

CLAIMS

What is claimed is:

1. A method implemented by a network node (10) for signaling to a wireless communication device (14) which precoders in a codebook are restricted from being used, the method characterized by:

generating (210) codebook subset restriction signaling that, for each of one or more groups of precoders, jointly restricts the precoders in the group by restricting a certain component that the precoders in the group have in common, and wherein the precoders are Kronecker product precoders and wherein the signaling is rank-agnostic so as to restrict precoders irrespective of their transmission rank; and

sending (220) the generated signaling from the network node (10) to the wireless communication device (14).

2. A method implemented by a wireless communication device (14) for decoding signaling from a network node (10) indicating which precoders in a codebook are restricted from being used, the method characterized by:

receiving (400) codebook subset restriction signaling that, for each of one or more groups of precoders, jointly restricts the precoders in the group by restricting a certain component that the precoders in the group have in common, and wherein the precoders are Kronecker product precoders and wherein the signaling is rank-agnostic so as to restrict precoders irrespective of their transmission rank; and

decoding (410) the received signaling as jointly restricting precoders in each of the one or more groups of precoders.

~~3.~~ The method of any of claims 1-2, wherein the certain component comprises a beam precoder.

~~4.~~ The method of any of claims 1-3, wherein a precoder comprising one or more beam precoders is restricted if at least one of its one or more beam precoders is restricted.

~~5.~~ The method of any of claims 3-4, wherein a beam precoder is a Kronecker product of different beamforming vectors associated with different dimensions of a multi-dimensional antenna array.

Deleted: 3. The method of any of claims 1-2, wherein the codebook subset restriction signaling is rank-agnostic signaling that jointly restricts the precoders in a group without regard to the precoders' transmission rank. ¶ 4

Deleted: 3

Deleted: 5

Deleted: 4

Deleted: 6

Deleted: 4

Deleted: 5

- 6. The method of claim 5, wherein the different beamforming vectors comprise Discrete Fourier Transform (DFT) vectors.
- 7. The method of any of claims 3-6, wherein a beam precoder is a beamforming vector used to transmit on a particular layer of a multi-layer transmission, wherein different scaled versions of that beamforming vector are transmitted on different polarizations;
- 8. The method of any of claims 3-6, wherein a beam precoder is a beamforming vector used to transmit on:
 - multiple different layers of a multi-layer transmission;
 - multiple different layers of a multi-layer transmission, wherein the layers are sent on orthogonal polarizations; or
 - a particular layer and on a particular polarization.
- 9. The method of any of claims 1-8, wherein the codebook subset restriction signaling comprises a bitmap, with different bits in the bitmap respectively dedicated to indicating whether or not different beam precoders are restricted from being used.
- 10. The method of any of claims 1-9, wherein a beam precoder is a Kronecker product of first and second beamforming vectors with first and second indices, wherein the first and second beamforming vectors are associated with different dimensions of a multi-dimensional antenna array, and wherein the codebook subset restriction signaling jointly restricts the precoders in a group of precoders that have the same pair of values for the first and second indices.
- 11. The method of any of claims 1-10, wherein each precoder comprises one or more beam precoders, wherein each beam precoder comprises multiple different components corresponding to different dimensions of a multi-dimensional antenna array, and wherein said certain component comprises a component of a beam precoder.
- 12. The method of any of claims 1-11, wherein the codebook subset restriction signaling jointly restricts the precoders in a group of precoders that transmit at least in part towards a certain angular pointing direction, by restricting a certain component which has that angular pointing direction.
- 13. A network node (10, 600) for signaling to a wireless communication device (14, 800) which precoders in a codebook are restricted from being used, the network node (10, 600) configured to:

Deleted: 7
Deleted: 6
Deleted: 5
Deleted: 4
Deleted: 7
Deleted: 9
Deleted: 4
Deleted: 7
Deleted: 10
Deleted: 9
Deleted: 11
Deleted: 9
Deleted: 12
Deleted: 3
Deleted: 13
Deleted: 12

Deleted: 14. A method implemented by a network node (10) for signaling to a wireless communication device (14) which precoders in a codebook are restricted from being used, the method characterized by ¶ identifying (110) one or more reference configurations for the group, wherein each reference configuration is one of different possible configurations that restrict different subgroups of precoders in the group from being used; ¶ identifying (120), from the different possible configurations for the group, an actual configuration to be signaled for the group; and ¶ generating (130) signaling to indicate the actual configuration for the group, by generating the signaling as a bit pattern whose length depends on (i) whether the actual configuration matches one of the one or more reference configurations and/or (ii) which reference configuration the actual configuration matches; and ¶ sending (160) the generated signaling to the wireless communication device (14). ¶

15. A method implemented by a wireless communication device (14) for decoding signaling from a network node (10) indicating which precoders in a codebook are restricted from being used, the method characterized by ¶ receiving (300) signaling from the network node (10); ¶ for each of one or more groups of precoders in the codebook; ¶ identifying (320) one or more reference configurations for the group, wherein each reference configuration is one of different possible configurations that restrict different subgroups of precoders in the group from being used; ¶ identifying (320) a bit pattern defined for signaling each reference configuration, and a length of that bit pattern; and ¶ detecting (330) an actual configuration signaled for the group, by detecting in the signaling a bit pattern whose length depends on (i) whether the actual configuration matches one of the one or more reference configurations and/or (ii) which reference configuration the actual configuration matches. ¶

16. The method of any of claims 14-15, wherein the signaling is a short bit pattern when the actual configuration matches any one of the one or more reference configurations and is a long bit pattern when the actual configuration does not match any of the one or more reference configurations, wherein a long bit pattern has more bits than a short bit pattern. ¶

17. The method of claim 16, wherein the one or more reference configurations for at least one of the one or more groups comprise a single reference configuration, and wherein different long bit patterns are respectively defined for signaling different configurations other than the single reference configuration. ¶

generate codebook subset restriction signaling that, for each of one or more groups of precoders, jointly restricts the precoders in the group by restricting a certain component that the precoders in the group have in common, and wherein the precoders are Kronecker product precoders and wherein the signaling is rank-agnostic so as to restrict precoders irrespective of their transmission rank; and send the generated signaling from the network node (10, 600) to the wireless communication device (14, 800).

5

14. The network node of claim 13, configured to perform the method of any of claims 3-12.

10

15. A network node (10, 600) for signaling to a wireless communication device (14, 800) which precoders in a codebook are restricted from being used, the network node (10, 600) characterized by:

15

a generating module (640) for generating codebook subset restriction signaling that, for each of one or more groups of precoders, jointly restricts the precoders in the group by restricting a certain component that the precoders in the group have in common, and wherein the precoders are Kronecker product precoders and wherein the signaling is rank-agnostic so as to restrict precoders irrespective of their transmission rank; and

20

a sending module (650) for sending the generated signaling from the network node (10, 600) to the wireless communication device (14, 800).

16. A wireless communication device (14, 800) for decoding signaling from a network node (10, 600) indicating which precoders in a codebook are restricted from being used, the wireless communication device (14, 800) configured to:

25

receive codebook subset restriction signaling that, for each of one or more groups of precoders, jointly restricts the precoders in the group by restricting a certain component that the precoders in the group have in common, and wherein the precoders are Kronecker product precoders and wherein the signaling is rank-agnostic so as to restrict precoders irrespective of their transmission rank; and decode the received signaling as jointly restricting precoders in each of the one or more groups of precoders.

30

17. The wireless communication device of claim 16, configured to perform the method of any of claims 3-12.

35

Deleted: 25

Deleted: 24

Deleted: 13

Deleted: 26

Deleted: 27

Deleted: 28

Deleted: 27

Deleted: 13

18. A wireless communication device (14, 800) for decoding signaling from a network node (10, 600) indicating which precoders in a codebook are restricted from being used, the wireless communication device (14, 800) characterized by:

- 5 a receiving module (840) for receiving codebook subset restriction signaling that, for each of one or more groups of precoders, jointly restricts the precoders in the group by restricting a certain component that the precoders in the group have in common, and wherein the precoders are Kronecker product precoders and wherein the signaling is rank-agnostic so as to restrict precoders irrespective of their transmission rank; and
- 10 a decoding module (850) for decoding the received signaling as jointly restricting precoders in each of the one or more groups of precoders.

19. A computer program comprising instructions which, when executed by at least one processor of a node (10, 14), causes the node (10, 14) to carry out the method of any of embodiments 1-12.

