER). When communicating with the PMA, the PCS uses a ten-bit wide, synchronous data path, which conveys ten-bit code-groups. At the PMA Service Interface, code-group alignment and MAC packet delimiting are made possible by embedding special non-data code-groups in the transmitted

Exhibit C-11

	code-group stream. The PCS provides the functions necessary to map packets between the GMII format and the PMA Service Interface format. The PCS Transmit process continuously generates code-groups based upon the TXD <7:0>, TX_EN, and TX_ER signals on the GMII, sending them immediately to the PMA Service Interface via the PMA_UNITDATA.request primitive. The PCS Transmit process generates the GMII signal COL based on whether a reception is occurring simultaneously with transmission. Additionally, it generates the internal flag, transmitting, for use by the Carrier Sense process. The PCS Transmit process monitors the Auto-Negotiation process xmit flag to determine whether to transmit data or reconfigure the link. The Carrier Sense process controls the GMII signal CRS (see Figure 36-8). The PCS Synchronization process continuously accepts code-groups via the PMA_UNITDATA.indicate primitive and conveys received code-groups to the PCS Receive process via the SYNC_UNITDATA.indicate primitive. The PCS Synchronization process sets the sync_status flag to indicate whether the PMA is functioning dependably (as well as can be determined without exhaustive error-rate analysis). The PCS Receive process continuously accepts code-groups via the SYNC_UNITDATA.indicate primitive. The PCS Receive process monitors these code-groups and generates RXD <7:0>, RX_DV, and RX_ER on the GMII, and the internal flag, receiving, used by the Carrier Sense and Transmit processes. The PCS Auto-Negotiation process sets the xmit flag to inform the PCS Transmit process to either transmit normal idles interspersed with packets as requested by the GMII or to reconfigure the link. The PCS Auto- Negotiation process is specified in Clause 37. 36.2.3 Use of code-groups The PCS maps GMII signals into ten-bit code groups, and vice versa, using an 8B/10B block coding scheme. Implicit in the definition of a code-group is an establishment of code-group boundaries by a PMA code-group alignment function as specified in 36.3.2.4. Code-groups are unobservable and have no meaning outside the PCS. The PCS functions ENCODE and DECODE generate, manipulate, and interpret codegroups as provided by the rules in 36.2.4. 36.2.4 8B/10B transmission code The PCS uses a transmission code to improve the transmission characteristics of information to be transferred across the link. The encodings defined by the transmission code ensure that sufficient transitions are present in the PHY bit stream to make clock recovery possible at the receiver. Such encoding also greatly increases the likelihood of detecting any single or multiple bit errors that may occur during transmission and reception of information. In addition, some of the special code-groups of the transmission code contain a distinct and easily recognizable bit pattern that assists a receiver in achieving code-group alignment on the incoming PHY bit stream. The 8B/10B
--	---

Exhibit C-11

transmission code specified for use in this standard has a high transition density, is a run-length-limited code, and is dc-balanced. The transition density of the 8B/10B symbols ranges from 3 to 8 transitions per symbol.

The definition of the 8B/10B transmission code in this standard is identical to that specified in ANSI X3.230-1994 (FC-PH), Clause 11. The relationship of code-group bit positions to PMA and other PCS constructs is illustrated in Figure 36-3.




Figure 36-3—PCS reference diagram

36.2.4.1 Notation conventions

8B/10B transmission code uses letter notation for describing the bits of an unencoded information octet and a single control variable. Each bit of the unencoded information octet contains either a binary zero or a binary one. A control variable, Z, has either the value D or the value K. When the control variable associated with an unencoded information octet contains the value D, the associated encoded code-group is referred to as a data code-group. When the control variable associated with an unencoded information octet contains the value K, the associated encoded code-group is referred to as a special code-group. The bit notation of A,B,C,D,E,F,G,H for an unencoded information octet is used in the description of the 8B/10B transmission code. The bits A,B,C,D,E,F,G,H are translated to bits a,b,c,d,e,f,g,h,j of 10-bit transmission code-groups. 8B/10B code-group bit assignments are illustrated in Figure 36-3. Each valid codegroup has been given a name using the following

Exhibit C-11

		convention: /Dx.y/ for the 256 valid data code-groups, and /Kx.y/ for special control code-groups, where x is the decimal value of bits EDCBA, and y is the decimal value of bits HGF.” 802.3-1998, p927-928. To the extent that Plaintiff contends that 802.3-1998 does not explicitly or inherently disclose this element, or render obvious this element, this element is obvious in combination with one or more references and/or with the knowledge of a person of ordinary skill in the art, as set forth in Defendant’s Invalidity Contentions.
19(a)	an identification component that identifies a packet type of a packet of symbols; and	802.3-1998 discloses and/or renders obvious this limitation. <i>See, e.g.:</i> “24.2.2 Functional requirements The PCS comprises the Transmit, Receive, and Carrier Sense functions for 100BASE-T. In addition, the collisionDetect signal required by the MAC (COL on the MII) is derived from the PMA code-bit stream. The PCS shields the Reconciliation sublayer (and MAC) from the specific nature of the underlying channel. Specifically for receiving, the 100BASE-X PCS passes to the MII a sequence of data nibbles derived from incoming code-groups, each comprised of five code-bits, received from the medium. Code-group alignment and MAC packet delimiting is performed by embedding special non-data code-groups. The MII uses a nibble- wide, synchronous data path, with packet delimiting being provided by separate TX_EN and RX_DV signals. The PCS provides the functions necessary to map these two views of the exchanged data. The process is reversed for transmit. The following provides a detailed specification of the functions performed by the PCS, which comprise five parallel processes (Transmit, Transmit Bits, Receive, Receive Bits, and Carrier Sense). Figure 24-4 includes a functional block diagram of the PCS. The Receive Bits process accepts continuous code-bits via the PMA_UNITDATA.indicate primitive. Receive monitors these bits and generates RXD <3:0>, RX_DV and RX_ER on the MII, and the internal flag, receiving, used by the Carrier Sense and Transmit processes. The Transmit process generates continuous code-groups based upon the TXD <3:0>, TX_EN, and TX_ER signals on the MII. These code-groups are transmitted by Transmit Bits via the PMA_UNITDATA.request primitive.

Exhibit C-11

		The Transmit process generates the MII signal COL based on whether a reception is occurring simultaneously with transmission. Additionally, it generates the internal flag, transmitting, for use by the Carrier Sense process. The Carrier Sense process asserts the MII signal CRS when either transmitting or receiving is TRUE. Both the Transmit and Receive processes monitor link_status via the PMA_LINK.indicate primitive, to account for potential link failure conditions. 24.2.2.1 Code-groups The PCS maps four-bit nibbles from the MII into five-bit code-groups, and vice versa, using a 4B/5B block coding scheme. A code-group is a consecutive sequence of five code-bits interpreted and mapped by the PCS. Implicit in the definition of a code-group is an establishment of code-group boundaries by an alignment function within the PCS Receive process. It is important to note that, with the sole exception of the SSD, which is used to achieve alignment, code-groups are undetectable and have no meaning outside the 100BASE-X physical protocol data unit, called a “stream.” The coding method used, derived from ISO/IEC 9314-1, provides a) Adequate codes (32) to provide for all Data code-groups (16) plus necessary control code-groups; b) Appropriate coding efficiency (4 data bits per 5 code-bits; 80%) to effect a 100 Mb/s Physical Layer interface on a 125 Mb/s physical channel as provided by FDDI PMDs; and c) Sufficient transition density to facilitate clock recovery (when not scrambled). Table 24-1 specifies the interpretation assigned to each five bit code-group, including the mapping to the nibble- wide (TXD or RXD) Data signals on the MII. The 32 code-groups are divided into four categories, as shown. For clarity in the remainder of this clause, code-group names are shown between /slashes/. Code-group sequences are shown in succession, e.g.: /1/2/.... The indicated code-group mapping is identical to ISO/IEC 9314-1: 1989, with four exceptions: a) The FDDI term symbol is avoided in order to prevent confusion with other 100BASE-T terminology. In general, the term code-group is used in its place. b) The /S/ and /Q/ code-groups are not used by 100BASE-X and are interpreted as INVALID. c) The /R/ code-group is used in 100BASE-X as the second code-group of the End-of-Stream delimiter rather than to indicate a Reset condition. d) The /H/ code-group is used to propagate receive errors rather than to indicate the Halt Line State. 24.2.2.1.1 Data code-groups A Data code-group conveys one nibble of arbitrary data between the MII and the PCS. The sequence of Data code-groups is arbitrary, where any Data code-group can be followed by any other Data code-group. Data code-groups are coded and decoded but not interpreted by the PCS. Successful
--	--	---

Exhibit C-11

	decoding of Data code-groups depends on proper receipt of the Start-of-Stream delimiter sequence, as defined in Table 24-1. 24.2.2.1.2 Idle code-groups The Idle code-group (/I/) is transferred between streams. It provides a continuous fill pattern to establish and maintain clock synchronization. Idle code-groups are emitted from, and interpreted by, the PCS. 24.2.2.1.3 Control code-groups The Control code-groups are used in pairs (/J/K/, /T/R/) to delimit MAC packets. Control code-groups are emitted from, and interpreted by, the PCS. 24.2.2.1.4 Start-of-Stream Delimiter (/J/K/) A Start-of-Stream Delimiter (SSD) is used to delineate the boundary of a data transmission sequence and to authenticate carrier events. The SSD is unique in that it may be recognized independently of previously established code-group boundaries. The Receive function within the PCS uses the SSD to establish code-group boundaries. A SSD consists of the sequence /J/K/. On transmission, the first 8 bits of the MAC preamble are replaced by the SSD, a replacement that is reversed on reception. 24.2.2.1.5 End-of-Stream delimiter (/T/R/) An End-of-Stream delimiter (ESD) terminates all normal data transmissions. Unlike the SSD, an ESD cannot be recognized independent of previously established code-group boundaries. An ESD consists of the sequence /T/R/.
--	--

Exhibit C-11

Table 24-1—4B/5B code-groups

	PCS code-group [4:0] 4 3 2 1 0	Name	MH (TXD/RXD) <3:0> 3 2 1 0	Interpretation
D A T A	1 1 1 1 0	0	0 0 0 0	Data 0
	0 1 0 0 1	1	0 0 0 1	Data 1
	1 0 1 0 0	2	0 0 1 0	Data 2
	1 0 1 0 1	3	0 0 1 1	Data 3
	0 1 0 1 0	4	0 1 0 0	Data 4
	0 1 0 1 1	5	0 1 0 1	Data 5
	0 1 1 1 0	6	0 1 1 0	Data 6
	0 1 1 1 1	7	0 1 1 1	Data 7
	1 0 0 1 0	8	1 0 0 0	Data 8
	1 0 0 1 1	9	1 0 0 1	Data 9
	1 0 1 1 0	A	1 0 1 0	Data A
	1 0 1 1 1	B	1 0 1 1	Data B
	1 1 0 1 0	C	1 1 0 0	Data C
	1 1 0 1 1	D	1 1 0 1	Data D
	1 1 1 0 0	E	1 1 1 0	Data E
	1 1 1 0 1	F	1 1 1 1	Data F
C O N T R O L	1 1 1 1 1	I	undefined	IDLE; used as inter-stream fill code
	1 1 0 0 0	J	0 1 0 1	Start-of-Stream Delimiter, Part 1 of 2; always used in pairs with K
	1 0 0 0 1	K	0 1 0 1	Start-of-Stream Delimiter, Part 2 of 2; always used in pairs with J
	0 1 1 0 1	T	undefined	End-of-Stream Delimiter, Part 1 of 2; always used in pairs with R
I N V A L I D	0 0 1 1 1	R	undefined	End-of-Stream Delimiter, Part 2 of 2; always used in pairs with T
	0 0 1 0 0	H	Undefined	Transmit Error; used to force signaling errors
	0 0 0 0 0	V	Undefined	Invalid code
	0 0 0 0 1	V	Undefined	Invalid code
	0 0 0 1 0	V	Undefined	Invalid code
	0 0 0 1 1	V	Undefined	Invalid code
	0 0 1 0 1	V	Undefined	Invalid code
	0 0 1 1 0	V	Undefined	Invalid code
	0 1 0 0 0	V	Undefined	Invalid code
	0 1 1 0 0	V	Undefined	Invalid code
	1 0 0 0 0	V	Undefined	Invalid code
	1 1 0 0 1	V	Undefined	Invalid code

24.2.2.1.6 Invalid code-groups

Exhibit C-11

		The /H/ code-group indicates that the PCS's client wishes to indicate a Transmit Error to its peer entity. The normal use of this indicator is for repeaters to propagate received errors. Transmit Error code-groups are emitted from the PCS, at the request of the PCS's client through the use of the TX_ER signal, as described in 24.2.4.2. The presence of any invalid code-group on the medium, including /H/, denotes a collision artifact or an error condition. Invalid code-groups are not intentionally transmitted onto the medium by DTE's. The PCS indicates the reception of an Invalid code-group on the MII through the use of the RX_ER signal, as described in 24.2.4.4.” 802.3-1998, p616-619.
--	--	--

Exhibit C-11




Figure 24-11—Receive state diagram

802.3-1998, p628.

“24.3.1.1 PMA_TYPE.indicate

Exhibit C-11

		This primitive is generated by the PMA to indicate the nature of the PMA instantiation. The purpose of this primitive is to allow clients to support connections to the various types of 100BASE-T PMA entities in a generalized manner. 24.3.1.1.1 Semantics of the service primitive PMA_TYPE.indicate (pma_type) The pma_type parameter for use with a 100BASE-X PMA is “X”. 802.3-1998, p629. “36.2.2 Functions within the PCS The PCS comprises the PCS Transmit, Carrier Sense, Synchronization, PCS Receive, and Auto-Negotiation processes for 1000BASE-X. The PCS shields the Reconciliation sublayer (and MAC) from the specific nature of the underlying channel. When communicating with the GMII, the PCS uses an octet-wide, synchronous data path, with packet delimiting being provided by separate transmit control signals (TX_EN and TX_ER) and receive control signals (RX_DV and RX_ER). When communicating with the PMA, the PCS uses a ten-bit wide, synchronous data path, which conveys ten-bit code-groups. At the PMA Service Interface, code-group alignment and MAC packet delimiting are made possible by embedding special non-data code-groups in the transmitted code-group stream. The PCS provides the functions necessary to map packets between the GMII format and the PMA Service Interface format. The PCS Transmit process continuously generates code-groups based upon the TXD <7:0>, TX_EN, and TX_ER signals on the GMII, sending them immediately to the PMA Service Interface via the PMA_UNITDATA.request primitive. The PCS Transmit process generates the GMII signal COL based on whether a reception is occurring simultaneously with transmission. Additionally, it generates the internal flag, transmitting, for use by the Carrier Sense process. The PCS Transmit process monitors the Auto-Negotiation process xmit flag to determine whether to transmit data or reconfigure the link. The Carrier Sense process controls the GMII signal CRS (see Figure 36-8). The PCS Synchronization process continuously accepts code-groups via the PMA_UNITDATA.indicate primitive and conveys received code-groups to the PCS Receive process via the SYNC_UNITDATA.indicate primitive. The PCS Synchronization process sets the sync_status flag to indicate whether the PMA is functioning dependably (as well as can be determined without exhaustive error-rate analysis). The PCS Receive process continuously accepts code-groups via the SYNC_UNITDATA.indicate primitive. The PCS Receive process monitors these code-groups and generates RXD <7:0>, RX_DV, and RX_ER on the GMII, and the internal flag, receiving, used by the Carrier Sense and Transmit processes. The
--	--	---

Exhibit C-11

	PCS Auto-Negotiation process sets the xmit flag to inform the PCS Transmit process to either transmit normal idles interspersed with packets as requested by the GMII or to reconfigure the link. The PCS Auto- Negotiation process is specified in Clause 37. 36.2.3 Use of code-groups The PCS maps GMII signals into ten-bit code groups, and vice versa, using an 8B/10B block coding scheme. Implicit in the definition of a code-group is an establishment of code-group boundaries by a PMA code-group alignment function as specified in 36.3.2.4. Code-groups are unobservable and have no meaning outside the PCS. The PCS functions ENCODE and DECODE generate, manipulate, and interpret codegroups as provided by the rules in 36.2.4. 36.2.4 8B/10B transmission code The PCS uses a transmission code to improve the transmission characteristics of information to be transferred across the link. The encodings defined by the transmission code ensure that sufficient transitions are present in the PHY bit stream to make clock recovery possible at the receiver. Such encoding also greatly increases the likelihood of detecting any single or multiple bit errors that may occur during transmission and reception of information. In addition, some of the special code-groups of the transmission code contain a distinct and easily recognizable bit pattern that assists a receiver in achieving code-group alignment on the incoming PHY bit stream. The 8B/10B transmission code specified for use in this standard has a high transition density, is a run-length-limited code, and is dc-balanced. The transition density of the 8B/10B symbols ranges from 3 to 8 transitions per symbol. The definition of the 8B/10B transmission code in this standard is identical to that specified in ANSI X3.230-1994 (FC-PH), Clause 11. The relationship of code-group bit positions to PMA and other PCS constructs is illustrated in Figure 36-3.
--	--

Exhibit C-11




Figure 36-3—PCS reference diagram

36.2.4.1 Notation conventions


8B/10B transmission code uses letter notation for describing the bits of an unencoded information octet and a single control variable. Each bit of the unencoded information octet contains either a binary zero or a binary one. A control variable, Z, has either the value D or the value K. When the control variable associated with an unencoded information octet contains the value D, the associated encoded code-group is referred to as a data code-group. When the control variable associated with an unencoded information octet contains the value K, the associated encoded code-group is referred to as a special code-group. The bit notation of A,B,C,D,E,F,G,H for an unencoded information octet is used in the description of the 8B/10B transmission code. The bits A,B,C,D,E,F,G,H are translated to bits a,b,c,d,e,i,f,g,h,j of 10-bit transmission code-groups. 8B/10B code-group bit assignments are illustrated in Figure 36-3. Each valid codegroup has been given a name using the following convention: /Dx.y/ for the 256 valid data code-groups, and /Kx.y/ for special control code-groups, where x is the decimal value of bits EDCBA, and y is the decimal value of bits HGF.”

802.3-1998, p927-928.

“36.2.5.2.2 Receive The PCS shall implement its Receive process as depicted in Figures 36-7a and 36-7b, including compliance with the associated state variables as specified in 36.2.5.1. The PCS

Exhibit C-11


Receive process continuously passes RXD<7:0> and sets the RX_DV and RX_ER signals to the GMII based on the received code-group from the PMA.



NOTE—Outgoing arcs leading to labeled polygons flow offpage to corresponding incoming arcs leading from labeled circles on Figure 36-7b, and vice versa.

Figure 36-7a—PCS receive state diagram, part a

Exhibit C-11



NOTE—Outgoing arcs leading to labeled polygons flow offpage to corresponding incoming arcs leading from labeled circles on Figure 36-7a, and vice versa.

Figure 36-7b—PCS receive state diagram, part b

When the Auto-Negotiation process xmit flag indicates CONFIGURATION or IDLE, the PCS Receive process continuously passes /C/ and /I/ ordered sets and rx_Config_Reg<D15:D0> to the Auto-Negotiation process.