20. A carrier containing the computer program of embodiment 19, wherein the carrier is one of an electronic signal, optical signal, radio signal, or computer readable storage medium.

20

Deleted: 29

Deleted: ¶

30. A network node (10, 500) for signaling to a wireless communication device which precoders in a codebook are restricted from being used, the network node (10, 500) configured to ¶
 for each of one or more groups of precoders in the codebook, ¶
 identify one or more reference configurations for the group, wherein each reference configuration is one of different possible configurations that restrict different subgroups of precoders in the group from being used, ¶
 identify, from the different possible configurations for the group, an actual configuration to be signaled for the group, and ¶
 generate signaling to indicate the actual configuration for the group, by generating the signaling as a bit pattern whose length depends on (i) whether the actual configuration matches one of the one or more reference configurations and/or (ii) which reference configuration the actual configuration matches, and ¶
 send the generated signaling to the wireless communication device. ¶

31. The network node of claim 30, configured to perform the method of any of claims 16-23 ¶

32. A network node (10, 600) for signaling to a wireless communication device (14, 700) which precoders in a codebook are restricted from being used, the network node (10, 600) characterized by, ¶
 for each of one or more groups of precoders in the codebook, ¶
 a reference configuration identifying module (540) for identifying one or more reference configurations for the group, wherein each reference configuration is one of different possible configurations that restrict different subgroups of precoders in the group from being used, ¶
 an actual configuration identifying module (550) for identifying, from the different possible configurations for the group, an actual configuration to be signaled for the group, and ¶
 a generating module (560) for generating signaling to indicate the actual configuration for the group, by generating the signaling as a bit pattern whose length depends on (i) whether the actual configuration matches one of the one or more reference configurations and/or (ii) which reference configuration the actual configuration matches, and ¶
 a sending module (570) for sending the generated signaling to the wireless communication device (14, 700). ¶

33. The network node of claim 32, configured to perform the method of any of claims 16-23 ¶

34. A wireless communication device (14, 700) for decoding signaling from a network node (10, 500) indicating which precoders in a codebook are restricted from being used, the wireless communication device (14, 700) configured to, ¶
 receive signaling from the network node (10, 500), ¶
 for each of one or more groups of precoders in the codebook, ¶
 identify one or more reference configurations for the group, wherein each reference configuration is one of different possible configurations that restrict different subgroups of precoders in the group from being used, ¶
 identify a bit pattern defined for signaling each reference configuration, and a length of that bit pattern, and ¶
 detect an actual configuration signaled for the group, by detecting in the signaling a bit pattern whose length depends on (i) whether the actual configuration matches one of the one or more reference configurations and/or (ii) which reference configuration the actual configuration matches. ¶

35. The wireless communication device of claim 34, configured to perform the method of any of claims 16-23 ¶

Deleted: 37

Deleted: 23

Deleted: 28

Deleted: 37

P45698-WO1 / 4015-9362

CLAIMS

What is claimed is:

1. A method implemented by a network node (10) for signaling to a wireless communication device (14) which precoders in a codebook are restricted from being used, the method characterized by:
- 5 generating (210) codebook subset restriction signaling that, for each of one or more groups of precoders, jointly restricts the precoders in the group by restricting a certain component that the precoders in the group have in common, and wherein
- 10 the precoders are Kronecker product precoders and wherein the signaling is rank-agnostic so as to restrict precoders irrespective of their transmission rank; and
- sending (220) the generated signaling from the network node (10) to the wireless communication device (14).
- 15
2. A method implemented by a wireless communication device (14) for decoding signaling from a network node (10) indicating which precoders in a codebook are restricted from being used, the method characterized by:
- 20 receiving (400) codebook subset restriction signaling that, for each of one or more groups of precoders, jointly restricts the precoders in the group by restricting a certain component that the precoders in the group have in common, and wherein the precoders are Kronecker product precoders and wherein the signaling is rank-agnostic so as to restrict precoders irrespective of their transmission rank; and
- 25 decoding (410) the received signaling as jointly restricting precoders in each of the one or more groups of precoders.
3. The method of any of claims 1-2, wherein the certain component comprises a beam precoder.
- 30
4. The method of any of claims 1-3, wherein a precoder comprising one or more beam precoders is restricted if at least one of its one or more beam precoders is restricted.
5. The method of any of claims 3-4, wherein a beam precoder is a Kronecker product of
- 35 different beamforming vectors associated with different dimensions of a multi-dimensional antenna array.

P45698-WO1 / 4015-9362

6. The method of claim 5, wherein the different beamforming vectors comprise Discrete Fourier Transform (DFT) vectors.
7. The method of any of claims 3-6, wherein a beam precoder is a beamforming vector used to transmit on a particular layer of a multi-layer transmission, wherein different scaled versions of that beamforming vector are transmitted on different polarizations;
8. The method of any of claims 3-6, wherein a beam precoder is a beamforming vector used to transmit on:
- multiple different layers of a multi-layer transmission;
- multiple different layers of a multi-layer transmission, wherein the layers are sent on orthogonal polarizations; or
- a particular layer and on a particular polarization.
9. The method of any of claims 1-8, wherein the codebook subset restriction signaling comprises a bitmap, with different bits in the bitmap respectively dedicated to indicating whether or not different beam precoders are restricted from being used.
10. The method of any of claims 1-8, wherein a beam precoder is a Kronecker product of first and second beamforming vectors with first and second indices, wherein the first and second beamforming vectors are associated with different dimensions of a multi-dimensional antenna array, and wherein the codebook subset restriction signaling jointly restricts the precoders in a group of precoders that have the same pair of values for the first and second indices.
11. The method of any of claims 1-2, wherein each precoder comprises one or more beam precoders, wherein each beam precoder comprises multiple different components corresponding to different dimensions of a multi-dimensional antenna array, and wherein said certain component comprises a component of a beam precoder.
12. The method of any of claims 1-11, wherein the codebook subset restriction signaling jointly restricts the precoders in a group of precoders that transmit at least in part towards a certain angular pointing direction, by restricting a certain component which has that angular pointing direction.
13. A network node (10, 600) for signaling to a wireless communication device (14, 800) which precoders in a codebook are restricted from being used, the network node (10, 600) configured to:

P45698-WO1 / 4015-9362

5 generate codebook subset restriction signaling that, for each of one or more groups of precoders, jointly restricts the precoders in the group by restricting a certain component that the precoders in the group have in common, and wherein the precoders are Kronecker product precoders and wherein the signaling is rank-agnostic so as to restrict precoders irrespective of their transmission rank; and send the generated signaling from the network node (10, 600) to the wireless communication device (14, 800).

10 14. The network node of claim 13, configured to perform the method of any of claims 3-12.

15 15. A network node (10, 600) for signaling to a wireless communication device (14, 800) which precoders in a codebook are restricted from being used, the network node (10, 600) characterized by:

20 a generating module (640) for generating codebook subset restriction signaling that, for each of one or more groups of precoders, jointly restricts the precoders in the group by restricting a certain component that the precoders in the group have in common, and wherein the precoders are Kronecker product precoders and wherein the signaling is rank-agnostic so as to restrict precoders irrespective of their transmission rank; and

25 a sending module (650) for sending the generated signaling from the network node (10, 600) to the wireless communication device (14, 800).

30 16. A wireless communication device (14, 800) for decoding signaling from a network node (10, 600) indicating which precoders in a codebook are restricted from being used, the wireless communication device (14, 800) configured to:

35 receive codebook subset restriction signaling that, for each of one or more groups of precoders, jointly restricts the precoders in the group by restricting a certain component that the precoders in the group have in common, and wherein the precoders are Kronecker product precoders and wherein the signaling is rank-agnostic so as to restrict precoders irrespective of their transmission rank; and decode the received signaling as jointly restricting precoders in each of the one or more groups of precoders.

17. The wireless communication device of claim 16, configured to perform the method of any of claims 3-12.

P45698-WO1 / 4015-9362

18. A wireless communication device (14, 800) for decoding signaling from a network node (10, 600) indicating which precoders in a codebook are restricted from being used, the wireless communication device (14, 800) characterized by:

5 a receiving module (840) for receiving codebook subset restriction signaling that, for each of one or more groups of precoders, jointly restricts the precoders in the group by restricting a certain component that the precoders in the group have in common, and wherein the precoders are Kronecker product precoders and wherein the signaling is rank-agnostic so as to restrict precoders irrespective of their transmission rank; and

10 a decoding module (850) for decoding the received signaling as jointly restricting precoders in each of the one or more groups of precoders.

19. A computer program comprising instructions which, when executed by at least one processor of a node (10, 14), causes the node (10, 14) to carry out the method of any of embodiments 1-12.

20. A carrier containing the computer program of embodiment 19, wherein the carrier is one of an electronic signal, optical signal, radio signal, or computer readable storage medium.

20

Document code: WFEE

United States Patent and Trademark Office
Sales Receipt for Accounting Date: 01/31/2018

CHAYWOOD	SALE	#00000003	Mailroom Dt:	06/17/2016	15105648
		01	FC : 1642	480.00	OP

Document code: WFEE

United States Patent and Trademark Office
Sales Receipt for Accounting Date: 01/31/2018

CHAYWOOD	ADJ #00000003	Mailroom Dt: 06/17/2016	
	Seq No: 7992	Sales Acctg Dt: 06/17/2016	15105648
	02 FC : 1632	-600.00	OP

Doc code: IDS

Doc description: Information Disclosure Statement (IDS) Filed

PTO/SB/08a (03-15)

Approved for use through 07/31/2016. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)	Application Number	15105648
	Filing Date	2016-06-17
	First Named Inventor	Faxér
	Art Unit	
	Examiner Name	
	Attorney Docket Number	4015-9595 / P45698-US2

U.S.PATENTS							Remove
Examiner Initial*	Cite No	Patent Number	Kind Code ¹	Issue Date	Name of Patentee or Applicant of cited Document	Pages, Columns, Lines where Relevant Passages or Relevant Figures Appear	
	1						

If you wish to add additional U.S. Patent citation information please click the Add button.

Add

U.S.PATENT APPLICATION PUBLICATIONS							Remove
Examiner Initial*	Cite No	Publication Number	Kind Code ¹	Publication Date	Name of Patentee or Applicant of cited Document	Pages, Columns, Lines where Relevant Passages or Relevant Figures Appear	
	1	20130163687	A1	2013-06-27	Jing et al.		
	2	20140016549	A1	2014-01-16	Novlan et al.		
	3	20130229980	A1	2013-09-05	Wernersson et al.		
	4	20100223237	A1	2010-09-02	Mishra et al.		

If you wish to add additional U.S. Published Application citation information please click the Add button.

Add

FOREIGN PATENT DOCUMENTS								Remove
Examiner Initial*	Cite No	Foreign Document Number ³	Country Code ²ⁱ	Kind Code ⁴	Publication Date	Name of Patentee or Applicant of cited Document	Pages, Columns, Lines where Relevant Passages or Relevant Figures Appear	T ⁵

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)	Application Number		15105648	
	Filing Date		2016-06-17	
	First Named Inventor	Faxér		
	Art Unit			
	Examiner Name			
	Attorney Docket Number		4015-9595 / P45698-US2	

	1						
--	---	--	--	--	--	--	--

If you wish to add additional Foreign Patent Document citation information please click the Add button

NON-PATENT LITERATURE DOCUMENTS

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T ⁵
	1	ALAMELDEEN, A., et al., "Frequent Pattern Compression: A Significance-Based Compression Scheme for L2 Caches", Technical Report #1500, 2004-05-01, pp. 1-15, University of Wisconsin	
	2	THOMAS, M. et al., "Elements of Information Theory", Chapter 3, Asymptotic Equipartition Property", 2006-01-01, pp. 57-62, Second edition, John Wiley & Sons, Inc.	

If you wish to add additional non-patent literature document citation information please click the Add button

EXAMINER SIGNATURE

Examiner Signature		Date Considered	
--------------------	--	-----------------	--

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

¹ See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. ² Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). ³ For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. ⁴ Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. ⁵ Applicant is to place a check mark here if English language translation is attached.

**INFORMATION DISCLOSURE
STATEMENT BY APPLICANT**
(Not for submission under 37 CFR 1.99)

Application Number	15105648		
Filing Date	2016-06-17		
First Named Inventor	Faxér		
Art Unit			
Examiner Name			
Attorney Docket Number	4015-9595 / P45698-US2		

CERTIFICATION STATEMENT

Please see 37 CFR 1.97 and 1.98 to make the appropriate selection(s):

That each item of information contained in the information disclosure statement was first cited in any communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(1).

OR

That no item of information contained in the information disclosure statement was cited in a communication from a foreign patent office in a counterpart foreign application, and, to the knowledge of the person signing the certification after making reasonable inquiry, no item of information contained in the information disclosure statement was known to any individual designated in 37 CFR 1.56(c) more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(2).

See attached certification statement.

The fee set forth in 37 CFR 1.17 (p) has been submitted herewith.

A certification statement is not submitted herewith.

SIGNATURE

A signature of the applicant or representative is required in accordance with CFR 1.33, 10.18. Please see CFR 1.4(d) for the form of the signature.

Signature	/Justin J. Leonard, Reg. No. 60986/	Date (YYYY-MM-DD)	2016-08-02
Name/Print	Justin Leonard	Registration Number	60986

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 1 hour to complete, including gathering, preparing and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. **DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these records.
2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.
3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.
4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).
5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.
6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).
7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.
8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.
9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

ASYMPTOTIC EQUIPARTITION PROPERTY

In information theory, the analog of the law of large numbers is the asymptotic equipartition property (AEP). It is a direct consequence of the weak law of large numbers. The *law of large numbers* states that for independent, identically distributed (i.i.d.) random variables, $\frac{1}{n} \sum_{i=1}^n X_i$ is close to its expected value EX for large values of n . The AEP states that $\frac{1}{n} \log \frac{1}{p(X_1, X_2, \dots, X_n)}$ is close to the entropy H , where X_1, X_2, \dots, X_n are i.i.d. random variables and $p(X_1, X_2, \dots, X_n)$ is the probability of observing the sequence X_1, X_2, \dots, X_n . Thus, the probability $p(X_1, X_2, \dots, X_n)$ assigned to an observed sequence will be close to 2^{-nH} .

This enables us to divide the set of all sequences into two sets, the *typical set*, where the sample entropy is close to the true entropy, and the *nontypical set*, which contains the other sequences. Most of our attention will be on the typical sequences. Any property that is proved for the typical sequences will then be true with high probability and will determine the average behavior of a large sample.

First, an example. Let the random variable $X \in \{0, 1\}$ have a probability mass function defined by $p(1) = p$ and $p(0) = q$. If X_1, X_2, \dots, X_n are i.i.d. according to $p(x)$, the probability of a sequence x_1, x_2, \dots, x_n is $\prod_{i=1}^n p(x_i)$. For example, the probability of the sequence $(1, 0, 1, 1, 0, 1)$ is $p^4 q^2$. Clearly, it is not true that all 2^n sequences of length n have the same probability.

However, we might be able to predict the probability of the sequence that we actually observe. We ask for the probability $p(X_1, X_2, \dots, X_n)$ of the outcomes X_1, X_2, \dots, X_n , where X_1, X_2, \dots are i.i.d. $\sim p(x)$. This is insidiously self-referential, but well defined nonetheless. Apparently, we are asking for the probability of an event drawn according to the same

Elements of Information Theory, Second Edition. By Thomas M. Cover and Joy A. Thomas
Copyright © 2006 John Wiley & Sons, Inc.

probability distribution. Here it turns out that $p(X_1, X_2, \dots, X_n)$ is close to 2^{-nH} with high probability.

We summarize this by saying, "Almost all events are almost equally surprising." This is a way of saying that

$$\Pr\{(X_1, X_2, \dots, X_n) : p(X_1, X_2, \dots, X_n) = 2^{-n(H \pm \epsilon)}\} \approx 1 \quad (3.1)$$

if X_1, X_2, \dots, X_n are i.i.d. $\sim p(x)$.

In the example just given, where $p(X_1, X_2, \dots, X_n) = p^{\sum x_i} q^{n - \sum x_i}$, we are simply saying that the number of 1's in the sequence is close to np (with high probability), and all such sequences have (roughly) the same probability $2^{-nH(p)}$. We use the idea of convergence in probability, defined as follows:

Definition (Convergence of random variables). Given a sequence of random variables, X_1, X_2, \dots , we say that the sequence X_1, X_2, \dots converges to a random variable X :

1. In probability if for every $\epsilon > 0$, $\Pr\{|X_n - X| > \epsilon\} \rightarrow 0$
2. In mean square if $E(X_n - X)^2 \rightarrow 0$
3. With probability 1 (also called *almost surely*) if $\Pr\{\lim_{n \rightarrow \infty} X_n = X\} = 1$

3.1 ASYMPTOTIC EQUIPARTITION PROPERTY THEOREM

The asymptotic equipartition property is formalized in the following theorem.