Exhibit C-11

		36.2.5.2.3 State variable function carrier_detect(x) The detection of carrier on the underlying channel is used both by the MAC (via the GMII CRS signal and the Reconciliation sublayer) for deferral purposes, and by the PCS Transmit process for collision detection. A carrier event, signaled by the assertion of receiving, is indicated by the detection of a difference between the received code-group and /K28.5/ as specified in 36.2.5.1.4. A carrier event is in error if it does not start with an SPD. The PCS Receive process performs this function by continuously monitoring incoming code-groups for specific patterns that indicate non-/I/ activity such as SPD. The detection of an SPD carrier event causes the PCS to substitute the value (01010101) for the SPD, set RXD<7:0> to this value, and assert RX_DV. The pattern substituted for the SPD is consistent with the preamble pattern expected by the MAC. The detection of a non-SPD carrier event (false carrier) causes the PCS to substitute the value (00001110) for the code-group received, set RXD<7:0> to this value, and assert RX_ER. 36.2.5.2.4 Code-group stream decoding Subsequent to the detection of an SPD carrier event, the PCS Receive process performs the DECODE function on the incoming code-groups, passing decoded data to the GMII, including those corresponding to the remainder of the MAC preamble and SFD. The GMII signal RX_ER is asserted upon decoding any codegroup following the SPD that neither is a valid /D/ code-group nor follows the EPD rules in 36.2.4.14.1. Packets are terminated with an EPD as specified in 36.2.4.14. The PCS Receive process performs the check_end function to preserve the ability of the MAC to properly delimit the FCS at the end of a packet. Detection of /T/R/R/ or /T/R/K28.5/ by the check_end function denotes normal (i.e. non-error) packet termination. Detection of /R/R/R/ by the check_end function denotes packet termination with error and Carrier_Extend processing. Detection of /K28.5/D/K28.5/ by the check_end function denotes packet termination with error. Detection of /K28.5/(D21.5 or D2.2)/D0.0 by the check_end function denotes packet termination with error.” 802.3-1998, p947-952. To the extent that Plaintiff contends that 802.3-1998 does not explicitly or inherently disclose this element, or render obvious this element, this element is obvious in combination with one or more references and/or with the knowledge of a person of ordinary skill in the art, as set forth in Defendant’s Invalidity Contentions.
19(b)	a transmission component that transmits a	802.3-1998 discloses and/or renders obvious this limitation. <i>See, e.g.:</i>

Exhibit C-11

synchronization symbol that corresponds to the identified packet type, the transmitted synchronization symbol providing synchronization information and each packet type having a different synchronization symbol and permitting an external receiving node to properly align with a synchronization primitive to be correctly aligned on a symbol boundary, and that transmits the symbols of the packet.	“24.2.2 Functional requirements The PCS comprises the Transmit, Receive, and Carrier Sense functions for 100BASE-T. In addition, the collisionDetect signal required by the MAC (COL on the MII) is derived from the PMA code-bit stream. The PCS shields the Reconciliation sublayer (and MAC) from the specific nature of the underlying channel. Specifically for receiving, the 100BASE-X PCS passes to the MII a sequence of data nibbles derived from incoming code-groups, each comprised of five code-bits, received from the medium. Code-group alignment and MAC packet delimiting is performed by embedding special non-data code-groups. The MII uses a nibble- wide, synchronous data path, with packet delimiting being provided by separate TX_EN and RX_DV signals. The PCS provides the functions necessary to map these two views of the exchanged data. The process is reversed for transmit. The following provides a detailed specification of the functions performed by the PCS, which comprise five parallel processes (Transmit, Transmit Bits, Receive, Receive Bits, and Carrier Sense). Figure 24-4 includes a functional block diagram of the PCS. The Receive Bits process accepts continuous code-bits via the PMA_UNITDATA.indicate primitive. Receive monitors these bits and generates RXD <3:0>, RX_DV and RX_ER on the MII, and the internal flag, receiving, used by the Carrier Sense and Transmit processes. The Transmit process generates continuous code-groups based upon the TXD <3:0>, TX_EN, and TX_ER signals on the MII. These code-groups are transmitted by Transmit Bits via the PMA_UNITDATA.request primitive. The Transmit process generates the MII signal COL based on whether a reception is occurring simultaneously with transmission. Additionally, it generates the internal flag, transmitting, for use by the Carrier Sense process. The Carrier Sense process asserts the MII signal CRS when either transmitting or receiving is TRUE. Both the Transmit and Receive processes monitor link_status via the PMA_LINK.indicate primitive, to account for potential link failure conditions. 24.2.2.1 Code-groups The PCS maps four-bit nibbles from the MII into five-bit code-groups, and vice versa, using a 4B/5B block coding scheme. A code-group is a consecutive sequence of five code-bits interpreted and mapped by the PCS. Implicit in the definition of a code-group is an establishment of code-group boundaries by an alignment function within the PCS Receive process. It is important to note that, with the sole exception of the SSD, which is used to achieve alignment, code-groups are undetectable and have no meaning outside the 100BASE-X physical protocol data unit, called a “stream.” The coding method used, derived from ISO/IEC 9314-1, provides
--	---

Exhibit C-11

		a) Adequate codes (32) to provide for all Data code-groups (16) plus necessary control code-groups; b) Appropriate coding efficiency (4 data bits per 5 code-bits; 80%) to effect a 100 Mb/s Physical Layer interface on a 125 Mb/s physical channel as provided by FDDI PMDs; and c) Sufficient transition density to facilitate clock recovery (when not scrambled). Table 24-1 specifies the interpretation assigned to each five bit code-group, including the mapping to the nibble- wide (TXD or RXD) Data signals on the MII. The 32 code-groups are divided into four categories, as shown. For clarity in the remainder of this clause, code-group names are shown between /slashes/. Code-group sequences are shown in succession, e.g.: /1/2/.... The indicated code-group mapping is identical to ISO/IEC 9314-1: 1989, with four exceptions: a) The FDDI term symbol is avoided in order to prevent confusion with other 100BASE-T terminology. In general, the term code-group is used in its place. b) The /S/ and /Q/ code-groups are not used by 100BASE-X and are interpreted as INVALID. c) The /R/ code-group is used in 100BASE-X as the second code-group of the End-of-Stream delimiter rather than to indicate a Reset condition. d) The /H/ code-group is used to propagate receive errors rather than to indicate the Halt Line State. 24.2.2.1.1 Data code-groups A Data code-group conveys one nibble of arbitrary data between the MII and the PCS. The sequence of Data code-groups is arbitrary, where any Data code-group can be followed by any other Data code-group. Data code-groups are coded and decoded but not interpreted by the PCS. Successful decoding of Data code-groups depends on proper receipt of the Start-of-Stream delimiter sequence, as defined in Table 24-1. 24.2.2.1.2 Idle code-groups The Idle code-group (/I/) is transferred between streams. It provides a continuous fill pattern to establish and maintain clock synchronization. Idle code-groups are emitted from, and interpreted by, the PCS. 24.2.2.1.3 Control code-groups The Control code-groups are used in pairs (/J/K/, /T/R/) to delimit MAC packets. Control code-groups are emitted from, and interpreted by, the PCS. 24.2.2.1.4 Start-of-Stream Delimiter (/J/K/) A Start-of-Stream Delimiter (SSD) is used to delineate the boundary of a data transmission sequence and to authenticate carrier events. The SSD is unique in that it may be recognized independently of previously established code-group boundaries. The Receive function within the PCS uses the SSD to establish code-group boundaries. A SSD consists of the sequence /J/K/. On
--	--	--

Exhibit C-11

		transmission, the first 8 bits of the MAC preamble are replaced by the SSD, a replacement that is reversed on reception. 24.2.2.1.5 End-of-Stream delimiter (/T/R/) An End-of-Stream delimiter (ESD) terminates all normal data transmissions. Unlike the SSD, an ESD cannot be recognized independent of previously established code-group boundaries. An ESD consists of the sequence /T/R/.
--	--	---

Exhibit C-11

Table 24-1—4B/5B code-groups

	PCS code-group [4:0] 4 3 2 1 0	Name	MH (TXD/RXD) <3:0> 3 2 1 0	Interpretation
D A T A	1 1 1 1 0	0	0 0 0 0	Data 0
	0 1 0 0 1	1	0 0 0 1	Data 1
	1 0 1 0 0	2	0 0 1 0	Data 2
	1 0 1 0 1	3	0 0 1 1	Data 3
	0 1 0 1 0	4	0 1 0 0	Data 4
	0 1 0 1 1	5	0 1 0 1	Data 5
	0 1 1 1 0	6	0 1 1 0	Data 6
	0 1 1 1 1	7	0 1 1 1	Data 7
	1 0 0 1 0	8	1 0 0 0	Data 8
	1 0 0 1 1	9	1 0 0 1	Data 9
	1 0 1 1 0	A	1 0 1 0	Data A
	1 0 1 1 1	B	1 0 1 1	Data B
	1 1 0 1 0	C	1 1 0 0	Data C
	1 1 0 1 1	D	1 1 0 1	Data D
	1 1 1 0 0	E	1 1 1 0	Data E
	1 1 1 0 1	F	1 1 1 1	Data F
C O N T R O L	1 1 1 1 1	I	undefined	IDLE; used as inter-stream fill code
	1 1 0 0 0	J	0 1 0 1	Start-of-Stream Delimiter, Part 1 of 2; always used in pairs with K
	1 0 0 0 1	K	0 1 0 1	Start-of-Stream Delimiter, Part 2 of 2; always used in pairs with J
	0 1 1 0 1	T	undefined	End-of-Stream Delimiter, Part 1 of 2; always used in pairs with R
I N V A L I D	0 0 1 1 1	R	undefined	End-of-Stream Delimiter, Part 2 of 2; always used in pairs with T
	0 0 1 0 0	H	Undefined	Transmit Error; used to force signaling errors
	0 0 0 0 0	V	Undefined	Invalid code
	0 0 0 0 1	V	Undefined	Invalid code
	0 0 0 1 0	V	Undefined	Invalid code
	0 0 0 1 1	V	Undefined	Invalid code
	0 0 1 0 1	V	Undefined	Invalid code
	0 0 1 1 0	V	Undefined	Invalid code
	0 1 0 0 0	V	Undefined	Invalid code
	0 1 1 0 0	V	Undefined	Invalid code
	1 0 0 0 0	V	Undefined	Invalid code
	1 1 0 0 1	V	Undefined	Invalid code

24.2.2.1.6 Invalid code-groups

Exhibit C-11

The /H/ code-group indicates that the PCS's client wishes to indicate a Transmit Error to its peer entity. The normal use of this indicator is for repeaters to propagate received errors. Transmit Error code-groups are emitted from the PCS, at the request of the PCS's client through the use of the TX_ER signal, as described in 24.2.4.2. The presence of any invalid code-group on the medium, including /H/, denotes a collision artifact or an error condition. Invalid code-groups are not intentionally transmitted onto the medium by DTE's. The PCS indicates the reception of an Invalid code-group on the MII through the use of the RX_ER signal, as described in 24.2.4.4.”

802.3-1998, p616-619.

“24.2.2.4 Mapping between MII and PMA

Figure 24-6 depicts the mapping of the nibble-wide data path of the MII to the five-bit-wide code-groups (internal to the PCS) and the code-bit path of the PMA interface.




Figure 24-6—PCS reference diagram

802.3-1998, p620.

“24.2.4.1 Transmit Bits

Transmit Bits is responsible for taking code-groups prepared by the Transmit process and transmitting them to the PMA using PMA_UNITDATA.request, the frequency of which determines the transmit clock. Transmit deposits these code-groups in tx_bits with Transmit Bits signaling completion of a code-group transmission with sentCodeGroup.indicate. The PCS shall implement the Transmit Bits process as depicted in Figure 24-7 including compliance with the associated state variables as specified in 24.2.3.