Theorem 3.1.1 (AEP). If X_1, X_2, \dots are i.i.d. $\sim p(x)$, then

$$-\frac{1}{n} \log p(X_1, X_2, \dots, X_n) \rightarrow H(X) \quad \text{in probability.} \quad (3.2)$$

Proof: Functions of independent random variables are also independent random variables. Thus, since the X_i are i.i.d., so are $\log p(X_i)$. Hence, by the weak law of large numbers,

$$-\frac{1}{n} \log p(X_1, X_2, \dots, X_n) = -\frac{1}{n} \sum_i \log p(X_i) \quad (3.3)$$

$$\rightarrow -E \log p(X) \quad \text{in probability} \quad (3.4)$$

$$= H(X), \quad (3.5)$$

which proves the theorem. \square

Definition The typical set $A_\epsilon^{(n)}$ with respect to $p(x)$ is the set of sequences $(x_1, x_2, \dots, x_n) \in \mathcal{X}^n$ with the property

$$2^{-n(H(X)+\epsilon)} \leq p(x_1, x_2, \dots, x_n) \leq 2^{-n(H(X)-\epsilon)}, \quad (3.6)$$

As a consequence of the AEP, we can show that the set $A_\epsilon^{(n)}$ has the following properties:

Theorem 3.1.2

1. If $(x_1, x_2, \dots, x_n) \in A_\epsilon^{(n)}$, then $H(X) - \epsilon \leq -\frac{1}{n} \log p(x_1, x_2, \dots, x_n) \leq H(X) + \epsilon$.
2. $\Pr\{A_\epsilon^{(n)}\} > 1 - \epsilon$ for n sufficiently large.
3. $|A_\epsilon^{(n)}| \leq 2^{n(H(X)+\epsilon)}$, where $|A|$ denotes the number of elements in the set A .
4. $|A_\epsilon^{(n)}| \geq (1 - \epsilon)2^{n(H(X)-\epsilon)}$ for n sufficiently large.

Thus, the typical set has probability nearly 1, all elements of the typical set are nearly equiprobable, and the number of elements in the typical set is nearly 2^{nH} .

Proof: The proof of property (1) is immediate from the definition of $A_\epsilon^{(n)}$. The second property follows directly from Theorem 3.1.1, since the probability of the event $(X_1, X_2, \dots, X_n) \in A_\epsilon^{(n)}$ tends to 1 as $n \rightarrow \infty$. Thus, for any $\delta > 0$, there exists an n_0 such that for all $n \geq n_0$, we have

$$\Pr\left\{\left|-\frac{1}{n} \log p(X_1, X_2, \dots, X_n) - H(X)\right| < \epsilon\right\} > 1 - \delta. \quad (3.7)$$

Setting $\delta = \epsilon$, we obtain the second part of the theorem. The identification of $\delta = \epsilon$ will conveniently simplify notation later.

To prove property (3), we write

$$1 = \sum_{\mathbf{x} \in \mathcal{X}^n} p(\mathbf{x}) \quad (3.8)$$

$$\geq \sum_{\mathbf{x} \in A_\epsilon^{(n)}} p(\mathbf{x}) \quad (3.9)$$

$$\geq \sum_{\mathbf{x} \in A_\epsilon^{(n)}} 2^{-n(H(X)+\epsilon)} \quad (3.10)$$

$$= 2^{-n(H(X)+\epsilon)} |A_\epsilon^{(n)}|. \quad (3.11)$$

Copyright © 2002, Apple Computer, Inc. All rights reserved. See also Apple Computer, Inc. (2002) "The AEP and the Law of Large Numbers," Apple Computer, Inc., Cupertino, CA.

where the second inequality follows from (3.6). Hence

$$|A_\epsilon^{(n)}| \leq 2^{n(H(X)+\epsilon)} \tag{3.12}$$

Finally, for sufficiently large n , $\Pr\{A_\epsilon^{(n)}\} > 1 - \epsilon$, so that

$$1 - \epsilon < \Pr\{A_\epsilon^{(n)}\} \tag{3.13}$$

$$\leq \sum_{x \in A_\epsilon^{(n)}} 2^{-n(H(X)-\epsilon)} \tag{3.14}$$

$$\equiv 2^{-n(H(X)-\epsilon)} |A_\epsilon^{(n)}|, \tag{3.15}$$

where the second inequality follows from (3.6). Hence,

$$|A_\epsilon^{(n)}| \geq (1 - \epsilon) 2^{n(H(X)-\epsilon)}, \tag{3.16}$$

which completes the proof of the properties of $A_\epsilon^{(n)}$. □

3.2 CONSEQUENCES OF THE AEP: DATA COMPRESSION

Let X_1, X_2, \dots, X_n be independent, identically distributed random variables drawn from the probability mass function $p(x)$. We wish to find short descriptions for such sequences of random variables. We divide all sequences in \mathcal{X}^n into two sets: the typical set $A_\epsilon^{(n)}$ and its complement, as shown in Figure 3.1.




FIGURE 3.1. Typical sets and source coding.

$$= \sum_{x^n \in A_\epsilon^{(n)}} p(x^n)l(x^n) + \sum_{x^n \in A_\epsilon^{(n)c}} p(x^n)l(x^n) \quad (3.18)$$

$$\leq \sum_{x^n \in A_\epsilon^{(n)}} p(x^n)(n(H + \epsilon) + 2) + \sum_{x^n \in A_\epsilon^{(n)c}} p(x^n)(n \log |\mathcal{X}| + 2) \quad (3.19)$$

$$= \Pr\{A_\epsilon^{(n)}\} (n(H + \epsilon) + 2) + \Pr\{A_\epsilon^{(n)c}\} (n \log |\mathcal{X}| + 2) \quad (3.20)$$

$$\leq n(H + \epsilon) + \epsilon n(\log |\mathcal{X}|) + 2 \quad (3.21)$$

$$= n(H + \epsilon'), \quad (3.22)$$

where $\epsilon' = \epsilon + \epsilon \log |\mathcal{X}| + \frac{2}{n}$ can be made arbitrarily small by an appropriate choice of ϵ followed by an appropriate choice of n . Hence we have proved the following theorem.

Theorem 3.2.1 *Let X^n be i.i.d. $\sim p(x)$. Let $\epsilon > 0$. Then there exists a code that maps sequences x^n of length n into binary strings such that the mapping is one-to-one (and therefore invertible) and*

$$E\left[\frac{1}{n}l(X^n)\right] \leq H(X) + \epsilon \quad (3.23)$$

for n sufficiently large.

Thus, we can represent sequences X^n using $nH(X)$ bits on the average.

3.3 HIGH-PROBABILITY SETS AND THE TYPICAL SET

From the definition of $A_\epsilon^{(n)}$, it is clear that $A_\epsilon^{(n)}$ is a fairly small set that contains most of the probability. But from the definition, it is not clear whether it is the smallest such set. We will prove that the typical set has essentially the same number of elements as the smallest set, to first order in the exponent.

Definition For each $n = 1, 2, \dots$, let $B_\delta^{(n)} \subset \mathcal{X}^n$ be the smallest set with

$$\Pr\{B_\delta^{(n)}\} \geq 1 - \delta. \quad (3.24)$$

Frequent Pattern Compression: A Significance-Based Compression Scheme for L2 Caches

Alaa R. Alamefdeen and David A. Wood
 Computer Sciences Department, University of Wisconsin-Madison
 {alaa, david}@cs.wisc.edu

Abstract

With the widening gap between processor and memory speeds, memory system designers may find cache compression beneficial to increase cache capacity and reduce off-chip bandwidth. Most hardware compression algorithms fall into the dictionary-based category, which depend on building a dictionary and using its entries to encode repeated data values. Such algorithms are effective in compressing large data blocks and files. Cache lines, however, are typically short (32-256 bytes), and a per-line dictionary places a significant overhead that limits the compressibility and increases decompression latency of such algorithms. For such short lines, significance-based compression is an appealing alternative.

We propose and evaluate a simple significance-based compression scheme that has a low compression and decompression overhead. This scheme, Frequent Pattern Compression (FPC) compresses individual cache lines on a word-by-word basis by storing common word patterns in a compressed format accompanied with an appropriate prefix. For a 64-byte cache line, compression can be completed in three cycles and decompression in five cycles, assuming 12 FO4 gate delays per cycle. We propose a compressed cache design in which data is stored in a compressed form in the L2 caches, but are uncompressed in the L1 caches. L2 cache lines are compressed to predetermined sizes that never exceed their original size to reduce decompression overhead. This simple scheme provides comparable compression ratios to more complex schemes that have higher cache hit latencies.

1 Introduction

As semiconductor technology continues to improve, the rising disparity between processor and memory speed increasingly dominates performance. Modern processors use two or more levels of cache memories to reduce effective memory latency and bandwidth. Effectively using the limited on-chip cache resources becomes increasingly important as memory latencies continue to increase relative to processor speeds. Cache compression has previously been proposed to improve performance, since compressing data stored in on-chip caches increases their effective capacity, potentially reducing misses.

Most previous proposals in hardware cache or memory compression (Section 2) are hardware implementations of dictionary-based software compression algorithms (e.g., LZ77 [32]). Such hardware dictionary-based schemes depend mainly on maintaining a per-block dictionary and encoding words (or bytes) that match in the dictionary, while keeping words (bytes) that do not match in their original form with an appropriate prefix.

Schemes such as the Block-Referential Compression with Lookahead (BRCL) used in the IBM MXT memory compression depend on having long enough lines / pages to increase the overall compression ratio [14]. BRCL provides a

good compression ratio for 1K-byte or longer blocks. However, cache lines are typically much shorter and BRCL does not perform as well for shorter lines. In addition, decompression latency is high, since the parallel implementation of BRCL decompresses data at a speed of 8 bytes per cycle [26], or 8 cycles for a 64-byte cache line. The X-Match compression scheme [18] tries to compress more data with a small dictionary by allowing partial matches of data words to dictionary entries. Frequent-value cache designs [29, 31] achieve better compression for cache lines by constructing a single dictionary (the Frequent-Value Cache, FVC) for the whole cache, which increases the chance of a single word to be found and compressed. These designs are based on the observation that a few cache values are frequent and thus can be compressed to a fewer number of bits. However, a large FVC requires an increased decompression latency due to the increased FVC access time.

Significance-based compression is based on the observation that most data types (e.g., 32-bit integers) can be stored in a fewer number of bits than the maximum allowed. For example, sign-bit extension is a commonly implemented technique to store small integers (e.g., 8-bit) into 32-bit words, while all the information in the word is stored in the least-significant few bits. In contrast with dictionary-based compression schemes, significance-based compression [9, 11, 12] does not incur a per-line dictionary overhead, which makes it more suitable for the typically-short cache lines. In addition, compression and decompression hardware is faster than dictionary-based encoding and decoding. However, compressibility can be significantly impaired for long cache lines.

In this document, we propose a significance-based compression scheme that provides reasonable compressibility for the typically short cache lines with a relatively fast hardware compression and decompression (Section 3). This scheme, the Frequent Pattern Compression (FPC) compresses a cache line on a word-by-word basis. For each word, FPC detects whether it falls into one of the patterns that can be stored in a smaller number of bits, and stores it in a compressed form with an appropriate prefix. We discuss the implementation of a hardware decompression pipeline that decompresses a 64-byte cache line in five cycles (Section 4). We evaluate this scheme and compare it with other hardware compression schemes in Section 5.

2 Related Work

Several researchers used hardware-based compression to increase effective memory size, reduce memory address and data bandwidth, and increase effective cache size.

IBM's Memory Compression. IBM's MXT technology [26] employs real-time main-memory content compression that can be used to effectively double the main memory capacity without a significant added cost. It was first implemented in the Pinnacle chip [25], a single-chip memory controller. Franaszek, et al. [13], described the design of a compressed random access memory (C-RAM), which formed the basis for the memory organization for the MXT technology, and studied the optimal line size for such an organization. Data in main memory is compressed using a hardware parallelized derivative of the Lempel-Ziv (LZ77) sequential algorithm [32]. This parallel algorithm, Parallel Block-Referential Compression with Directory Sharing, divides the input data block (1 KB in MXT) into sub-blocks (four 256-byte sub-blocks), and cooperatively constructs dictionaries while compressing all sub-blocks in parallel [14]. MXT is shown to have a negligible performance penalty compared to standard memory, and memory contents for many applications and web servers can be compressed by a factor of two to one [1].

Other Hardware Memory Compression Designs. Kjelson, et al. [18], demonstrated that hardware main memory compression is feasible and worthwhile. They used the X-Match hardware compression algorithm that maintains a dictionary and replaces each input data element (whose size is fixed at four bytes) with a shorter code in case of a total or partial match with a dictionary entry. Communication bandwidth is reduced by "compacting" cache-to-memory address streams [12] or data streams [11]. Benini, et al. [8], propose a data compression/decompression scheme to reduce memory traffic in general purpose processor systems. Data is stored uncompressed in the cache, and compressed on the fly when transferred to memory. Memory-to-cache traffic is also decompressed on the fly. They used a differential compression scheme described in [7] that is based on the assumption that it is likely for data words in the same cache line to have some bits in common. Zhang and Gupta [30] introduce a class of common-prefix and narrow-data transformations for general-purpose programs that compress 32-bit addresses and integer words into 15-bit entities. They implemented these transformations by augmenting six data compression extension (DCX) instructions to the MIPS instruction set.

Cache Compression and Related Designs. Lee, et al. [21, 19, 20], propose a compressed memory hierarchy model that selectively compresses L2 cache and memory blocks if they can be reduced to half their original size. Their selective compressed memory system (SCMS) use a hardware implementation of the X-RL compression algorithm [18], a variant of the X-Match algorithm that gives a special treatment for runs of zeros. They propose several techniques to hide decompression overhead, including parallel decompression, selective adaptive compression for blocks that can be compressed to below a certain threshold, and the use of a decompression buffer to be accessed on L1 misses in parallel with L2 access. Ahn, et al. [2], propose several improvements on the X-RL technique that capture common values. Chen, et al. [10], propose a scheme that dynamically partitions the cache into sections of different compressibility, and they use a variant of the LZ compression algorithm. Pomerene, et al. [22], used a shadow directory scheme with more address tags than data blocks to improve upon LRU replacement.

Frequent-Value-Based Compression. Yang and Gupta [28] found out from an analysis of the SPECint95 benchmarks that a small number of distinct values occupy a large fraction of memory access values. This value locality phenomenon enabled them to design energy-efficient caches [27] and data compressed caches [29]. In their compressed cache design, each line in the L1 cache can be either one uncompressed line or two lines compressed to at least half their original sizes based on frequent values [29]. Zhang, et al., designed a value-centric data cache design called the frequent value cache (FVC) [31], which is a small direct-mapped cache dedicated to holding frequent benchmark values. They showed that augmenting a direct mapped cache with a small frequent value cache can greatly reduce the cache miss rate.

Significance-Based Compression. Farrens and Park [12] make use of the fact that many address references transferred between processor and memory have redundant information in their high-order (most significant) portions. They cached these high order bits in a group of dynamically allocated base registers and only transferred small register indexes rather than the high-order address bits between the processor and memory. Citron and Rudolph [11] store common high-order bits in address and data words in a table and transfer only an index plus the low order bits between the processor and memory. Canal, et al. [9], proposed a scheme that compresses data, addresses and instruc-

tions into their significant bytes while maintaining a two or three extension bits to maintain significant byte positions. They use this method to reduce dynamic power consumption in a processor pipeline. Kant and Iyer [16] studied the compressibility properties of address and data transfers in commercial workloads, and report that the high-order bits can be predicted with high accuracy in address transfers but with less accuracy for data transfers.

3 Frequent Pattern Compression (FPC)

We propose a compression scheme that builds on significance-based compression schemes [9, 11, 12]. It is also based on the observation that some data patterns are frequent and also compressible to a fewer number of bits. For example, many small-value integers can be stored in 4, 8 or 16 bits, but are normally stored in a full 32-bit word (or 64-bits for 64-bit architectures). These values are frequent enough to merit special treatment, and storing them in a more compact form can increase the cache capacity. In addition, special treatment is also given to runs of zeros since they are very frequent, which is similar to the special treatment in X-RL [18]. The insight behind FPC is that we want to get most of the benefits of dictionary-based schemes, while keeping the per-line overhead at a minimum.