24.2.4.2 Transmit

Exhibit C-11

The Transmit process sends code-groups to the PMA via tx_bits and the Transmit Bits process. When initially invoked, and between streams (delimited by TX_EN on the MII), the Transmit process sources continuous Idle code-groups (/I/) to the PMA. Upon the assertion of TX_EN by the MII, the Transmit process passes an SSD (/J/K/) to the PMA, ignoring the TXD <3:0> nibbles during these two code-group times. Following the SSD, each TXD <3:0> nibble is encoded into a five-bit code-group until TX_EN is deasserted. If, while TX_EN is asserted, the TX_ER signal is asserted, the Transmit process passes Transmit Error code-groups (/H/) to the PMA. Following the de-assertion of TX_EN, an ESD (/T/R/) is generated, after which the transmission of Idle code-groups is resumed by the IDLE state. Collision detection is implemented by noting the occurrence of carrier receptions during transmissions, following the model of 10BASE-T. The indication of link_status ¹ OK by the PMA at any time causes an immediate transition to the IDLE state and supersedes any other Transmit process operations. The PCS shall implement the Transmit process as depicted in Figure 24-8 including compliance with the associated state variables as specified in 24.2.3.




Figure 24-7—Transmit Bits state diagram

24.2.4.3 Receive Bits The Receive Bits process collects code-bits from the PMA interface passing them to the Receive process via rx_bits. rx_bits [9:0] represents a sliding, 10-bit window on the

Exhibit C-11

PMA code-bits, with newly received code-bits from the PMA (rx_code-bit) being shifted into rx_bits [0]. This is depicted in Figure 24-9. Bits are collected serially until Receive indicates alignment by asserting RX_DV, after which Receive Bits signals Receive for every five code-bits accumulated. Serial processing resumes with the de-assertion of RX_DV. The PCS shall implement the Receive Bits process as depicted in Figure 24-10 including compliance with the associated state variables as specified in 24.2.3.

24.2.4.4 Receive The Receive process state machine can be viewed as comprising two sections: prealigned and aligned. In the prealigned states, IDLE, CARRIER DETECT, and CONFIRM K, the Receive process is waiting for an indication of channel activity followed by a SSD. After successful alignment, the incoming code-groups are decoded while waiting for stream termination.” 802.3-1998, p623-624.

“24.2.4.4.2 Code-group alignment

After channel activity is detected, the Receive process first aligns the incoming code-bits on code-group boundaries for subsequent data decoding. This is achieved by scanning the rx_bits vector for a SSD (/J/K/). The MII RX_DV signal remains deasserted during this time, which ensures that the Reconciliation sublayer will ignore any signals on RXD <3:0>. Detection of the SSD causes the Receive process to enter the START OF STREAM J state.




Figure 24-9—Receive Bits reference diagram

Well-formed streams contain SSD (/J/K/) in place of the first eight preamble bits. In the event that something else is sensed immediately following detection of carrier, a False Carrier Indication is signaled to the MII by asserting RX_ER and setting RXD to 1110 while RX_DV remains deasserted. The associated carrier event, as terminated by 10 ONEs, is otherwise ignored.” 802.3-1998, p625-626.

“24.3.1.1 PMA_TYPE.indicate

Exhibit C-11

	This primitive is generated by the PMA to indicate the nature of the PMA instantiation. The purpose of this primitive is to allow clients to support connections to the various types of 100BASE-T PMA entities in a generalized manner. 24.3.1.1.1 Semantics of the service primitive PMA_TYPE.indicate (pma_type) The pma_type parameter for use with a 100BASE-X PMA is “X”. 802.3-1998, p629. “36.2.2 Functions within the PCS The PCS comprises the PCS Transmit, Carrier Sense, Synchronization, PCS Receive, and Auto-Negotiation processes for 1000BASE-X. The PCS shields the Reconciliation sublayer (and MAC) from the specific nature of the underlying channel. When communicating with the GMII, the PCS uses an octet-wide, synchronous data path, with packet delimiting being provided by separate transmit control signals (TX_EN and TX_ER) and receive control signals (RX_DV and RX_ER). When communicating with the PMA, the PCS uses a ten-bit wide, synchronous data path, which conveys ten-bit code-groups. At the PMA Service Interface, code-group alignment and MAC packet delimiting are made possible by embedding special non-data code-groups in the transmitted code-group stream. The PCS provides the functions necessary to map packets between the GMII format and the PMA Service Interface format. The PCS Transmit process continuously generates code-groups based upon the TXD <7:0>, TX_EN, and TX_ER signals on the GMII, sending them immediately to the PMA Service Interface via the PMA_UNITDATA.request primitive. The PCS Transmit process generates the GMII signal COL based on whether a reception is occurring simultaneously with transmission. Additionally, it generates the internal flag, transmitting, for use by the Carrier Sense process. The PCS Transmit process monitors the Auto-Negotiation process xmit flag to determine whether to transmit data or reconfigure the link. The Carrier Sense process controls the GMII signal CRS (see Figure 36-8). The PCS Synchronization process continuously accepts code-groups via the PMA_UNITDATA.indicate primitive and conveys received code-groups to the PCS Receive process via the SYNC_UNITDATA.indicate primitive. The PCS Synchronization process sets the sync_status flag to indicate whether the PMA is functioning dependably (as well as can be determined without exhaustive error-rate analysis). The PCS Receive process continuously accepts code-groups via the SYNC_UNITDATA.indicate primitive. The PCS Receive process monitors these code-groups and generates RXD <7:0>, RX_DV, and RX_ER on the GMII, and the internal flag, receiving, used by the Carrier Sense and Transmit processes. The
--	---

Exhibit C-11

	PCS Auto-Negotiation process sets the xmit flag to inform the PCS Transmit process to either transmit normal idles interspersed with packets as requested by the GMII or to reconfigure the link. The PCS Auto- Negotiation process is specified in Clause 37. 36.2.3 Use of code-groups The PCS maps GMII signals into ten-bit code groups, and vice versa, using an 8B/10B block coding scheme. Implicit in the definition of a code-group is an establishment of code-group boundaries by a PMA code-group alignment function as specified in 36.3.2.4. Code-groups are unobservable and have no meaning outside the PCS. The PCS functions ENCODE and DECODE generate, manipulate, and interpret codegroups as provided by the rules in 36.2.4. 36.2.4 8B/10B transmission code The PCS uses a transmission code to improve the transmission characteristics of information to be transferred across the link. The encodings defined by the transmission code ensure that sufficient transitions are present in the PHY bit stream to make clock recovery possible at the receiver. Such encoding also greatly increases the likelihood of detecting any single or multiple bit errors that may occur during transmission and reception of information. In addition, some of the special code-groups of the transmission code contain a distinct and easily recognizable bit pattern that assists a receiver in achieving code-group alignment on the incoming PHY bit stream. The 8B/10B transmission code specified for use in this standard has a high transition density, is a run-length-limited code, and is dc-balanced. The transition density of the 8B/10B symbols ranges from 3 to 8 transitions per symbol. The definition of the 8B/10B transmission code in this standard is identical to that specified in ANSI X3.230-1994 (FC-PH), Clause 11. The relationship of code-group bit positions to PMA and other PCS constructs is illustrated in Figure 36-3.
--	--

Exhibit C-11




Figure 36-3—PCS reference diagram

36.2.4.1 Notation conventions

8B/10B transmission code uses letter notation for describing the bits of an unencoded information octet and a single control variable. Each bit of the unencoded information octet contains either a binary zero or a binary one. A control variable, Z, has either the value D or the value K. When the control variable associated with an unencoded information octet contains the value D, the associated encoded code-group is referred to as a data code-group. When the control variable associated with an unencoded information octet contains the value K, the associated encoded code-group is referred to as a special code-group. The bit notation of A,B,C,D,E,F,G,H for an unencoded information octet is used in the description of the 8B/10B transmission code. The bits A,B,C,D,E,F,G,H are translated to bits a,b,c,d,e,i,f,g,h,j of 10-bit transmission code-groups. 8B/10B code-group bit assignments are illustrated in Figure 36-3. Each valid codegroup has been given a name using the following convention: /Dx.y/ for the 256 valid data code-groups, and /Kx.y/ for special control code-groups, where x is the decimal value of bits EDCBA, and y is the decimal value of bits HGF.

36.2.4.2 Transmission order

Code-group bit transmission order is illustrated in Figure 36-3. Code-groups within multi-code-group ordered_sets (as specified in Table 36-3) are transmitted sequentially beginning with the special code-group used to distinguish the ordered_set (e.g., /K28.5/) and proceeding code-group

Exhibit C-11

by code-group from left to right within the definition of the ordered_set until all code-groups of the ordered_set are transmitted.

The first code-group of every multi-code-group ordered_set is transmitted in an even-numbered code-group position counting from the first code-group after a reset or power-on. Subsequent code-groups continuously alternate as odd and even-numbered code-groups. The contents of a packet are transmitted sequentially beginning with the ordered_set used to denote the Start_of_Packet (the SPD delimiter) and proceeding code-group by code-group from left to right within the definition of the packet until the ordered_set used to denote the End_of_Packet (the EPD delimiter) is transmitted.

36.2.4.3 Valid and invalid code-groups

Table 36-1a defines the valid data code-groups (D code-groups) of the 8B/10B transmission code. Table 36- 2 defines the valid special code-groups (K code-groups) of the code. The tables are used for both generating valid code-groups (encoding) and checking the validity of received code-groups (decoding). In the tables, each octet entry has two columns that represent two (not necessarily different) code-groups. The two columns correspond to the valid code-group based on the current value of the running disparity (Current RD – or Current RD +). Running disparity is a binary parameter with either the value negative (–) or the value positive (+). Annex 36B provides several 8B/10B transmission code running disparity calculation examples.”

802.3-1998, p927-929.

Table 36-1a—Valid data code-groups

Code Group Name	Octet Value	Octet Bits HGFEDCBA	Current RD –	Current RD +
			abcdei fghj	abcdei fghj
D0.0	00	000 00000	100111 0100	011000 1011
D1.0	01	000 00001	011101 0100	100010 1011
D2.0	02	000 00010	101101 0100	010010 1011
D3.0	03	000 00011	110001 1011	110001 0100
D4.0	04	000 00100	110101 0100	001010 1011
D5.0	05	000 00101	101001 1011	101001 0100
D6.0	06	000 00110	011001 1011	011001 0100
D7.0	07	000 00111	111000 1011	000111 0100
D8.0	08	000 01000	111001 0100	000110 1011
D9.0	09	000 01001	100101 1011	100101 0100
D10.0	0A	000 01010	010101 1011	010101 0100
D11.0	0B	000 01011	110100 1011	110100 0100
D12.0	0C	000 01100	001101 1011	001101 0100
D13.0	0D	000 01101	101100 1011	101100 0100
D14.0	0E	000 01110	011100 1011	011100 0100
D15.0	0F	000 01111	010111 0100	101000 1011
D16.0	10	000 10000	011011 0100	100100 1011

“

Exhibit C-11

...