The Frequent Pattern Compression (FPC) compresses / decompresses on a cache line basis. Each cache line is divided into 32-bit words (e.g., 16 words for a 64-byte line). Each 32-bit word is encoded as a 3-bit prefix plus data. Table 1 shows the different patterns corresponding to each prefix.

Each word in the cache line is encoded into a compressed format if it matches any of the patterns in the first six rows of Table 1. These patterns are: a zero run (one or more all-zero words), 4-bit sign-extended (including one-word zero runs), one byte sign-extended, one halfword sign-extended, one halfword padded with a zero halfword, two byte-sign-extended halfwords, and a word consisting of repeated bytes (e.g. "0x20202020", or similar patterns that can be used for data initialization). These patterns are selected based on their high frequency in many of our integer and commercial benchmarks. A word that doesn't match any of these categories is stored in its original 32-bit format. All prefix values as well as the zero-run length data bits are stored at the beginning of the line to speed up decompression.

3.1 Segmented Frequent Pattern Compression (S-FPC)

To exploit compression, the L2 cache must be able to pack more compressed cache lines than uncompressed lines into the same space. One approach is to decouple the cache access, adding a level of indirection between the address

Table 1. Frequent Pattern Encoding

Prefix	Pattern Encoded	Data Size
000	Zero Run	3 bits (for runs up to 8 zeros)
001	4-bit sign-extended	4 bits
010	One byte sign-extended	8 bits
011	halfword sign-extended	16 bits
100	halfword padded with a zero halfword	The nonzero halfword (16 bits)
101	Two halfwords, each a byte sign-extended	The two bytes (16 bits)
110	word consisting of repeated bytes	8 bits
111	Uncompressed word	Original Word (32 bits)

tag and the data storage. Seznec's decoupled sector cache does this on a per-set basis to improve the utilization of sector (or sub-block) caches [23]. Hallnor and Reinhardt's Indirect-Index Cache (IIC) does this across the whole cache, allowing fully-associative placement and a software managed replacement policy [15]. Lee, et al.'s selective compressed caches use this technique to allow two compressed cache blocks to occupy the space required for one uncompressed block [21, 19, 20]. Decoupled access is simpler if we serially access the cache tags before the data. Fortunately, this is increasingly necessary to limit power dissipation [17].

In theory, a cache line can be compressed into any number of bits. This can be achieved in a completely decoupled design across the whole cache (e.g., IIC). However, such design adds more complexity to cache management. In our compressed cache design, the decoupled variable segment cache [5], each cache line is stored as a group of one or more 8-byte *segments*. For example, a 64-byte line can be stored in 1-8 segments. A compressed line is padded with zeros till its size becomes a multiple of the segment size, and these extra zeros (that do not correspond to any tags) are ignored during decompression. While this approach doesn't permit high compression ratios for some cache lines (e.g., all zero lines), it allows for a more practical and faster implementation of cache accesses.

4 Compression and Decompression

We propose a compressed cache design in which data is stored uncompressed in the level-1 caches and compressed in the level-2 caches [5]. This helps reduce many of the costly L2 cache misses that hinder performance, while not affecting the common case of an L1 hit. However, such a design adds the overhead of compressing or decompressing cache lines when moved between the two levels. FPC allows a relatively fast implementations of both of these functions.

Compression. Cache line compression occurs when data is written back from the L1 to the L2 cache. A cache line is compressed easily using a simple circuit that checks each word (in parallel) for pattern matches. If a word matches any of the seven compressible patterns, a simple encoder circuit is used to encode the word into its most compact form. If no match was found, the whole word is stored with the prefix '111'. This can be performed in one cycle, assuming 12 FO4 delays. For zero runs, we need to detect such runs of consecutive zeros, and increment the data value of the first occurrence to represent their count. Since zero runs are limited in our design to eight zeros, this can be implemented in a single cycle using a simple multiplexer/adder circuit.

Cache line compression can be implemented in a memory pipeline, by allocating three pipeline stages on the L1-to-L2 write path (one for pattern matching, one for zero run encoding, and one for gathering the compressed line). A small victim cache that contains a few entries in both compressed and uncompressed form can be used to hide the compression latency on L1 writebacks.

Decompression. Cache line decompression occurs when data is read from the L2 to the L1 caches. This is a frequent event for most benchmarks whose working sets do not fit in the L1 cache. Decompression latency is critical since it is directly added to the L2 hit latency. Decompression is a slower process than compression, since prefixes for all words in the line have to be accessed in series, because each prefix is used to determine the length of its corresponding encoded word and therefore the starting location of all the subsequent compressed words. Figure 1 presents a sche-

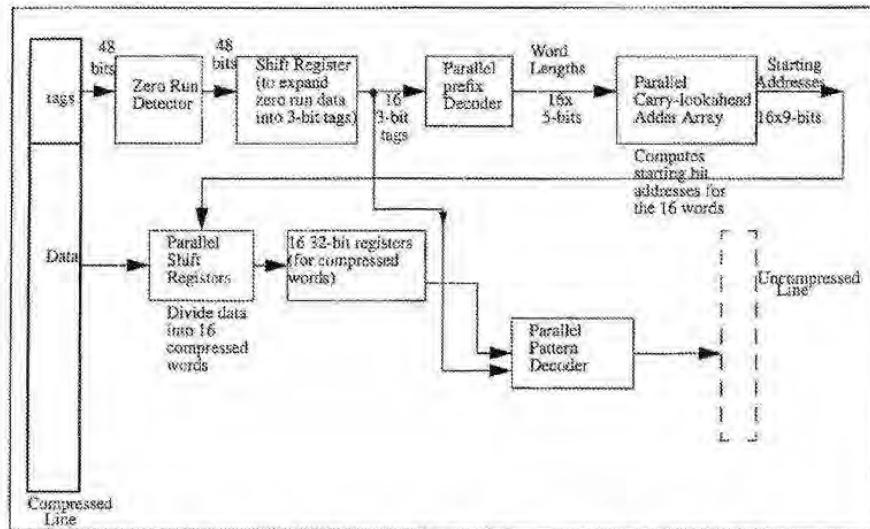


Figure 1. Cache line decompression pipeline for a 64-byte (16-word) cache line.

This is a five-stage pipeline used to decompress a compressed cache line, where each stage contains 12 FO4 gate delays or less. The first pipeline stage (containing the zero run detector, shift register and parallel prefix decoder) decodes the prefix array to determine the length in bits of each word. The second and third stages (Parallel Prefix adder array) compute the starting bit address for each data word by adding the length fields of the preceding words in a hierarchical fashion. The fourth stage (parallel shift registers) contains 16 registers each of which is shifted by the starting address of its word. The fifth and last stage contains the pattern decoder, which decodes the content of each 32-bit register into an uncompressed word according to its corresponding prefix.

matic diagram for a five-stage hardware pipeline that can be used to decompress 64-byte cache lines. Each pipeline stage is 12 FO4 delays or less, assuming the parallel resources required are available for the parallel adder, shift register and pattern decoder. Assuming one processor cycle requires 12 FO4 gate delays, this means that the decompression latency is limited to 5 processor cycles.

5 Evaluation

We evaluate our FPC scheme in terms of its achieved compressibility compared to other compression schemes. We show compression results for our frequent patterns, and demonstrate that zero runs are the most frequent. We also analyze the performance of segmented compression, and the effect of restricting compressed lines to segment boundaries on compression ratios.

5.1 Workloads

To evaluate our design against alternative schemes, we used several multi-threaded commercial workloads from the Wisconsin Commercial Workload Suite [3]. We also used six of the SPEC [24] benchmarks, three from the integer suite (SPECint2000) and three from the floating point suite (SPECfp2000). All of these workloads run under the




Figure 2. Compression ratios (original size / compressed size) for XRL, BRCL, FPC and gzip




Figure 3. Compression ratios for segmented and variable-length FPC

In most practical cache designs, cache lines cannot occupy any arbitrary number of bits. Restricting the compressed line sizes to a certain subset of all possible lengths (as we do in our segmented design) partially reduces compressibility. To assess the loss in compressibility, we compare the compression ratio from our Segmented Frequent Pattern Compression scheme (Segmented-FPC) against the compression ratio from the Frequent Pattern Compression scheme assuming variable-length lines are possible (Maximum-FPC).

Figure 3 shows the compression ratios from the two schemes for our ten benchmarks. The simple scheme (Segmented-FPC) has compression ratios of 1.7-2.4 for the three SPECint2000 benchmarks, 1.0-1.3 for the three SPECfp2000 benchmarks, 1.4-2.1 for the four commercial benchmarks. OLTP had the lowest compression ratio among our set of commercial benchmarks, since its data is randomly generated. A real OLTP application would have much less randomness, and thus have a higher compression ratio.




Figure 4. Frequent Pattern Histogram

Compression ratios are higher for all integer benchmarks (1.4-2.4) than floating point benchmarks. For example, only 0.4% of all cache lines in *swim* are compressible. This is because of the nature of floating point benchmarks where floating point numbers might not fit any of the frequent patterns. However, some benefit is still possible for benchmarks with lots of zero words. Segmented-FPC achieved most of the compression benefit from variable-length lines of Maximum-FPC.

5.3 Which Patterns Are Frequent?

Frequent Pattern Compression (FPC) is built on the observation that some word patterns are more frequent than others. We experimented with cache snapshots for our different benchmarks to come up with a reasonable set of frequent patterns (described in Table 1). Figure 4 shows the relative frequency of incompressible words, zero words and words compressible to 4, 8 and 16 bits. The 4-, 8-, and 16-bit patterns are present with various frequencies across our integer and commercial benchmarks. Unfortunately, most of the words in floating point benchmarks are incompressible with FPC, since our patterns are mainly integer patterns.

As Figure 4 demonstrates, zero words are the most frequent compressible pattern across all benchmarks, which is why some compression techniques (e.g., X-RI) specifically optimize for runs of zeros. Figure 5 shows the average number of zeros in a zero run for our set of benchmarks. Except for *equake* and *jbb*, the average zero run length for all benchmarks is greater than two. In developing the FPC schema, we had two options to compress zeros. The first was to have a prefix for each zero word with no data. The second was to encode zero runs with a single prefix and save the length of the run in the data part corresponding to that prefix. However, since most zero runs have more than two (and in most cases three) words, the additional compressibility justifies having special treatment for zero runs.




Figure 5. Average number of words in a zero run for our ten benchmarks




Figure 6. FPC Compression ratios for segment sizes (1 bit to 32 bytes)

5.4 Analysis of Segmented Frequent Pattern Compression

In designing a practical compressed cache implementation, selecting a specific base segment size is critical. A compressed line can only be stored in a size that is an integer multiple of the base segment size. Smaller segments allow for higher compression ratios. On the other hand, larger segments decrease the cache design complexity. Cache design should balance the tradeoff between these two conflicting issues. We selected a base segment size of 8-bytes (i.e., up to 8 segments for 64-byte lines) in our Segmented FPC design.

Figure 6 shows the sensitivity of our compression schemes to the base segment size. The four bars for each benchmark represent compression ratios if we have two possible sizes, i.e., an uncompressed line occupying two segments (32-byte segments), four (16-byte segments), eight (8-byte segments, which is the same as Segmented-FPC in




Figure 7. Segment Length Histogram: Percentage of cache lines compressed into 1-8 Segments

Figure 3), and all possible sizes (1-bit segments, the same as Maximum-FPC). Our 8-byte-segment design increases the compression ratio by up to 52% vs. 32-byte segments, and up to 19% vs. 16-byte segments. Figure 7 shows the percentage of lines that can be compressed into 1-8 segments. We show a more detailed distribution in Figure 8, demonstrating the cumulative distribution of compressed cache line sizes (1-512 bits) for our ten benchmarks, as well as the 25th, 50th and 75th percentiles.

6 Conclusion

Cache designers might consider using cache compression to increase cache capacity and reduce off-chip bandwidth. In this document, we propose and evaluate a simple significance-based compression scheme suitable for cache lines, since it has a low compression and decompression overhead. This scheme, Frequent Pattern Compression (FPC) compresses individual cache lines on a word-by-word basis by storing common word patterns in a compressed format accompanied with an appropriate prefix. This simple scheme provides comparable compression ratios to more complex schemes that have higher cache hit latencies.

References

- [1] Bulent Abali, Hubertus Franke, Xiaowei Shen, Dan E. Poff, and T. Basil Smith. Performance of hardware compressed main memory. In *Proceedings of the Seventh IEEE Symposium on High-Performance Computer Architecture*, pages 73–81, January 2003.
- [2] Edward Ahn, Seung-Moon Yoo, and Sung-Mo Steve Kang. Effective Algorithms for Cache-level Compression. In *Proceedings of the 2001 Conference on Great Lakes Symposium on VLSI*, pages 89–92, 2001.
- [3] Alaa R. Alameldeen, Milo M. K. Martin, Carl J. Mayer, Kevin E. Moore, Min Xu, Daniel J. Sorin, Mark D. Hill, and David A. Wood. Simulating a 52M Commercial Server on a 33K PC. *IEEE Computer*, 36(2):50–57, February 2003.
- [4] Alaa R. Alameldeen and David A. Wood. Variability in Architectural Simulations of Multi-threaded Workloads. In *Proceedings of the Ninth IEEE Symposium on High-Performance Computer Architecture*, pages 7–18, February 2003.
- [5] Alaa R. Alameldeen and David A. Wood. Adaptive Cache Compression for High-Performance Processors. In *Proceedings of the 31st Annual International Symposium on Computer Architecture*, June 2004.

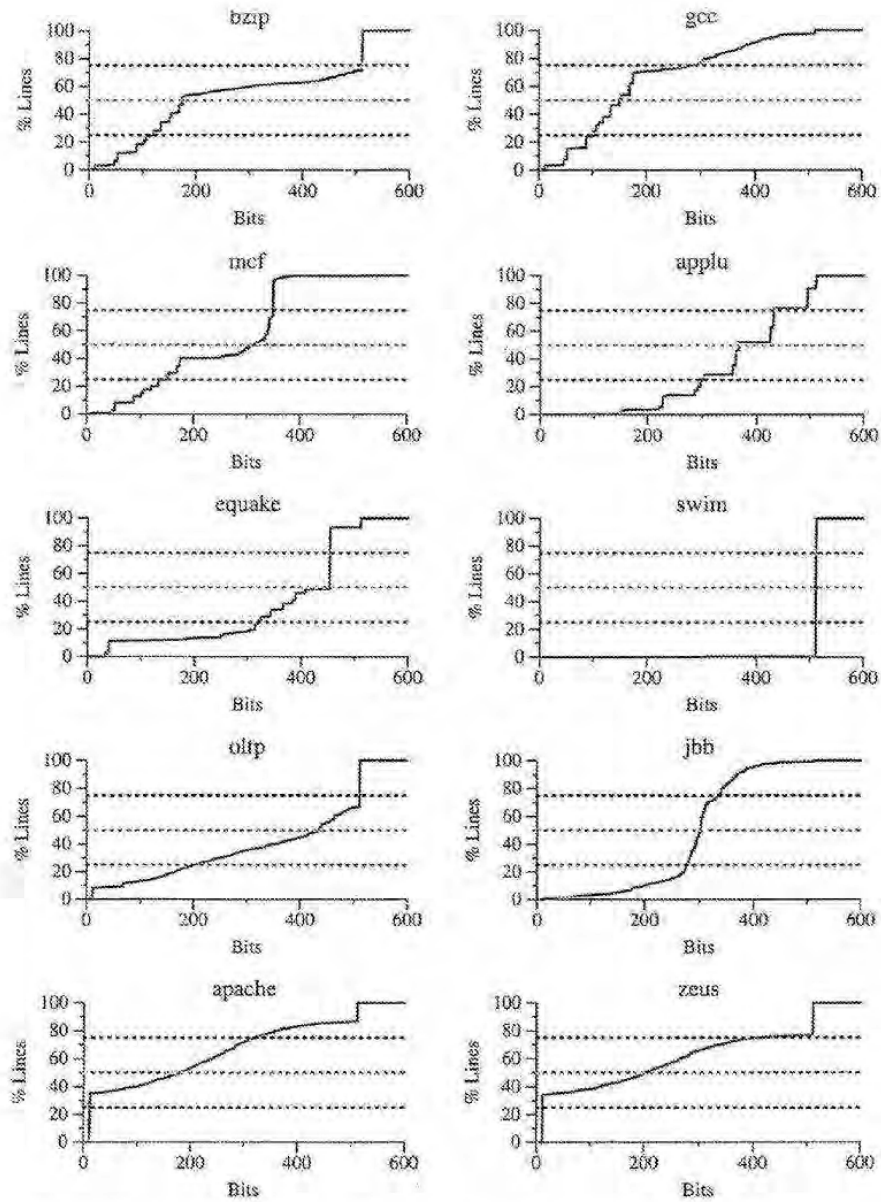


Figure 8. Cumulative Distribution of Compressed Line Lengths (1 to 512 bits). These graphs highlight the 25th, 50th and 75th percentile values.