Table 36-1e—Valid data code-groups

Code Group Name	Octet Value	Octet Bits HGF EDCBA	Current RD –	Current RD +
			abedei fghj	abedei fghj
D16.7	F0	111 10000	011011 0001	100100 1110
D17.7	F1	111 10001	100011 0111	100011 0001
D18.7	F2	111 10010	010011 0111	010011 0001
D19.7	F3	111 10011	110010 1110	110010 0001
D20.7	F4	111 10100	001011 0111	001011 0001
D21.7	F5	111 10101	101010 1110	101010 0001
D22.7	F6	111 10110	011010 1110	011010 0001
D23.7	F7	111 10111	111010 0001	000101 1110
D24.7	F8	111 11000	110011 0001	001100 1110
D25.7	F9	111 11001	100110 1110	100110 0001
D26.7	FA	111 11010	010110 1110	010110 0001
D27.7	FB	111 11011	110110 0001	001001 1110
D28.7	FC	111 11100	001110 1110	001110 0001
D29.7	FD	111 11101	101110 0001	010001 1110
D30.7	FE	111 11110	011110 0001	100001 1110
D31.7	FF	111 11111	101011 0001	010100 1110
(concluded)				

Table 36-2—Valid special code-groups

Code Group Name	Octet Value	Octet Bits HGF EDCBA	Current RD –	Current RD +	Notes
			abedei fghj	abedei fghj	
K28.0	1C	000 11100	001111 0100	110000 1011	1
K28.1	3C	001 11100	001111 1001	110000 0110	1,2
K28.2	5C	010 11100	001111 0101	110000 1010	1
K28.3	7C	011 11100	001111 0011	110000 1100	1
K28.4	9C	100 11100	001111 0010	110000 1101	1
K28.5	BC	101 11100	001111 1010	110000 0101	2
K28.6	DC	110 11100	001111 0110	110000 1001	1
K28.7	FC	111 11100	001111 1000	110000 0111	1,2
K23.7	F7	111 10111	111010 1000	000101 0111	
K27.7	FB	111 11011	110110 1000	001001 0111	
K29.7	FD	111 11101	101110 1000	010001 0111	
K30.7	FE	111 11110	011110 1000	100001 0111	
NOTE 1—Reserved.					
NOTE 2—Contains a comma.					

36.2.4.7 Ordered_sets Eight ordered_sets, consisting of a single special code-group or combinations of special and data codegroups are specifically defined. Ordered_sets which include /K28.5/ provide the ability to obtain bit and code-group synchronization and establish ordered_set

Exhibit C-11

alignment (see 36.2.4.9 and 36.3.2.4). Ordered_sets provide for the delineation of a packet and synchronization between the transmitter and receiver circuits at opposite ends of a link. Table 36-3 lists the defined ordered_sets.

Table 36-3—Defined ordered_sets

Code	Ordered_Set	Number of Code-Groups	Encoding
/C/	Configuration		Alternating /C1/ and /C2/
/C1/	Configuration 1	4	/K28.5/D21.5/Config_Reg ^a
/C2/	Configuration 2	4	/K28.5/D2.2/Config_Reg ^a
/I/	IDLE		Correcting /I1/, Preserving /I2/
/I1/	IDLE 1	2	/K28.5/D5.6/
/I2/	IDLE 2	2	/K28.5/D16.2/
	Encapsulation		
/R/	Carrier_Extend	1	/K23.7/
/S/	Start_of_Packet	1	/K27.7/
/T/	End_of_Packet	1	/K29.7/
/V/	Error_Propagation	1	/K30.7/

^aTwo data code-groups representing the Config_Reg value.

802.3-1998, p931-936.

“36.2.5.2.6 Synchronization

The PCS shall implement the Synchronization process as depicted in Figure 36-9 including compliance with the associated state variables as specified in 36.2.5.1. The Synchronization process is responsible for determining whether the underlying receive channel is ready for operation. Failure of the underlying channel typically causes the PMA client to suspend normal actions. A receiver that is in the LOSS_OF_SYNC state and that has acquired bit synchronization attempts to acquire code-group synchronization via the Synchronization process. Code-group synchronization is acquired by the detection of three ordered_sets containing commas in their leftmost bit positions without intervening invalid code-group errors. Upon acquisition of code-group synchronization, the receiver enters the SYNC_ACQUIRED_1 state. Acquisition of synchronization ensures the alignment of multi-code-group ordered_sets to even-numbered code-group boundaries. Once synchronization is acquired, the Synchronization process tests received code-groups in sets of four code-groups and employs multiple sub-states, effecting hysteresis, to move between the SYNC_ACQUIRED_1 and LOSS_OF_SYNC states. The condition sync_status=FAIL existing for

Exhibit C-11

		ten ms or more causes the PCS Auto-Negotiation process to begin and the PCS Transmit process to begin transmission of /C/. Upon reception of three matching /C/s from the link partner, the PCS Auto-Negotiation process begins. The internal signal receiving is de-asserted in the PCS Receive process LINK_FAILED state when sync_status=FAIL and a code-group is received.” 802.3-1998, p953. “36.3.2.4 Code-group alignment In the event the PMA sublayer detects a comma+ within the incoming rx_bit stream, it may realign its current code-group boundary, if necessary, to that of the received comma+ as shown in Figure 36-3. This process is referred to in this document as code-group alignment. The code-group alignment function shall be operational when the EN_CDET signal is active (see 36.3.3.1). During the code-group alignment process, the PMA sublayer may delete or modify up to four, but shall delete or modify no more than four, ten-bit code-groups in order to align the correct receive clock and code-group containing the comma+. This process is referred to as code-group slipping. In addition, the PMA sublayer is permitted to realign the current code-group boundary upon receipt of a comma- pattern.” 802.3-1998, p956. “Code-group alignment occurs in the PMA Receive function, if enabled by EN_CDET, when a comma pattern occurs in the PHY bit stream. Upon recognition of the comma pattern, the PMA Receive function outputs the ten-bit code-group containing the comma on rx_code-group<9:0> with the alignment specified in Figure 36-3, and clocked on the rising edge of PMA_RX_CLK<1>.” 802.3-1998, p957. To the extent that Plaintiff contends that 802.3-1998 does not explicitly or inherently disclose this element, or render obvious this element, this element is obvious in combination with one or more references and/or with the knowledge of a person of ordinary skill in the art, as set forth in Defendant’s Invalidity Contentions.
Claim 22		
22[p]	The communications device of claim 19	802.3-1998 discloses, expressly or inherently, the transmitting device of claim 19 under any proposed construction.

Exhibit C-11

		See exemplary disclosures and corresponding discussion for claim 19.
22	wherein the synchronization symbol is transmitted before transmitting the symbols of the packet.	802.3-1998 discloses and/or renders obvious this limitation. <i>See, e.g.:</i> “24.2.4.1 Transmit Bits Transmit Bits is responsible for taking code-groups prepared by the Transmit process and transmitting them to the PMA using PMA_UNITDATA.request, the frequency of which determines the transmit clock. Transmit deposits these code-groups in tx_bits with Transmit Bits signaling completion of a code-group transmission with sentCodeGroup.indicate. The PCS shall implement the Transmit Bits process as depicted in Figure 24-7 including compliance with the associated state variables as specified in 24.2.3. 24.2.4.2 Transmit The Transmit process sends code-groups to the PMA via tx_bits and the Transmit Bits process. When initially invoked, and between streams (delimited by TX_EN on the MII), the Transmit process sources continuous Idle code-groups (/I/) to the PMA. Upon the assertion of TX_EN by the MII, the Transmit process passes an SSD (/J/K/) to the PMA, ignoring the TXD <3:0> nibbles during these two code-group times. Following the SSD, each TXD <3:0> nibble is encoded into a five-bit code-group until TX_EN is deasserted. If, while TX_EN is asserted, the TX_ER signal is asserted, the Transmit process passes Transmit Error code-groups (/H/) to the PMA. Following the de-assertion of TX_EN, an ESD (/T/R/) is generated, after which the transmission of Idle code-groups is resumed by the IDLE state. Collision detection is implemented by noting the occurrence of carrier receptions during transmissions, following the model of 10BASE-T. The indication of link_status ¹ OK by the PMA at any time causes an immediate transition to the IDLE state and supersedes any other Transmit process operations. The PCS shall implement the Transmit process as depicted in Figure 24-8 including compliance with the associated state variables as specified in 24.2.3.

Exhibit C-11




Figure 24-7—Transmit Bits state diagram

24.2.4.3 Receive Bits The Receive Bits process collects code-bits from the PMA interface passing them to the Receive process via rx_bits. rx_bits [9:0] represents a sliding, 10-bit window on the PMA code-bits, with newly received code-bits from the PMA (rx_code-bit) being shifted into rx_bits [0]. This is depicted in Figure 24-9. Bits are collected serially until Receive indicates alignment by asserting RX_DV, after which Receive Bits signals Receive for every five code-bits accumulated. Serial processing resumes with the de-assertion of RX_DV. The PCS shall implement the Receive Bits process as depicted in Figure 24-10 including compliance with the associated state variables as specified in 24.2.3.

24.2.4.4 Receive The Receive process state machine can be viewed as comprising two sections: prealigned and aligned. In the prealigned states, IDLE, CARRIER DETECT, and CONFIRM K, the Receive process is waiting for an indication of channel activity followed by a SSD. After successful alignment, the incoming code-groups are decoded while waiting for stream termination.

Exhibit C-11




Figure 24-8—Transmit state diagram

802.3-1998, p623-625.

Exhibit C-11

	“36.2.5.2.1 Transmit The PCS Transmit process is depicted in two state diagrams: PCS Transmit ordered_set and PCS Transmit code-group. The PCS shall implement its Transmit process as depicted in Figures 36-5 and 36-6, including compliance with the associated state variables as specified in 36.2.5.1. The Transmit ordered_set process continuously sources ordered_sets to the Transmit code-group process. When initially invoked, and when the Auto-Negotiation process xmit flag indicates CONFIGURATION, the Auto-Negotiation process is invoked. When the Auto-Negotiation process xmit flag indicates IDLE, and between packets (as delimited by the GMII), /I/ is sourced. Upon the assertion of TX_EN by the GMII when the Auto-Negotiation process xmit flag indicates DATA, the SPD ordered_set is sourced. Following the SPD, /D/ code-groups are sourced until TX_EN is deasserted. Following the de-assertion of TX_EN, EPD ordered_sets are sourced. If TX_ER is asserted when TX_EN is deasserted and carrier extend error is not indicated by TXD, /R/ ordered_sets are sourced for as many GTX_CLK periods as TX_ER is asserted with a delay of two GTX_CLK periods to first source the /T/ and /R/ ordered sets. If carrier extend error is indicated by TXD during carrier extend, /V/ ordered_sets are sourced. If TX_EN and TX_ER are both deasserted, the /R/ ordered_set may be sourced, after which the sourcing of /I/ is resumed. If, while TX_EN is asserted, the TX_ER signal is asserted, the /V/ ordered_set is sourced except when the SPD ordered set is selected for sourcing. Collision detection is implemented by noting the occurrence of carrier receptions during transmissions, following the models of 10BASE-T and 100BASE-X. The Transmit code-group process continuously sources tx_code-group<9:0> to the PMA based on the ordered_sets sourced to it by the Transmit ordered_set process. The Transmit code-group process determines the proper code-group to source based on even/odd-numbered code-group alignment, running disparity requirements, and ordered_set format.
--	--

Exhibit C-11




Figure 36-5—PCS transmit ordered_set state diagram

Exhibit C-11




Exhibit C-11

		802.3-1998, p947-949. To the extent that Plaintiff contends that 802.3-1998 does not explicitly or inherently disclose this element, or render obvious this element, this element is obvious in combination with one or more references and/or with the knowledge of a person of ordinary skill in the art, as set forth in Defendant's Invalidity Contentions.
Claim 25		
25[p]	The communications device of claim 19	802.3-1998 discloses, expressly or inherently, the transmitting device of claim 19 under any proposed construction. <i>See exemplary disclosures and corresponding discussion for claim 19.</i>
25	wherein the packet types include a control packet.	802.3-1998 discloses and/or renders obvious this limitation. <i>See, e.g.:</i> “24.2.2 Functional requirements The PCS comprises the Transmit, Receive, and Carrier Sense functions for 100BASE-T. In addition, the collisionDetect signal required by the MAC (COL on the MII) is derived from the PMA code-bit stream. The PCS shields the Reconciliation sublayer (and MAC) from the specific nature of the underlying channel. Specifically for receiving, the 100BASE-X PCS passes to the MII a sequence of data nibbles derived from incoming code-groups, each comprised of five code-bits, received from the medium. Code-group alignment and MAC packet delimiting is performed by embedding special non-data code-groups. The MII uses a nibble- wide, synchronous data path, with packet delimiting being provided by separate TX_EN and RX_DV signals. The PCS provides the functions necessary to map these two views of the exchanged data. The process is reversed for transmit. The following provides a detailed specification of the functions performed by the PCS, which comprise five parallel processes (Transmit, Transmit Bits, Receive, Receive Bits, and Carrier Sense). Figure 24-4 includes a functional block diagram of the PCS. The Receive Bits process accepts continuous code-bits via the PMA_UNITDATA.indicate primitive. Receive monitors these bits and generates RXD <3:0>, RX_DV and RX_ER on the MII, and the internal flag, receiving, used by the Carrier Sense and Transmit processes. The Transmit process generates continuous code-groups based upon the TXD <3:0>, TX_EN, and TX_ER signals

Exhibit C-11

		on the MII. These code-groups are transmitted by Transmit Bits via the PMA_UNITDATA.request primitive. The Transmit process generates the MII signal COL based on whether a reception is occurring simultaneously with transmission. Additionally, it generates the internal flag, transmitting, for use by the Carrier Sense process. The Carrier Sense process asserts the MII signal CRS when either transmitting or receiving is TRUE. Both the Transmit and Receive processes monitor link_status via the PMA_LINK.indicate primitive, to account for potential link failure conditions. 24.2.2.1 Code-groups The PCS maps four-bit nibbles from the MII into five-bit code-groups, and vice versa, using a 4B/5B block coding scheme. A code-group is a consecutive sequence of five code-bits interpreted and mapped by the PCS. Implicit in the definition of a code-group is an establishment of code-group boundaries by an alignment function within the PCS Receive process. It is important to note that, with the sole exception of the SSD, which is used to achieve alignment, code-groups are undetectable and have no meaning outside the 100BASE-X physical protocol data unit, called a “stream.” The coding method used, derived from ISO/IEC 9314-1, provides a) Adequate codes (32) to provide for all Data code-groups (16) plus necessary control code-groups; b) Appropriate coding efficiency (4 data bits per 5 code-bits; 80%) to effect a 100 Mb/s Physical Layer interface on a 125 Mb/s physical channel as provided by FDDI PMDs; and c) Sufficient transition density to facilitate clock recovery (when not scrambled). Table 24-1 specifies the interpretation assigned to each five bit code-group, including the mapping to the nibble- wide (TXD or RXD) Data signals on the MII. The 32 code-groups are divided into four categories, as shown. For clarity in the remainder of this clause, code-group names are shown between /slashes/. Code-group sequences are shown in succession, e.g.: /1/2/.... The indicated code-group mapping is identical to ISO/IEC 9314-1: 1989, with four exceptions: a) The FDDI term symbol is avoided in order to prevent confusion with other 100BASE-T terminology. In general, the term code-group is used in its place. b) The /S/ and /Q/ code-groups are not used by 100BASE-X and are interpreted as INVALID. c) The /R/ code-group is used in 100BASE-X as the second code-group of the End-of-Stream delimiter rather than to indicate a Reset condition. d) The /H/ code-group is used to propagate receive errors rather than to indicate the Halt Line State. 24.2.2.1.1 Data code-groups A Data code-group conveys one nibble of arbitrary data between the MII and the PCS. The sequence of Data code-groups is arbitrary, where any Data code-group can be followed by any other Data
--	--	---

Exhibit C-11

		code-group. Data code-groups are coded and decoded but not interpreted by the PCS. Successful decoding of Data code-groups depends on proper receipt of the Start-of-Stream delimiter sequence, as defined in Table 24-1. 24.2.2.1.2 Idle code-groups The Idle code-group (/I/) is transferred between streams. It provides a continuous fill pattern to establish and maintain clock synchronization. Idle code-groups are emitted from, and interpreted by, the PCS. 24.2.2.1.3 Control code-groups The Control code-groups are used in pairs (/J/K/, /T/R/) to delimit MAC packets. Control code-groups are emitted from, and interpreted by, the PCS. 24.2.2.1.4 Start-of-Stream Delimiter (/J/K/) A Start-of-Stream Delimiter (SSD) is used to delineate the boundary of a data transmission sequence and to authenticate carrier events. The SSD is unique in that it may be recognized independently of previously established code-group boundaries. The Receive function within the PCS uses the SSD to establish code-group boundaries. A SSD consists of the sequence /J/K/. On transmission, the first 8 bits of the MAC preamble are replaced by the SSD, a replacement that is reversed on reception. 24.2.2.1.5 End-of-Stream delimiter (/T/R/) An End-of-Stream delimiter (ESD) terminates all normal data transmissions. Unlike the SSD, an ESD cannot be recognized independent of previously established code-group boundaries. An ESD consists of the sequence /T/R/.
--	--	--

Exhibit C-11

Table 24-1—4B/5B code-groups

	PCS code-group [4:0] 4 3 2 1 0	Name	MH (TXD/RXD) <3:0> 3 2 1 0	Interpretation
D A T A	1 1 1 1 0	0	0 0 0 0	Data 0
	0 1 0 0 1	1	0 0 0 1	Data 1
	1 0 1 0 0	2	0 0 1 0	Data 2
	1 0 1 0 1	3	0 0 1 1	Data 3
	0 1 0 1 0	4	0 1 0 0	Data 4
	0 1 0 1 1	5	0 1 0 1	Data 5
	0 1 1 1 0	6	0 1 1 0	Data 6
	0 1 1 1 1	7	0 1 1 1	Data 7
	1 0 0 1 0	8	1 0 0 0	Data 8
	1 0 0 1 1	9	1 0 0 1	Data 9
	1 0 1 1 0	A	1 0 1 0	Data A
	1 0 1 1 1	B	1 0 1 1	Data B
	1 1 0 1 0	C	1 1 0 0	Data C
	1 1 0 1 1	D	1 1 0 1	Data D
	1 1 1 0 0	E	1 1 1 0	Data E
	1 1 1 0 1	F	1 1 1 1	Data F
C O N T R O L	1 1 1 1 1	I	undefined	IDLE; used as inter-stream fill code
	1 1 0 0 0	J	0 1 0 1	Start-of-Stream Delimiter, Part 1 of 2; always used in pairs with K
	1 0 0 0 1	K	0 1 0 1	Start-of-Stream Delimiter, Part 2 of 2; always used in pairs with J
	0 1 1 0 1	T	undefined	End-of-Stream Delimiter, Part 1 of 2; always used in pairs with R
I N V A L I D	0 0 1 1 1	R	undefined	End-of-Stream Delimiter, Part 2 of 2; always used in pairs with T
	0 0 1 0 0	H	Undefined	Transmit Error; used to force signaling errors
	0 0 0 0 0	V	Undefined	Invalid code
	0 0 0 0 1	V	Undefined	Invalid code
	0 0 0 1 0	V	Undefined	Invalid code
	0 0 0 1 1	V	Undefined	Invalid code
	0 0 1 0 1	V	Undefined	Invalid code
	0 0 1 1 0	V	Undefined	Invalid code
	0 1 0 0 0	V	Undefined	Invalid code
	0 1 1 0 0	V	Undefined	Invalid code
	1 0 0 0 0	V	Undefined	Invalid code
	1 1 0 0 1	V	Undefined	Invalid code