- [6] Paul Barford and Mark Crovella. Generating Representative Web Workloads for Network and Server Performance Evaluation. In *Proceedings of the 1998 ACM Sigmetrics Conference on Measurement and Modeling of Computer Systems*, pages 151–160, June 1998.
- [7] Luca Benini, Davide Bruni, Alberto Macii, and Enrico Macii. Hardware-Assisted Data Compression for Energy Minimization in Systems with Embedded Processors. In *Proceedings of the IEEE 2002 Design Automation and Test in Europe*, pages 449–453, 2002.
- [8] Luca Benini, Davide Bruni, Bruno Ricco, Alberto Macii, and Enrico Macii. An Adaptive Data Compression Scheme for Memory Traffic Minimization in Processor-Based Systems. In *Proceedings of the IEEE International Conference on Circuits and Systems, ICCAS-02*, pages 866–869, May 2002.
- [9] Ramon Canal, Antonio Gonzalez, and James E. Smith. Very Low Power Pipelines Using Significance Compression. In *Proceedings of the 33rd Annual IEEE/ACM International Symposium on Microarchitecture*, pages 181–190, December 2000.
- [10] David Chen, Enoch Peserico, and Larry Rudolph. A Dynamically Partitionable Compressed Cache. In *Proceedings of the Singapore-MIT Alliance Symposium*, January 2003.
- [11] Daniel Citron and Larry Rudolph. Creating a Wider Bus Using Caching Techniques. In *Proceedings of the First IEEE Symposium on High-Performance Computer Architecture*, pages 90–99, February 1995.
- [12] Matthew Farrens and Arvin Park. Dynamic Base Register Caching: A Technique for Reducing Address Bus Width. In *Proceedings of the 18th Annual International Symposium on Computer Architecture*, pages 128–137, May 1991.
- [13] P.A. Franaszek and J.T. Robinson. On Internal Organization in Compressed Random-Access Memories. *IBM Journal of Research and Development*, 45(2):259–270, March 2001.
- [14] Peter Franaszek, John Robinson, and Joy Thomas. Parallel Compression with Cooperative Dictionary Construction. In *Proceedings of the Data Compression Conference, DCC'96*, pages 200–209, March 1996.
- [15] Erik G. Hallnor and Steven K. Reinhardt. A Fully Associative Software-Managed Cache Design. In *Proceedings of the 27th Annual International Symposium on Computer Architecture*, pages 107–116, June 2000.
- [16] Krishna Kant and Ravi Iyer. Compressibility Characteristics of Address/Data transfers in Commercial Workloads. In *Proceedings of the Fifth Workshop on Computer Architecture Evaluation Using Commercial Workloads*, pages 59–67, February 2002.
- [17] R. E. Kessler. The Alpha 21264 Microprocessor. *IEEE Micro*, 19(2):24–36, March/April 1999.
- [18] Morten Kjølseto, Mark Gooch, and Simon Jones. Design and Performance of a Main Memory Hardware Data Compressor. In *Proceedings of the 22nd EUROMICRO Conference*, 1996.
- [19] Jang-Soo Lee, Won-Kee Hong, and Shin-Dug Kim. Design and Evaluation of a Selective Compressed Memory System. In *Proceedings of International Conference on Computer Design (ICCD'99)*, pages 184–191, October 1999.
- [20] Jang-Soo Lee, Won-Kee Hong, and Shin-Dug Kim. An On-chip Cache Compression Technique to Reduce Decompression Overhead and Design Complexity. *Journal of Systems Architecture: the EUROMICRO Journal*, 46(15):1365–1382, December 2000.
- [21] Jang-Soo Lee, Won-Kee Hong, and Shin-Dug Kim. Adaptive Methods to Minimize Decompression Overhead for Compressed On-chip Cache. *International Journal of Computers and Application*, 25(2), January 2003.
- [22] J. Pomerene, T. Puzak, R. Rechtschaffen, and P. Sparacio. Prefetching System for a Cache Having a Second Directory for Sequentially Accessed Blocks, February 1989. U.S. Patent 4,807,110.
- [23] André Seznot. Decoupled Sectored Caches. *IEEE Transactions on Computers*, 46(2):210–215, February 1997.
- [24] Systems Performance Evaluation Cooperation. SPEC Benchmarks. <http://www.spec.org>.
- [25] R. Brett Tremaine, T. Basil Smith, Mike Wazlowski, David Har, Kwok-Ken Mak, and Sujith Aravamreddy. Pin-nacle: IBM MXT in a Memory Controller Chip. *IEEE Micro*, 21(2):56–68, March/April 2001.
- [26] R.B. Tremaine, P.A. Franaszek, J.T. Robinson, C.O. Schulz, T.B. Smith, M.E. Wazlowski, and P.M. Bland. IBM Memory Expansion Technology (MXT). *IBM Journal of Research and Development*, 45(2):271–285, March 2001.
- [27] Jun Yang and Rajiv Gupta. Energy Efficient Frequent Value Data Cache Design. In *Proceedings of the 35th Annual IEEE/ACM International Symposium on Microarchitecture*, pages 197–207, November 2002.
- [28] Jun Yang and Rajiv Gupta. Frequent Value Locality and its Applications. *ACM Transactions on Embedded Computing Systems*, 1(1):79–105, November 2002.
- [29] Jun Yang, Youtao Zhang, and Rajiv Gupta. Frequent Value Compression in Data Caches. In *Proceedings of the 33rd Annual IEEE/ACM International Symposium on Microarchitecture*, pages 258–265, December 2000.

- [30] Youtao Zhang and Rajiv Gupta. Data Compression Transformations for Dynamically Allocated Data Structures. In *Proceedings of the International Conference on Compiler Construction (CC)*, pages 24–28, April 2002.
- [31] Youtao Zhang, Jun Yang, and Rajiv Gupta. Frequent Value Locality and Value-centric Data Cache Design. In *Proceedings of the Ninth International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 150–159, November 2000.
- [32] Jacob Ziv and Abraham Lempel. A Universal Algorithm for Sequential Data Compression. *IEEE Transactions on Information Theory*, 23(3):337–343, May 1977.

Electronic Acknowledgement Receipt	
EFS ID:	26525118
Application Number:	15105648
International Application Number:	
Confirmation Number:	5548
Title of Invention:	Codebook Subset Restriction Signaling
First Named Inventor/Applicant Name:	Sebastian FaxÃ©r
Customer Number:	24112
Filer:	Justin J. Leonard/Katya Fox
Filer Authorized By:	Justin J. Leonard
Attorney Docket Number:	4015-9595 / P45698-US2
Receipt Date:	02-AUG-2016
Filing Date:	
Time Stamp:	16:40:55
Application Type:	U.S. National Stage under 35 USC 371

Payment information:

Submitted with Payment	no
------------------------	----

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Transmittal Letter	Cover_letter.pdf	26493 /s/Justin Leonard 15105648 20160802 164055	no	1

Warnings:

Information:					
2	Information Disclosure Statement (IDS) Form (SB08)	IDS1.pdf	1035445 06700b13e1d66f423a7f8d11fd755c2a561d8fe5	no	4
Warnings:					
Information:					
3	Non Patent Literature	THOMAS.pdf	5918706 05cf7d5d6490a1125692a29db6c0ba4ee91a0c3	no	6
Warnings:					
Information:					
4	Non Patent Literature	ALAMELDEEN.pdf	18304462 115665131b6a0c0f880e8d1d46f8577b15643281f7	no	15
Warnings:					
Information:					
Total Files Size (in bytes):				25285106	
<p>This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.</p> <p><u>New Applications Under 35 U.S.C. 111</u> If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.</p> <p><u>National Stage of an International Application under 35 U.S.C. 371</u> If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.</p> <p><u>New International Application Filed with the USPTO as a Receiving Office</u> If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.</p>					