24.2.2.1.6 Invalid code-groups

Exhibit C-11

		The /H/ code-group indicates that the PCS's client wishes to indicate a Transmit Error to its peer entity. The normal use of this indicator is for repeaters to propagate received errors. Transmit Error code-groups are emitted from the PCS, at the request of the PCS's client through the use of the TX_ER signal, as described in 24.2.4.2. The presence of any invalid code-group on the medium, including /H/, denotes a collision artifact or an error condition. Invalid code-groups are not intentionally transmitted onto the medium by DTE's. The PCS indicates the reception of an Invalid code-group on the MII through the use of the RX_ER signal, as described in 24.2.4.4.” 802.3-1998, p616-619. “36.2.2 Functions within the PCS The PCS comprises the PCS Transmit, Carrier Sense, Synchronization, PCS Receive, and Auto-Negotiation processes for 1000BASE-X. The PCS shields the Reconciliation sublayer (and MAC) from the specific nature of the underlying channel. When communicating with the GMII, the PCS uses an octet-wide, synchronous data path, with packet delimiting being provided by separate transmit control signals (TX_EN and TX_ER) and receive control signals (RX_DV and RX_ER). When communicating with the PMA, the PCS uses a ten-bit wide, synchronous data path, which conveys ten-bit code-groups. At the PMA Service Interface, code-group alignment and MAC packet delimiting are made possible by embedding special non-data code-groups in the transmitted code-group stream. The PCS provides the functions necessary to map packets between the GMII format and the PMA Service Interface format. The PCS Transmit process continuously generates code-groups based upon the TXD <7:0>, TX_EN, and TX_ER signals on the GMII, sending them immediately to the PMA Service Interface via the PMA_UNITDATA.request primitive. The PCS Transmit process generates the GMII signal COL based on whether a reception is occurring simultaneously with transmission. Additionally, it generates the internal flag, transmitting, for use by the Carrier Sense process. The PCS Transmit process monitors the Auto-Negotiation process xmit flag to determine whether to transmit data or reconfigure the link. The Carrier Sense process controls the GMII signal CRS (see Figure 36-8). The PCS Synchronization process continuously accepts code-groups via the PMA_UNITDATA.indicate primitive and conveys received code-groups to the PCS Receive process via the SYNC_UNITDATA.indicate primitive. The PCS Synchronization process sets the sync_status flag to indicate whether the PMA is functioning dependably (as well as can be determined without exhaustive error-rate analysis). The PCS Receive process continuously accepts code-groups via the SYNC_UNITDATA.indicate primitive. The PCS Receive process monitors these code-groups and generates RXD <7:0>, RX_DV, and RX_ER on
--	--	---

Exhibit C-11

	the GMII, and the internal flag, receiving, used by the Carrier Sense and Transmit processes. The PCS Auto-Negotiation process sets the xmit flag to inform the PCS Transmit process to either transmit normal idles interspersed with packets as requested by the GMII or to reconfigure the link. The PCS Auto- Negotiation process is specified in Clause 37. 36.2.3 Use of code-groups The PCS maps GMII signals into ten-bit code groups, and vice versa, using an 8B/10B block coding scheme. Implicit in the definition of a code-group is an establishment of code-group boundaries by a PMA code-group alignment function as specified in 36.3.2.4. Code-groups are unobservable and have no meaning outside the PCS. The PCS functions ENCODE and DECODE generate, manipulate, and interpret codegroups as provided by the rules in 36.2.4. 36.2.4 8B/10B transmission code The PCS uses a transmission code to improve the transmission characteristics of information to be transferred across the link. The encodings defined by the transmission code ensure that sufficient transitions are present in the PHY bit stream to make clock recovery possible at the receiver. Such encoding also greatly increases the likelihood of detecting any single or multiple bit errors that may occur during transmission and reception of information. In addition, some of the special code-groups of the transmission code contain a distinct and easily recognizable bit pattern that assists a receiver in achieving code-group alignment on the incoming PHY bit stream. The 8B/10B transmission code specified for use in this standard has a high transition density, is a run-length-limited code, and is dc-balanced. The transition density of the 8B/10B symbols ranges from 3 to 8 transitions per symbol. The definition of the 8B/10B transmission code in this standard is identical to that specified in ANSI X3.230-1994 (FC-PH), Clause 11. The relationship of code-group bit positions to PMA and other PCS constructs is illustrated in Figure 36-3.
--	--

Exhibit C-11




Figure 36-3—PCS reference diagram

36.2.4.1 Notation conventions

8B/10B transmission code uses letter notation for describing the bits of an unencoded information octet and a single control variable. Each bit of the unencoded information octet contains either a binary zero or a binary one. A control variable, Z, has either the value D or the value K. When the control variable associated with an unencoded information octet contains the value D, the associated encoded code-group is referred to as a data code-group. When the control variable associated with an unencoded information octet contains the value K, the associated encoded code-group is referred to as a special code-group. The bit notation of A,B,C,D,E,F,G,H for an unencoded information octet is used in the description of the 8B/10B transmission code. The bits A,B,C,D,E,F,G,H are translated to bits a,b,c,d,e,i,f,g,h,j of 10-bit transmission code-groups. 8B/10B code-group bit assignments are illustrated in Figure 36-3. Each valid codegroup has been given a name using the following convention: /Dx.y/ for the 256 valid data code-groups, and /Kx.y/ for special control code-groups, where x is the decimal value of bits EDCBA, and y is the decimal value of bits HGF.”

802.3-1998, p927-928.

Exhibit C-11

Table 36-1a—Valid data code-groups				
Code Group Name	Octet Value	Octet Bits HGF EDCBA	Current RD –	Current RD +
			abcdei fghj	abcdei fghj
D0.0	00	000 00000	100111 0100	011000 1011
D1.0	01	000 00001	011101 0100	100010 1011
D2.0	02	000 00010	101101 0100	010010 1011
D3.0	03	000 00011	110001 1011	110001 0100
D4.0	04	000 00100	110101 0100	001010 1011
D5.0	05	000 00101	101001 1011	101001 0100
D6.0	06	000 00110	011001 1011	011001 0100
D7.0	07	000 00111	111000 1011	000111 0100
D8.0	08	000 01000	111001 0100	000110 1011
D9.0	09	000 01001	100101 1011	100101 0100
D10.0	0A	000 01010	010101 1011	010101 0100
D11.0	0B	000 01011	110100 1011	110100 0100
D12.0	0C	000 01100	001101 1011	001101 0100
D13.0	0D	000 01101	101100 1011	101100 0100
D14.0	0E	000 01110	011100 1011	011100 0100
D15.0	0F	000 01111	010111 0100	101000 1011
D16.0	10	000 10000	011011 0100	100100 1011

“

...

Exhibit C-11

Table 36-1e—Valid data code-groups

Code Group Name	Octet Value	Octet Bits HGF EDCBA	Current RD –	Current RD +
			abcdei fghj	abcdei fghj
D16.7	F0	111 10000	011011 0001	100100 1110
D17.7	F1	111 10001	100011 0111	100011 0001
D18.7	F2	111 10010	010011 0111	010011 0001
D19.7	F3	111 10011	110010 1110	110010 0001
D20.7	F4	111 10100	001011 0111	001011 0001
D21.7	F5	111 10101	101010 1110	101010 0001
D22.7	F6	111 10110	011010 1110	011010 0001
D23.7	F7	111 10111	111010 0001	000101 1110
D24.7	F8	111 11000	110011 0001	001100 1110
D25.7	F9	111 11001	100110 1110	100110 0001
D26.7	FA	111 11010	010110 1110	010110 0001
D27.7	FB	111 11011	110110 0001	001001 1110
D28.7	FC	111 11100	001110 1110	001110 0001
D29.7	FD	111 11101	101110 0001	010001 1110
D30.7	FE	111 11110	011110 0001	100001 1110
D31.7	FF	111 11111	101011 0001	010100 1110
(concluded)				

Table 36-2—Valid special code-groups

Code Group Name	Octet Value	Octet Bits HGF EDCBA	Current RD –	Current RD +	Notes
			abcdei fghj	abcdei fghj	
K28.0	1C	000 11100	001111 0100	110000 1011	1
K28.1	3C	001 11100	001111 1001	110000 0110	1,2
K28.2	5C	010 11100	001111 0101	110000 1010	1
K28.3	7C	011 11100	001111 0011	110000 1100	1
K28.4	9C	100 11100	001111 0010	110000 1101	1
K28.5	BC	101 11100	001111 1010	110000 0101	2
K28.6	DC	110 11100	001111 0110	110000 1001	1
K28.7	FC	111 11100	001111 1000	110000 0111	1,2
K23.7	F7	111 10111	111010 1000	000101 0111	
K27.7	FB	111 11011	110110 1000	001001 0111	
K29.7	FD	111 11101	101110 1000	010001 0111	
K30.7	FE	111 11110	011110 1000	100001 0111	
NOTE 1—Reserved.					
NOTE 2—Contains a comma.					

36.2.4.7 Ordered_sets Eight ordered_sets, consisting of a single special code-group or combinations of special and data codegroups are specifically defined. Ordered_sets which include /K28.5/ provide the ability to obtain bit and code-group synchronization and establish ordered_set alignment (see 36.2.4.9 and 36.3.2.4). Ordered_sets provide for the delineation of a packet and

Exhibit C-11

		synchronization between the transmitter and receiver circuits at opposite ends of a link. Table 36-3 lists the defined ordered_sets. Table 36-3—Defined ordered_sets <table><tr><th>Code</th><th>Ordered_Set</th><th>Number of Code-Groups</th><th>Encoding</th></tr><tr><td>/C/</td><td>Configuration</td><td></td><td>Alternating /C1/ and /C2/</td></tr><tr><td>/C1/</td><td>Configuration 1</td><td>4</td><td>/K28.5/D21.5/Config_Reg^a</td></tr><tr><td>/C2/</td><td>Configuration 2</td><td>4</td><td>/K28.5/D2.2/Config_Reg^a</td></tr><tr><td>/I/</td><td>IDLE</td><td></td><td>Correcting /I1/, Preserving /I2/</td></tr><tr><td>/I1/</td><td>IDLE 1</td><td>2</td><td>/K28.5/D5.6/</td></tr><tr><td>/I2/</td><td>IDLE 2</td><td>2</td><td>/K28.5/D16.2/</td></tr><tr><td></td><td>Encapsulation</td><td></td><td></td></tr><tr><td>/R/</td><td>Carrier_Extend</td><td>1</td><td>/K23.7/</td></tr><tr><td>/S/</td><td>Start_of_Packet</td><td>1</td><td>/K27.7/</td></tr><tr><td>/T/</td><td>End_of_Packet</td><td>1</td><td>/K29.7/</td></tr><tr><td>/V/</td><td>Error_Propagation</td><td>1</td><td>/K30.7/</td></tr></table> ^aTwo data code-groups representing the Config_Reg value. ” 802.3-1998, p931-936. To the extent that Plaintiff contends that 802.3-1998 does not explicitly or inherently disclose this element, or render obvious this element, this element is obvious in combination with one or more references and/or with the knowledge of a person of ordinary skill in the art, as set forth in Defendant’s Invalidity Contentions.	Code	Ordered_Set	Number of Code-Groups	Encoding	/C/	Configuration		Alternating /C1/ and /C2/	/C1/	Configuration 1	4	/K28.5/D21.5/Config_Reg ^a	/C2/	Configuration 2	4	/K28.5/D2.2/Config_Reg ^a	/I/	IDLE		Correcting /I1/, Preserving /I2/	/I1/	IDLE 1	2	/K28.5/D5.6/	/I2/	IDLE 2	2	/K28.5/D16.2/		Encapsulation			/R/	Carrier_Extend	1	/K23.7/	/S/	Start_of_Packet	1	/K27.7/	/T/	End_of_Packet	1	/K29.7/	/V/	Error_Propagation	1	/K30.7/
Code	Ordered_Set	Number of Code-Groups	Encoding																																															
/C/	Configuration		Alternating /C1/ and /C2/																																															
/C1/	Configuration 1	4	/K28.5/D21.5/Config_Reg ^a																																															
/C2/	Configuration 2	4	/K28.5/D2.2/Config_Reg ^a																																															
/I/	IDLE		Correcting /I1/, Preserving /I2/																																															
/I1/	IDLE 1	2	/K28.5/D5.6/																																															
/I2/	IDLE 2	2	/K28.5/D16.2/																																															
	Encapsulation																																																	
/R/	Carrier_Extend	1	/K23.7/																																															
/S/	Start_of_Packet	1	/K27.7/																																															
/T/	End_of_Packet	1	/K29.7/																																															
/V/	Error_Propagation	1	/K30.7/																																															
Claim 28																																																		
28[p]	The communications device of claim 19	802.3-1998 discloses, expressly or inherently, the transmitting device of claim 19 under any proposed construction. See exemplary disclosures and corresponding discussion for claim 19.																																																
28	wherein the synchronization primitive is made up of a sequence	802.3-1998 discloses and/or renders obvious this limitation. See, e.g.: “Table 24-1 specifies the interpretation assigned to each five bit code-group, including the mapping to the nibble- wide (TXD or RXD) Data signals on the MII. The 32 code-groups are divided into four categories, as shown. For clarity in the remainder of this clause, code-group names are shown																																																

Exhibit C-11

	of bit-string synchronization symbols.	between /slashes/. Code-group sequences are shown in succession, e.g.: /1/2/.... The indicated code-group mapping is identical to ISO/IEC 9314-1: 1989, with four exceptions: a) The FDDI term symbol is avoided in order to prevent confusion with other 100BASE-T terminology. In general, the term code-group is used in its place. b) The /S/ and /Q/ code-groups are not used by 100BASE-X and are interpreted as INVALID. c) The /R/ code-group is used in 100BASE-X as the second code-group of the End-of-Stream delimiter rather than to indicate a Reset condition. d) The /H/ code-group is used to propagate receive errors rather than to indicate the Halt Line State. 24.2.2.1.1 Data code-groups A Data code-group conveys one nibble of arbitrary data between the MII and the PCS. The sequence of Data code-groups is arbitrary, where any Data code-group can be followed by any other Data code-group. Data code-groups are coded and decoded but not interpreted by the PCS. Successful decoding of Data code-groups depends on proper receipt of the Start-of-Stream delimiter sequence, as defined in Table 24-1. 24.2.2.1.2 Idle code-groups The Idle code-group (/I/) is transferred between streams. It provides a continuous fill pattern to establish and maintain clock synchronization. Idle code-groups are emitted from, and interpreted by, the PCS. 24.2.2.1.3 Control code-groups The Control code-groups are used in pairs (/J/K/, /T/R/) to delimit MAC packets. Control code-groups are emitted from, and interpreted by, the PCS.
--	--	---

Exhibit C-11

Table 24-1—4B/5B code-groups

	PCS code-group [4:0] 4 3 2 1 0	Name	MII (TXD/RXD) <3:0> 3 2 1 0	Interpretation
D A T A	1 1 1 1 0	0	0 0 0 0	Data 0
	0 1 0 0 1	1	0 0 0 1	Data 1
	1 0 1 0 0	2	0 0 1 0	Data 2
	1 0 1 0 1	3	0 0 1 1	Data 3
	0 1 0 1 0	4	0 1 0 0	Data 4
	0 1 0 1 1	5	0 1 0 1	Data 5
	0 1 1 1 0	6	0 1 1 0	Data 6
	0 1 1 1 1	7	0 1 1 1	Data 7
	1 0 0 1 0	8	1 0 0 0	Data 8
	1 0 0 1 1	9	1 0 0 1	Data 9
	1 0 1 1 0	A	1 0 1 0	Data A
	1 0 1 1 1	B	1 0 1 1	Data B
	1 1 0 1 0	C	1 1 0 0	Data C
	1 1 0 1 1	D	1 1 0 1	Data D
	1 1 1 0 0	E	1 1 1 0	Data E
	1 1 1 0 1	F	1 1 1 1	Data F
C O N T R O L	1 1 1 1 1	I	undefined	IDLE; used as inter-stream fill code
	1 1 0 0 0	J	0 1 0 1	Start-of-Stream Delimiter, Part 1 of 2; always used in pairs with K
	1 0 0 0 1	K	0 1 0 1	Start-of-Stream Delimiter, Part 2 of 2; always used in pairs with J
	0 1 1 0 1	T	undefined	End-of-Stream Delimiter, Part 1 of 2; always used in pairs with R
I N V A L I D	0 0 1 1 1	R	undefined	End-of-Stream Delimiter, Part 2 of 2; always used in pairs with T
	0 0 1 0 0	H	Undefined	Transmit Error; used to force signaling errors
	0 0 0 0 0	V	Undefined	Invalid code
	0 0 0 0 1	V	Undefined	Invalid code
	0 0 0 1 0	V	Undefined	Invalid code
	0 0 0 1 1	V	Undefined	Invalid code
	0 0 1 0 1	V	Undefined	Invalid code
	0 0 1 1 0	V	Undefined	Invalid code
	0 1 0 0 0	V	Undefined	Invalid code
	0 1 1 0 0	V	Undefined	Invalid code
	1 0 0 0 0	V	Undefined	Invalid code
	1 1 0 0 1	V	Undefined	Invalid code

802.3-1998, p617-618.

Exhibit C-11

“36.2.4.5 Generating code-groups

The appropriate entry in either Table 36-1a or Table 36-2 is found for each octet for which a code-group is to be generated (encoded). The current value of the transmitter’s running disparity shall be used to select the code-group from its corresponding column. For each code-group transmitted, a new value of the running disparity is calculated. This new value is used as the transmitter’s current running disparity for the next octet to be encoded and transmitted.

Table 36-1a—Valid data code-groups

Code Group Name	Octet Value	Octet Bits HGF EDCBA	Current RD –	Current RD +
			abcdei fghj	abcdei fghj
D0.0	00	000 00000	100111 0100	011000 1011
D1.0	01	000 00001	011101 0100	100010 1011
D2.0	02	000 00010	101101 0100	010010 1011
D3.0	03	000 00011	110001 1011	110001 0100
D4.0	04	000 00100	110101 0100	001010 1011
D5.0	05	000 00101	101001 1011	101001 0100
D6.0	06	000 00110	011001 1011	011001 0100
D7.0	07	000 00111	111000 1011	000111 0100
D8.0	08	000 01000	111001 0100	000110 1011
D9.0	09	000 01001	100101 1011	100101 0100
D10.0	0A	000 01010	010101 1011	010101 0100
D11.0	0B	000 01011	110100 1011	110100 0100
D12.0	0C	000 01100	001101 1011	001101 0100
D13.0	0D	000 01101	101100 1011	101100 0100
D14.0	0E	000 01110	011100 1011	011100 0100
D15.0	0F	000 01111	010111 0100	101000 1011
D16.0	10	000 10000	011011 0100	100100 1011

...

Exhibit C-11

Table 36-1e—Valid data code-groups

Code Group Name	Octet Value	Octet Bits HGF EDCBA	Current RD –	Current RD +
			abcdei fghj	abcdei fghj
D16.7	F0	111 10000	011011 0001	100100 1110
D17.7	F1	111 10001	100011 0111	100011 0001
D18.7	F2	111 10010	010011 0111	010011 0001
D19.7	F3	111 10011	110010 1110	110010 0001
D20.7	F4	111 10100	001011 0111	001011 0001
D21.7	F5	111 10101	101010 1110	101010 0001
D22.7	F6	111 10110	011010 1110	011010 0001
D23.7	F7	111 10111	111010 0001	000101 1110
D24.7	F8	111 11000	110011 0001	001100 1110
D25.7	F9	111 11001	100110 1110	100110 0001
D26.7	FA	111 11010	010110 1110	010110 0001
D27.7	FB	111 11011	110110 0001	001001 1110
D28.7	FC	111 11100	001110 1110	001110 0001
D29.7	FD	111 11101	101110 0001	010001 1110
D30.7	FE	111 11110	011110 0001	100001 1110
D31.7	FF	111 11111	101011 0001	010100 1110
(concluded)				

Table 36-2—Valid special code-groups

Code Group Name	Octet Value	Octet Bits HGF EDCBA	Current RD –	Current RD +	Notes
			abcdei fghj	abcdei fghj	
K28.0	1C	000 11100	001111 0100	110000 1011	1
K28.1	3C	001 11100	001111 1001	110000 0110	1,2
K28.2	5C	010 11100	001111 0101	110000 1010	1
K28.3	7C	011 11100	001111 0011	110000 1100	1
K28.4	9C	100 11100	001111 0010	110000 1101	1
K28.5	BC	101 11100	001111 1010	110000 0101	2
K28.6	DC	110 11100	001111 0110	110000 1001	1
K28.7	FC	111 11100	001111 1000	110000 0111	1,2
K23.7	F7	111 10111	111010 1000	000101 0111	
K27.7	FB	111 11011	110110 1000	001001 0111	
K29.7	FD	111 11101	101110 1000	010001 0111	
K30.7	FE	111 11110	011110 1000	100001 0111	
NOTE 1—Reserved.					
NOTE 2—Contains a comma.					

802.3-1998, p930-935.

To the extent that Plaintiff contends that 802.3-1998 does not explicitly or inherently disclose this element, or render obvious this element, this element is obvious in combination with one or more

Exhibit C-11

		references and/or with the knowledge of a person of ordinary skill in the art, as set forth in Defendant's Invalidity Contentions.
--	--	--